

The Springer Series on Demographic Methods
and Population Analysis 41

André Grow
Jan Van Bavel *Editors*

Agent-Based Modelling in Population Studies

Concepts, Methods, and Applications

 Springer

The Springer Series on Demographic Methods and Population Analysis

Volume 41

Series Editor

Kenneth C. Land, Duke University

In recent decades, there has been a rapid development of demographic models and methods and an explosive growth in the range of applications of population analysis. This series seeks to provide a publication outlet both for high-quality textual and expository books on modern techniques of demographic analysis and for works that present exemplary applications of such techniques to various aspects of population analysis.

Topics appropriate for the series include:

- General demographic methods
- Techniques of standardization
- Life table models and methods
- Multistate and multiregional life tables, analyses and projections
- Demographic aspects of biostatistics and epidemiology
- Stable population theory and its extensions
- Methods of indirect estimation
- Stochastic population models
- Event history analysis, duration analysis, and hazard regression models
- Demographic projection methods and population forecasts
- Techniques of applied demographic analysis, regional and local population estimates and projections
- Methods of estimation and projection for business and health care applications
- Methods and estimates for unique populations such as schools and students

Volumes in the series are of interest to researchers, professionals, and students in demography, sociology, economics, statistics, geography and regional science, public health and health care management, epidemiology, biostatistics, actuarial science, business, and related fields.

More information about this series at <http://www.springer.com/series/6449>

André Grow • Jan Van Bavel
Editors

Agent-Based Modelling in Population Studies

Concepts, Methods, and Applications

 Springer

Editors

André Grow
Centre for Sociological Research
University of Leuven (KU Leuven)
Leuven, Belgium

Jan Van Bavel
Centre for Sociological Research
University of Leuven (KU Leuven)
Leuven, Belgium

ISSN 1389-6784

ISSN 2215-1990 (electronic)

The Springer Series on Demographic Methods and Population Analysis

ISBN 978-3-319-32281-0

ISBN 978-3-319-32283-4 (eBook)

DOI 10.1007/978-3-319-32283-4

Library of Congress Control Number: 2016942784

© Springer International Publishing Switzerland 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG Switzerland

Preface

This book contains a selection of papers that have been presented at the workshop ‘Recent Developments and Future Directions in Agent-Based Modelling in Population Studies’ that we organized at the University of Leuven (KU Leuven), Belgium, 18–19 September 2014. All papers have been revised after the workshop and accepted after having been peer-reviewed. The workshop was organized in the context of the project ‘Implications of the Shifting Gender Balance in Education for Reproductive Behaviour in Europe’ and received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement no. 312290 for the GENDER-BALL project, the Concerted Research Action ‘New Approaches to the Social Dynamics of Long Term Fertility Change’ (KU Leuven grant), and the Scientific Research Group Historical Demography (Research Foundation Flanders). We are grateful for the support from the financing institutions. We would also like to thank Francesco Billari for his inspiring concluding talk; the other members of the scientific committee, Koenraad Matthys, Geert Molenberghs, Giovanni Samaey, and Geert Verbeke, for their contributions to the workshop; and Martine Parton and Marina Franckx for their help with the organization. Finally, we would like to thank the reviewers for their help with assessing the contributions to this book.

Leuven, Belgium
2016

André Grow
Jan Van Bavel

Contents

Part I Perspectives on Agent-Based Modelling in Population Studies

- 1 **Introduction: Agent-Based Modelling as a Tool to Advance Evolutionary Population Theory** 3
Jan Van Bavel and André Grow
- 2 **Model-Based Demography: Towards a Research Agenda** 29
Daniel Courgeau, Jakub Bijak, Robert Franck, and Eric Silverman
- 3 **The Role of Social Interactions in Demography: An Agent-Based Modelling Approach** 53
Alexia Prskawetz

Part II Designing, Analysing, and Reporting Agent-Based Models

- 4 **Agent-Based Computational Demography and Microsimulation Using JAS-mine** 75
Matteo Richiardi and Ross E. Richardson
- 5 **Simulating Synthetic Life Courses of Individuals and Couples, and Mate Matching** 113
Sabine Zinn
- 6 **Using Survey Data for Agent-Based Modeling: Design and Challenges in a Model of Armed Conflict and Population Change** 159
Nathalie E. Williams, Michelle L. O'Brien, and Xiaozheng Yao

7 Regression Metamodels for Sensitivity Analysis in Agent-Based Computational Demography 185
 André Grow

8 Design and Analysis of Demographic Simulations 211
 Jason Hilton and Jakub Bijak

9 How to Describe Agent-Based Models in Population Studies? 237
 Jürgen Groeneveld, Anna Klabunde, Michelle L. O’Brien,
 and André Grow

Part III Modelling Decision Processes

10 The Decision to Emigrate: A Simulation Model Based on the Theory of Planned Behaviour 257
 Frans Willekens

11 Deciding to Disclose: A Decision Theoretic Agent Model of Pregnancy and Alcohol Misuse 301
 Jonathan Gray, Jakub Bijak, and Seth Bullock

Part IV Family Formation and Fertility

12 An Agent-Based Model of Sex Ratio at Birth Distortions 343
 Ridhi Kashyap and Francisco Villavicencio

13 Exploring the Role of Communication in Shaping Fertility Transition Patterns in Space and Time 369
 Sebastian Klüsener, Francesco Scalone, and Martin Dribe

14 Feedback Mechanisms in the Postponement of Fertility in Spain 405
 Daniel Ciganda and Francisco Villavicencio

Part V Health, Mortality, and Support in Old Age

15 Linking Income, Transfers, and Social Support in an Agent-Based Family Exchange Model 439
 Fanny A. Kluge and Tobias C. Vogt

16 Agent-Based Modelling to Inform Health Intervention Strategies: The Case of Severe Acute Malnutrition in Children in High-Burden Low-Income Countries 461
 Hedwig Deconinck, Carine Van Malderen, Niko Speybroeck,
 Jean Macq, and Jean-Christophe Chiem

17 Exploring Contingent Inequalities: Building the Theoretical Health Inequality Model 487
 Michael Wolfson, Steve Gribble, and Reed Beall

Contributors

Reed Beall University of Ottawa, Ottawa, ON, Canada

Jakub Bijak Department of Social Statistics and Demography, University of Southampton, Southampton, UK

Seth Bullock Department of Computer Science, University of Bristol, UK

Jean-Christophe Chiem Institut de recherche santé et société, Université catholique de Louvain, Brussels, Belgium

Daniel Ciganda Department of Political and Social Sciences, Pompeu Fabra University, Barcelona, Spain

Daniel Courgeau Institut national d'études démographiques, Paris, France

Hedwig Deconinck Institut de recherche santé et société, Université catholique de Louvain, Brussels, Belgium

Martin Dribe Centre for Economic Demography, Lund University, Lund, Sweden

Robert Franck Université catholique de Louvain, Louvain-la-Neuve, Belgium

Jonathan Gray Department of Social Statistics and Demography, University of Southampton, Southampton, UK

Steve Gribble University of Ottawa, Ottawa, ON, Canada

Jürgen Groeneveld Department of Ecological Modelling, Helmholtz Centre for Environmental Research – UFZ, Leipzig, Germany

André Grow Centre for Sociological Research, University of Leuven (KU Leuven), Leuven, Belgium

Jason Hilton University of Southampton, Southampton, UK

Ridhi Kashyap Department of Sociology and Nuffield College, University of Oxford, Oxford, UK

Max Planck Institute for Demographic Research, Rostock, Germany

Anna Klabunde Max Planck Institute for Demographic Research, Rostock, Germany

Sebastian Klüsener Max Planck Institute for Demographic Research, Rostock, Germany

Fanny A. Kluge Max Planck Institute for Demographic Research, Rostock, Germany

Jean Macq Institut de recherche santé et société, Université catholique de Louvain, Brussels, Belgium

Michelle L. O'Brien Department of Sociology, University of Washington, Seattle, WA, USA

Center for Studies in Demography and Ecology, University of Washington, Seattle, WA, USA

Alexia Prskawetz Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria

Wittgenstein Centre for Demography and Global Human Capital (IIASA, VID/ÖAW, WIC), Vienna, Austria

Ross E. Richardson Institute for New Economic Thinking at the Oxford Martin School, Mathematical Institute, University of Oxford, Oxford, UK

Matteo Richiardi Institute for New Economic Thinking at the Oxford Martin School, Mathematical Institute, University of Oxford, Oxford, UK

Nuffield College, Oxford, UK

Collegio Carlo Alberto, Moncalieri, Italy

Department of Economics and Statistics, University of Torino, Turin, Italy

Francesco Scalone Department of Statistical Sciences, University of Bologna, Bologna, Italy

Eric Silverman School of Computing, Teesside University, Middlesbrough, UK

Niko Speybroeck Institut de recherche santé et société, Université catholique de Louvain, Brussels, Belgium

Jan Van Bavel Centre for Sociological Research, University of Leuven (KU Leuven), Leuven, Belgium

Carine Van Malderen Institut de recherche santé et société, Université catholique de Louvain, Brussels, Belgium

Francisco Villavicencio Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark

Max-Planck Odense Center on the Biodemography of Aging, Odense, Denmark

Tobias C. Vogt Max Planck Institute for Demographic Research, Rostock, Germany

Frans Willekens Netherlands Interdisciplinary Demographic Institute (NIDI), The Hague, The Netherlands

Nathalie E. Williams Department of Sociology, University of Washington, Seattle, WA, USA

Jackson School of International Studies, University of Washington, Seattle, WA, USA

Center for Studies in Demography and Ecology, University of Washington, Seattle, WA, USA

Michael Wolfson University of Ottawa, Ottawa, ON, Canada

Xiaozheng Yao Virginia Tech University, Blacksburg, VA, USA

Sabine Zinn National Educational Panel Study, Leibniz Institute for Educational Trajectories, Bamberg, Germany

Max Planck Institute for Demographic Research, Rostock, Germany

Part I
Perspectives on Agent-Based Modelling
in Population Studies

Chapter 1

Introduction: Agent-Based Modelling as a Tool to Advance Evolutionary Population Theory

Jan Van Bavel and André Grow

1.1 Introduction

Demography has for long and repeatedly been described as a field that is rich in methods but poor in theories (Burch 2003a; De Bruijn 1999; Tabutin 2007; Vance 1952). While there has been a lot of methodological advancement, the field has made less progress in generating widely accepted theories that explain trends in fertility, mortality, migration, or other aspects of population. Of course, to the extent that the dynamics of human populations are governed by the same kind of forces as other social processes, demographers can and do borrow theories from other social sciences. However, to the extent that important aspects of population processes really are a reality *sui generis*, the field would strongly benefit from more theory development.

More than 10 years ago, Billari et al. (2003) recommended agent-based modelling (ABM) as a tool to advance population theory. While a number of ABM-contributions have been published in the mainstream demographic journals since then, ABM still has not become a standard tool in every demographer's kit and the advancement of population theory through ABM still remains limited. Ironically, Billari et al. (2003, p. 3) already pointed out an important factor hindering the widespread application of ABM in population studies: the lack of theories. ABM proceeds by implementing theoretical rules of behaviour, decision-making, and interaction in a simulation and then investigates the resulting patterns that emerge from this. So, on the one hand, in order to apply ABM, one needs theory; on the other hand, we want to apply ABM in order to develop the theories we are lacking.

J. Van Bavel (✉) • A. Grow
Centre for Sociological Research, University of Leuven (KU Leuven), Parkstraat 45,
3000 Leuven, Belgium
e-mail: jan.vanbavel@soc.kuleuven.be; andre.grow@soc.kuleuven.be

In this introduction, we want to highlight one reason for this ‘Catch-22’, namely the rather ‘closed’ concept of population that has been dominating the field. Most of the effort in demography has been devoted so far to the numerical monitoring of national population flows and structures. Not only has demography been preoccupied with empirical data and techniques for analysing these data (Burch 2003b), the field has also been preoccupied with data representing the populations of nation states. The requirement to have ‘nationally representative’ data has important advantages, but it has hindered creative theory development. The field should adopt a more ‘open’ approach to population to allow more flexible experimentation with theories. We will argue how ABM offers a tool to help bridging the gap between different approaches to the concept of population. Next, we illustrate some of the arguments with examples from chapters in this volume. Subsequently, we argue that evolutionary theory might be a particularly suitable theoretical framework for developing population theory aided by ABM.

1.2 Two Concepts of Population

In the course of its development over the nineteenth and twentieth century, an approach to population has come to dominate demography linked with a ‘closed’ concept of population. This approach prioritizes the descriptive coverage of nationally representative population indicators rather than understanding the underlying heterogeneity and processes (Kreager 2009, 2015a, b). Central to classic demography has been the accurate bookkeeping of humans in national populations. In the national demographic accounts, births and deaths represent the natural sources of population flow. Migration is considered from a national point of view as well, namely as outmigration from one country and immigration into another country. The basic demographic equation describes how both natural and migration flows affect the size and age structure of the population, and cohort component methods can be used to project it into the future. Getting the rates right is central to accurate national bookkeeping, depending on correctly counting the number of demographic events to put in the numerator and enumerating the relevant population in the denominator. The seminal work by Lotka and later developments in formal demography exemplify this ‘closed’ concept of population (Dublin and Lotka 1925; Schoen 2006).

This approach was closely connected with the rise of the idea of the nation state, where nations are defined by a delimited population, sharing territory, language, and historical experience (Kreager 2009, 2015a). It has been a powerful ally for the establishment (and national funding) of human demography as a field. The concept of the national population (and their smaller and larger scale derivatives) has stimulated demographers to develop ingenious methods to measure fertility and reproduction, mortality, and migration. It has inspired debates over things such as replacement level fertility, about whether the increase of the TFR from 1.5 to 1.6 in some country represents a quantum or a tempo shift, or about the impact

of immigration on the structure of the population. Such measures and debates have enhanced our insight into important issues like population growth and decline, the relative role played by quantum and tempo shifts in demographic trends, or population ageing. This is great progress and looked at it in a specific way, it could be called theoretical progress too (see Burch 2003a). Many key insights from demography are important for the management of nation states and their institutions, as the long-standing debate about below-replacement fertility and its relation to population ageing illustrates (Van Bavel 2010a). Accordingly, demography has become an important field “in service of the state” (Kreager 2015b, p. S32).

The fact that the national population has become the dominant point of reference does not imply that demographers have failed to investigate variation within countries (cf. Billari 2015; Courgeau et al. 2016). Notably towards the end of the twentieth century, demographers have increasingly adopted regression analysis as a tool to investigate how fertility, mortality, and migration co-vary with things such as education, wealth, or religion. Courgeau et al. (2016) discuss more in depth the advances made in demography, from studying national aggregates over individual level modelling towards multilevel event history analysis, and they argue that these advances may even be considered as paradigm shifts. Still, the national population remained the standard point of reference, with analyses being carried out preferably based on nationally representative samples, and comparative studies being carried out between nation states.

While the ‘closed’ concept of the national population has been very instrumental in the establishment of the discipline, the rather rigid approach may have hindered the creative development of population theories. Methodologically, the dominance of this nineteenth century concept of population is reflected in the heavy reliance of demographic studies on either (single- or multi-country) census or nationally representative survey data – to such an extent that sound studies of demographic processes might be rejected due to a lack of ‘representative data’. Similarly, theoretical work tends to be accepted as a serious scientific contribution only if its relevance could be shown, empirically, on a census or nationally representative sample (cf. Billari 2015). This is an extremely costly and inflexible requirement, discouraging creative experimentation with new ideas. It limits the room for more particularistic reasoning about how local conditions differentially affect certain groups and their relations with others (Kreager 2015a, p. 73, 2015b).

Kreager (2009, 2015a, b) has shown how the concept of the enumerated, national population as the standard point of reference got established at the expense of an alternative concept of population. In the alternative ‘open’ approach to population, the emphasis is not so much on enumerating all individuals who belong to the country, but rather on the processes and structures that emerge out of the interactions between heterogeneous individuals and their environments, embedded in social groups and networks. The main concern in the ‘open’ study of populations is understanding the processes and mechanisms that generate patterns of association between individuals, such as mating or social networks, and how these processes affect population change and heterogeneity. This alternative approach largely got lost in most of the mainstream work in human demography but it remained very

strong in population biology. The detailed observation of particular species in their specific habitat by Charles Darwin exemplifies the alternative, more ‘open’ concept of population, and this has remained the dominant population concept in the Modern Synthesis in biology (Mayr 1991, 2002). While the emphasis in demography has been on averaging demographic behaviour in rates (typically for national populations and subpopulations) and calculating their long-term stable population implications, the emphasis in population biology has been on heterogeneity and change. In the words of Ernst Mayr: “The populationist stresses the uniqueness of everything in the organic world. What is true for the human species – that no two individuals are alike – is equally true for all other species of animals and plants. [...] [F]or the populationist the type (average) is an abstraction and only the variation is real” (Mayr 1959 cited in Mayr 2002, p. 92).

In order to understand the past, present, and future dynamics of populations as networks of interactions, it is insufficient to survey and analyse statistically cross-sectional snapshots of samples of individuals and their characteristics. Alternative and complementary modes of observation are needed, including the kind of local, small-scale observations to which Darwin devoted much of his life (Kreager 2009), or the kind of in-depth studies of local communities common in historical demography (e.g., Kertzer and Hogan 1989; Tsuya et al. 2010). A more ‘open’ approach to population may also integrate insights from experimental research, as a particular form of local, typically small-scale observations but with particular strengths when it comes to drawing conclusions about causality.

A move towards a concept of population as a fundamentally open and dynamic network of interacting individuals also calls for methods to study these dynamics in a flexible way. ABM is a useful tool to help opening up the ‘closed’ approach to population that has dominated the field. This, in turn, will help us to develop and refine our theories of population processes. More precisely, ABM may help us to bridge the ‘open’ and ‘closed’ concepts of population in a way that we may benefit from the advantages of both approaches while acknowledging their respective limitations.

1.3 How Agent-Based Modelling May Bridge the Two Approaches to Population

Demography studies populations of individuals who interact in complex ways in different layers of cultural and social environments. It often investigates emergent regularities of such individual-level contextualized behaviour. ABM lends itself quite naturally to deal with this complexity (Courgeau et al. 2016): ABM is population oriented and applying ABM starts with imagining a population of individual agents. Here, we want to highlight how ABM may bridge the two concepts of population that we have just outlined. It can do this while maintaining a view on both the micro (individual) and the macro (aggregate) level. In this way,

it may help to address a major challenge in the development of population theory: “[H]ow to combine theoretical principles that operate at the local level with concepts of global population” (Hammel and Howell 1987, p. 142)? In order to see how this can work, it is useful to draw on the ‘macro-micro-macro model’ that is at the centre of the social mechanism approach to social theory (Coleman 1986, 1990; Hedström and Swedberg 1998; Hedström 2005) and which recently has been introduced to demography (Billari 2015).

1.3.1 *The Macro-Micro-Macro Model and Agent-Based Modelling*

The ‘macro-micro-macro model’ shown in Fig. 1.1 builds on the tradition of methodological individualism, in which social phenomena are viewed as the results of the actions of the individuals that make up the social system under consideration. Accordingly, proponents of the model argue that sound social science explanation should refer to these individuals and include explicit references to the causes and consequences of their actions (Hedström and Swedberg 1998, p. 12). In the model, explanations proceed in three steps. In the first step, an explanation indicates how the characteristics of the macro level affect the conditions and constraints that individuals face (situational mechanisms); in the second step, it indicates the way in which individuals assimilate these constraints and conditions in their behaviour (action-formation mechanisms); in the third step, it indicates how the actions and interactions of a large number of individuals bring about macro-level outcomes and social change (transformational mechanisms).

Applying this model to demography, Billari (2015) highlighted that the last step is the most novel and most important, but also the most challenging. It is most novel, because the first two steps have featured in existing demographic research. For example, the notion that the individual is affected by the characteristics of the macro

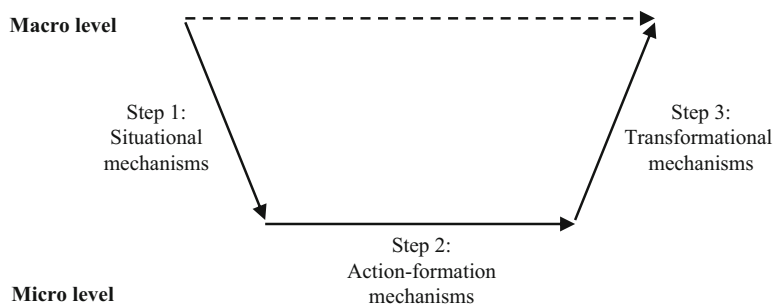


Fig. 1.1 The macro-micro-macro model in the social mechanism approach (adapted from Hedström and Swedberg 1998; Billari 2015)

level features prominently in the multilevel paradigm in demography; similarly, the notion that the constraints that individuals face affect their decisions and actions features in approaches such as life course analysis, in which the antecedents of people's behaviour lie in their own past (Courgeau et al. 2016). It is most important, because it reconnects the individual level with the macro level. It is most difficult, because the processes and dynamics by which individual interactions combine to generate macro-level outcomes can be very complex and this can make them very difficult to address with standard mathematical and statistical tools in demography.

For illustration, consider first the case of regression-based methods. The basic format for this approach is that an individual- or population-level outcome is "affected by" (Coleman 1986, p. 1328) a set of individual and contextual variables that are combined additively and linearly. Non-linearities in the relation between the different variables and the outcome are easily accommodated. Interaction effects between variables can be estimated too, in principle, but when several interactions between more than two variables are involved, then the model quickly becomes hard to handle. The goal of this approach is to find reliable statistical associations between the variables and the outcome, but it does not explicate the processes that underlie these associations. For example, a regression-based model can show that variation in the level of modernization is associated with variation in fertility rates across countries, but it does not explain how the two are connected through the actions of individuals. As noted above, multilevel models make it to some extent possible to model the way in which macro-level variables affect individuals, but even with this approach it is not possible to model the processes by which individual interactions feed back into the macro level. Williams et al. (2016) discuss additional problems that can arise from regression-based models. Similarly, econometric approaches, like instrumental variables and selection models, are geared toward isolating causal factors and assessing the extent to which the effect of variable X on variable Y testifies of 'true' causality. While this may sometimes be an important goal, also these models do not explicate the precise mechanism by which X and Y are connected.

Consider next mathematical models. An important question in demography is the role that social networks play for people's demographic decisions (Prskawetz 2016). The timing of entry into first marriage, for example, has been assumed to be affected by the number of peers who are already married. Hernes (1972) proposed a mathematical model that can show that the share of married individuals might indeed affect individuals' age-contingent probability to enter marriage. Yet, similar to regression models, this model does not explicate the processes and interactions by which individuals influence each other in their marital decisions (Hernes 1976). Even more, the model abstracts from social network structures that might exist in the population and that might affect the timing and spread of the diffusion process (cf. Cointet and Roth 2007).

For ABM, modelling transformational mechanisms and the interactions and network effects involved in this is at the core of the business. For example, as Klüsener et al. (2016) show, with ABM it is possible to implement socially and spatially segregated networks and this increases our ability to explain (spatial) diffusion

processes. This makes ABM a promising tool to facilitate theory development in population studies. A second important quality of ABM is its ability to show “the consequences of a few simple assumptions” (Axelrod 1997, p. 206). Given the computational power that is currently available on standard personal computers, and given the increasing access that many researchers have to grid computing, ABM makes it feasible to conduct simple as well as more complex thought experiments and to quantify the implications of different assumptions. ABM can therefore be used as a powerful computational laboratory to conduct simple as well as complex ‘what if’ thought experiments in a flexible but precise way. As pointed out by Prskawetz (2016) as well as Courgeau et al. (2016), with ABM, ‘toy models’ may therefore be employed to (pre-) test theories for which data are not easily obtained or not (yet) available at all. In this process researchers are not constrained to theorizing on the interplay between the individual and the population level alone. As the chapter by Wolfson et al. (2016) illustrates, ABM makes it possible to also consider all possible levels in between (e.g., schools, neighbourhoods, etc.). This flexible scalability (see Miller and Page 2007, pp. 85–86) enhances the opportunities for theory development, since theories can be developed at any level, while the implications of the theory for phenomena at other observational levels can then be computed.

This potential to conduct complex thought experiments does not mean that ABM should replace standard tools in demographic research and empirical data altogether. Quite to the contrary, we concur with Courgeau et al. (2016) that for making use of the full potential of ABM in demographic research, the connection with earlier modelling techniques is important. Such a connection enables to create empirically calibrated agent-models that have more realism and validity than purely theoretical simulation models.

1.3.2 Empirical Calibration

Calibrating agent-based models with empirical data is an important step in developing explanations of demographic change (Bijak et al. 2013; Courgeau et al. 2016; Hedström 2005). Such more advanced modes of computational experiments mixed with observational data can take the following basic form: (1) develop a theoretical model based on assumptions about individuals’ actions and interactions; (2) calibrate the model to match empirical data; and then, crucially, (3) conduct ‘what if’ experiments: what would happen if we leave out parameter X? How would things change if the empirical distributions would have been different than they are under actual conditions? And what if individuals behaved differently? In this way, after calibration, it is possible to perform counterfactual simulations that help advocating between possible alternative processes.

Many chapters in this volume provide examples of empirical model calibration. For example, Deconinck et al. (2016) draw on existing studies and expert knowledge

to calibrate their model of severe acute malnutrition in terms of important population parameters, thresholds, and decision processes. Kluge and Vogt (2016) use realistic demographic rates (e.g., age-specific death rates and transition into public pensions) to study how intra-familial transfers might help explaining observed patterns of old-age survival. Similarly, Williams et al. (2016) develop an agent-based model of the relation between armed conflicts and population change, in which key individual-level decision processes are implemented as probabilities estimated from a combination of survey and register data, which can be conceived as “evidence-based action rules” (Hedström 2005, p. 132).

An application where the combination of theory-driven ABM with empirical data may prove to become particularly useful for demography is in the field of population projections and forecasts – one of the key areas of applied demography, closely linked with nation state oriented demography using a ‘closed’ concept of population, as discussed above. Projections proceed by calculating the implications of a set of theoretical ‘what if’ assumptions about demographic rates; forecasts have the ambition to yield realistic predictions about actual population trends in the future. So far, forecasts as well as projections typically extrapolate macro-level trends without being based on clear theories about the underlying micro-level behaviour. ABM may help to improve this. While Prskawetz (2016) reminds us that explanation rather than prediction is the primary purpose of ABM, she still hints at how it can be used to improve demographic projections when she discusses the model presented in Aparicio Diaz et al. (2011) about the transition to parenthood. This highlights the potential of ABM to use theoretically informed simulation models to generate potential population trends rather than just relying on extrapolations of past and ongoing aggregate trends (Prskawetz 2016). ABM may also be instrumental in integrating classic scenario-based projections and more recent probabilistic approaches (Lutz and Goldstein 2004; Willekens 1990; Wilson and Rees 2005). Classic scenario-based projections are mechanistic and fail to quantify uncertainty. ABM offers the opportunity to really simulate scenarios while accounting for heterogeneity on the individual level, including random components and probability distributions, and to see what macro-level population patterns emerge. To facilitate this process, researchers can rely on advanced statistical tools that make it possible to systematically explore the uncertainty that exists in the outcomes of the agent-model, as illustrated in the methodological chapters by Hilton and Bijak (2016) and Grow (2016).

Evidently, increasing the integration of empirical data and existing methods with ABM will also pose new challenges in terms of the complexity of the modelling process. Richardi and Richardson (2016) provide an example of how some of these challenges can be overcome with a new software tool that makes it possible to combine micro-simulation with agent-based models and that allows easy data handling.

1.3.3 *Avoiding Potential Pitfalls*

In the end, any theory and model, as well as population forecasts, will have to be tested with empirical data. In the empirical testing of an ABM, it is important to keep in mind that the ability of the model to simulate (or ‘to grow’) an observed empirical pattern is far from sufficient proof of its validity (cf. also the chapter by Courgeau et al. 2016). It is *not* sufficient to show that an ABM can produce results that are compatible with some set of observed data because the model may contain so many parameters that it can be fitted to any set of data (Grimm et al. 2005). Or, as rightly pointed out by Smaldino and Schank (2011, p. 13), “if there are no empirical constraints on assumptions, almost any results can be generated from different decision rules by varying assumptions about the environmental structure”. The problem is that the model may be too flexible to draw firm conclusions about its validity; that there are too many degrees of freedom. The example provided by Smaldino and Schank nicely illustrates the issue: they show that very different models, involving very different but plausible decision-making rules involved in human mate choice, may all explain equally well the typical right-skewed distribution of age at marriage (i.e., the distribution that was also targeted by Billari et al. 2007 and Todd et al. 2005). The fact that a set of mechanisms implemented in an ABM is able to explain some patterns of empirical observations, even if all the available evidence has been used to calibrate the model, is therefore insufficient proof that these mechanisms actually generated these patterns.

The challenge of having ‘too many degrees of freedom’ may look like a limitation of ABM at first sight, but equivalent challenges apply to any kind of modelling. Conventional statistical models face similar challenges: very different models might fit the data equally well and a statistically significant ‘effect’ may actually be spurious, even when we have a plausible theory to portray it as a causal effect. Attacking ABM because it would claim to replace the role of empirical observation in the scientific endeavour (like Venturini et al. 2015 do) is therefore an attack on a straw man and misses the point. There is no antagonism between ABM and empirical observations. The one cannot replace the other; to rule out alternative scenarios and parameter values, scientists will still have to confront the model with empirical data that can help making the distinction.

There are two general ways in which a model can be confronted with empirical data to assess its validity. First, if there is a lack of empirical data, a model developed during the stage of theory formulation may guide subsequent data collection to advocate between model alternatives (cf. also Courgeau et al. 2016). In this step, ABM can also make the data collection effort more efficient, by sorting out potential candidate mechanisms *before* any data is collected. The chapter by Gray et al. (2016) in this volume illustrates this point. Given the lack of empirical data on how women decide whether or not to disclose their drinking behaviour to midwives, the authors explore several plausible decision mechanisms derived from existing decision theories. Their results suggest that there are characteristic differences in the results that the different decision models generate and this information can guide future data collection efforts to advocate between them.

Second, if empirical data does already exist, researchers can assess whether the model does not only reproduce the target outcome, but also other outcomes that were not in the focus of model development. This form of ‘pattern-oriented modelling’ (Grimm et al. 2005) aims at assessing the structural realism of models and helps to find the optimal zone of model complexity: addressing multiple patterns helps avoiding models that are too simple in structure and mechanism, or too complex and uncertain due to the high number of parameters. The agent-based model developed by Grow and Van Bavel (2015a) provides an example of such structural realism. The model was developed and calibrated with the goal to generate realistic patterns of educational assortative mating in the light of changing educational attainment in Western industrialized countries. Although patterns of divorce were not a target during the development of this model, Grow and Van Bavel (2015b) could show that it is able to also predict recent trends in divorce.

Summing up, the ability of a theory, implemented in an ABM, ‘to grow’ an empirically observed pattern or trend from the bottom up is insufficient to consider it a scientifically sound explanation. Such an ABM may be nothing more than ‘a good story’, while other stories may explain the empirical observations just as well. As always, the job on the to-do list then is to come up with clever ideas to set up a competition between different explanations and to collect new data that may differentiate between the right and the wrong story.

1.3.4 Bridging the Gap

We have emphasised that ABM allows linking the micro with the macro level and that its theory-based simulation approach allows computing the implications of hypothetical and empirically informed rules of action and interaction on the macro as well as the micro scale, and all scales in between. In practice, then, bridging ‘open’ and ‘closed’ concepts of population by means of ABM might work in two major steps. The first step consists of the in-depth study of actions and interactions in local populations, including rare events and exceptional instances as well as experiments. In combination with pre-existing theoretical frameworks and insights from earlier work, rules of behaviour and interaction (between individuals as well as with the environment) may be implemented in the simulation model. Already in this first step, ABM may be used for computational experimentation and to calibrate a model that is able to replicate (“grow”, Epstein 2006) the local observations. In a second step, ABM is used to simulate the micro- and macro-implications of hypothetical rules of action and interaction outside the original context. Part of the work involved in this second step, in order to make the jump towards quantification in a closed population, is to infuse the models with real-life observational data, which in demography will notably be information about the distribution in that closed population by variables such as age, sex, and education (see Grow and Van Bavel 2015a for an example), i.e. exogenously infusing the ABM with information from traditional demographic approaches in order to make

the model demographically realistic. It is only when the model is shown to work outside the context of where it was originally developed that its external validity can be demonstrated (cf. Hedström and Swedberg 1998).

To illustrate this process, consider psychological research that, more than any other field within the social sciences, has a long and rich tradition of conducting experimental studies to test theories. In combining ‘open’ and ‘closed’ population concepts, such experiments can be a first step to gain insights into individual behaviour and decision processes under controlled conditions. One shortcoming of such experiments is that they are often based on convenience samples, with undergraduate college students heavily over-represented in the data gathered, and focus on behaviour under sometimes unrealistic conditions. To avoid that the theoretical claims tested in such experiments hold only true for “the weirdest people in the world” (cf. Henrich et al. 2010) in the artificial context of the laboratory, a second step is needed. This second step does not just consist of collecting the same kind of samples as in the original experiments to check whether the predictions hold true in other samples as well. The true test of the theory is to study the patterns that the theory implies at other levels of observation and in the context of different populations. If the theory can correctly predict patterns at other levels of aggregation and for contexts in which the theory was not originally developed, this indicates the validity and structural realism of the model and underlying theory (Grimm et al. 2005; Hedström and Swedberg 1998).

To further illustrate this process, consider a specific example from the field of population studies. In today’s Western societies, in which feelings of mutual attraction are considered a key determinant of heterosexual marriage, knowledge about the characteristics that men and women prefer in each other is crucial to understand how observed marriage patterns come about (cf. Buss et al. 2001). Over the last years, research in sociology, psychology, and economics has devised ingenious ways to gain insights into these preferences, for example, by use of census and survey data (e.g., England and McClintock 2009), vignette studies (e.g., Greitemeyer 2007), and speed dating experiments and procedural data generated by online dating platforms (e.g., Skopek et al. 2011). Apart from census data and national representative surveys, none of these sources could be considered congruent with the ‘closed’ concept of population dominating in demography. Yet, as Grow and Van Bavel (2015a) have shown, the insights gained from such small scale and highly detailed studies can help formulating theories about mate search and processes involved in union formation. ABM makes it possible to compute the implications of these theories, which can then be compared with empirical data observed in another context than the one that first inspired the theories, namely national marriage markets.

1.4 Contributions to Agent-Based Modelling in this Book

The chapters in this book address many of the issues that we have outlined up to this point. The chapters in the section ‘Perspectives on Agent-Based Modelling in Population Studies’ discuss in more detail the tasks that lie ahead and the steps that need to be taken to connect the ‘closed’ and ‘open’ concepts of population. They also highlight the benefits that agent-based modelling might yield in this process. The integration of empirical data and ABM will require some methodological advancements and the chapters in the section ‘Designing, Analysing, and Reporting Agent-Based Models’ illustrate some of the most recent developments in this direction. As argued earlier, ABM requires theories about individual behaviour and the chapters in the section ‘Modelling Decision Processes’ illustrate in detail how existing theories can be adjusted and implemented in agent-based models. Finally, the chapters in the sections ‘Family Formation and Fertility’ and ‘Health, Mortality, and Support in Old Age’ provide applied examples of how ABM can be fruitfully used to study demographic phenomena. In this section, we briefly review each of the chapters.

1.4.1 *Perspectives on Agent-Based Modelling in Population Studies*

In Chap. 2, Courgeau et al. (2016) trace the methodological developments in demography over its 350-year history and suggest that the introduction of model-based approaches to the field, such as ABM, constitutes a paradigmatic shift. This shift results from an increased interest in individual behaviour and interactions in population research and the authors highlight that in contrast to ABM, the hitherto dominant methodological approaches do not make it possible to model the ‘two-way flow’ between the micro and the macro level. Yet, they also highlight that ABM should not be seen as an alternative to other, more empirical methods in demography. Instead, in their outline of a possible research agenda for model-based demography, they make the strong point that there needs to be a close connection between empirical research and ABM. This ensures that the insights into population dynamics that ABM might yield are firmly grounded in empirical evidence and are not based on arbitrary assumptions that are disconnected from reality.

In line with some of the views outlined by Courgeau et al. (2016), in Chap. 3 Prskawetz (2016) points out that there is increasing consensus in that individuals’ demographic decisions cannot be explained in isolation of the networks they are embedded in. She argues that ABM is particularly suitable to study such network effects from the bottom up and subsequently illustrates this capability of ABM with examples from her own work. Along the way, she discusses some of the central decisions that need to be taken when developing agent-based models; this will provide valuable guidance for novices to the field. In the last example, she also

highlights the capability of ABM to conduct ‘what if’ experiments and illustrates the usefulness of this possibility by showing how it can be used to assess potential policy implications.

1.4.2 Designing, Analysing, and Reporting Agent-Based Models

In Chap. 4, Richiardi and Richardson (2016) provide a step-by-step guide for a new open-source, Java-based simulation platform, JAS-mine, that makes it possible to easily combine aspects of micro-simulation models with aspects of agent-based models. The development of this platform was instigated by the observation that although micro-simulation and ABM have been developed with different goals (i.e. data-based forecasting based on probabilistic regression models vs. theory development and understanding with a focus on interactions between individuals), they also share important commonalities, such as that they are discrete-event simulations, are recursive, and that the states of individuals evolve over time. Both approaches have their unique strengths that JAS-mine aims to combine, while at the same time providing a convenient structure to separate the modelling process from the data recording process. Such developments in ABM software will greatly facilitate the grounding of agent-based models in empirical data.

In Chap. 5, Zinn (2016) illustrates how an integration of micro-simulation models and ABM, as addressed by Richiardi and Richardson (2016), can be achieved. As she points out, micro-simulation lends itself to conducting fine-grained population projections under the assumption that individuals do not interact with each other. If this is combined with ABM’s capability to model social relations and interactions, it becomes possible to model life courses of both individuals and couples at the same time. For this, Zinn relies on the ml-DEVS formalism and implements the model in the simulation framework JAMES II. Her exemplary analysis attests to the potential of this approach and her work provides a frame of reference for those interested in combining micro-simulation with ABM.

Next to having the technical possibility to infuse agent-based models with empirical data, it is important that the field develops ‘best practices’ as to how empirical data should be used. In Chap. 6, Williams et al. (2016) provide one of the first steps in this direction. Drawing on related research in geographic and land use sciences, the authors illustrate how various sources of information (in particular survey data) can be used to implement a detailed representation of a specific population, both in terms of structure and decision processes. They also illustrate how the resulting model can be used to conduct ‘what if’ experiments to gain deeper insights into the processes that underlie observed population changes.

Even if a model has been calibrated with empirical data, there often is uncertainty in terms of how different model aspects (i.e. different parameters) relate to model outcomes and under which conditions the model actually is able to reproduce

observed population patterns. One way to deal with this uncertainty is by the systematic design of simulation experiments combined with metamodels. In a nutshell, metamodels treat a simulation model as a black box and express the relation between model inputs and outputs by means of a statistical function. In Chap. 7, Grow (2016) illustrates the use of metamodels based on ordinary least squares regression analysis, whereas in Chap. 8 Hilton and Bijak (2016) illustrate the use of Gaussian process emulators. As the authors point out, metamodels based on regression analysis are an efficient tool to explore and describe input/output relations that can be described with polynomials. Gaussian process emulators make it possible to describe even more complex input/output relations and provide additional information about model uncertainty.

As Courgeau et al. (2016) argue, theory development by means of ABM will require explicit documentation of the way in which the simulation model was constructed and what assumptions guided this process. In Chap. 9, Groeneveld et al. (2016) review existing practices of model description in demographic research and come to the conclusion that so far no standard has emerged. After making the case that standardized descriptions can yield many benefits (e.g., enhanced replicability), they suggest the ODD+D standard as a possible candidate. Based on their experiences with an exemplary application to a demographic ABM, they also make recommendations as to how the standard could be adjusted to accommodate some aspects specific to demographic simulations.

1.4.3 Modelling Decision Processes

In developing agent-based models, researchers often have to draw on theories that were not developed with a procedural and dynamic focus. In Chap. 10, Willekens (2016) shows how existing theories from other fields of social research can be adjusted to better fit with the process-orientation of ABM. He uses the theory of planned behaviour to model the decisions that underlie international migration. For this, he extends the theory, so that it takes into account that the decision to migrate has a (random) processual character: the decision consists of several stages and it takes individuals time to transition from one stage to the other, contingent on systematic and random factors. He parametrizes the resulting simulation model with data from the Gallup World Poll 2005 and other sources and shows that it reproduces some stylized facts of international migration.

Agent-based models are often criticised for being based on ad hoc assumptions about individual behaviours and decision processes. In Chap. 11, Gray et al. (2016) address this issue by drawing on a long tradition of research in decision theory for modelling women's decision to disclose alcohol consumption during pregnancy to midwives. The authors frame the decision as a game theoretic problem in which both women and midwives are uncertain about the motivations and behaviours of each other. In the resulting signalling game, the authors compare four different decision models that differ in the complexity of the representation of the decision process

within individuals. The results of the simulation experiments show that the different rules lead to somewhat different outcomes and therefore also lead to different recommendations for ways to enhance disclosure by women. This highlights the need to collect additional detailed data in this area where empirical insights are so far limited.

1.4.4 Family Formation and Fertility

In Chap. 12, Kashyap and Villavicencio (2016) explore the mechanism that might explain the rise in the sex ratio at birth (measured as the number of males per 100 females) that has accompanied the fertility decline over the last decades in Asia and the Caucasus. Congruent with earlier theoretical research, the model conceptualises sex ratio imbalances as the result of an interplay between son preferences, technology diffusion, and fertility decline. Using UN data to validate the model in the contexts of South Korea and India, one of the central insights of this study is that even if son preferences would have declined, an increase in the sex ratio at birth can arise from an increase in the accessibility of techniques that make sex-selective abortion possible combined with a decrease in total fertility levels. An important strength of this study is its cross-national approach, that attests to the generality of the processes that are modelled.

In Chap. 13 Klüsener et al. (2016) study the role that socially and spatially structured communication and influence processes might have played in the historical fertility decline observed in Sweden between 1880 and 1900. The chapter illustrates how the creative use of available census and GIS data facilitates conducting ‘what if’ experiments that help to uncover some of the processes that might have contributed to observed changes in (historical) populations whose members (and their interactions) cannot be studied in depth anymore. The results suggest that their diffusion model can reproduce many of the spatiotemporal properties of the observed fertility decline. In Chap. 14, Ciganda and Villavicencio (2016) also explore the mechanisms that might have generated observed trends in fertility, but in a more recent time period (1944–2014) in Spain. The authors model these trends as the outcome of an interplay between educational expansion (increasing the average opportunity costs for having children), increasing economic uncertainty, and social influence processes. The model illustrates how effects from factors exogenous to the social interactions under consideration can be amplified by precisely these interactions.

1.4.5 Health, Mortality, and Support in Old Age

In Chap. 15, Kluge and Vogt (2016) employ the case of the German reunification in 1990 as a natural experiment to address the question whether the positive association

between income and old-age survival comes about through the goods and services that income can buy, or through third factors that affect both. In their modelling efforts, the authors focus on intra-familial exchange as a potential source of the observed association and draw on a variety of data sources for calibrating the model. Interactions occur within families and concern the exchange of income of parents for care from their children. The model can generate part of the observed changes in old-age survival in Eastern Germany after reunification and suggests that this increase might be partially caused by an increase in purchasing power and an increase in intra-familial exchanges.

In Chap. 16, Deconinck et al. (2016) show how ABM can be used to inform intervention strategies to reduce the effects of severe acute malnutrition. The authors highlight that the design and study of such interventions suffers from a lack of data and understanding of health system dynamics. They suggest that the theoretical, rule-based nature of ABM makes it possible to study factors that might potentially affect the effectiveness of interventions despite lack of data. For this, it is central to involve subject matter experts and practitioners in the model development process, to create accurate representations of the decision rules and interactions that occur in the actual system and to raise awareness among potential stakeholders.

In Chap. 17, finally, Wolfson et al. (2016) study the puzzling observation that in the US there exists an association between city-level income inequality and mortality, whereas no such association exists in Canada. Their main intuition is that this difference might be caused by the fact in US cities income segregation tends to be higher than in Canada. That is, in the US, there is more residential segregation in terms of income than in Canada and this might indirectly affect mortality rates through the properties of the communities (e.g., school quality) that feed back into the individual characteristics relevant for mortality (e.g., educational attainment). Using a simulation model that incorporates interactions between aspects of different layers of society (i.e. individuals, families, neighbourhoods, and cities), the authors find that their model is indeed able to generate patterns of mortality that are similar to those observed in reality, but for reasons that are different from what they expected.

1.5 Towards Evolutionary Population Theory

As we have indicated earlier, ABM is a useful method to help developing population thinking. The method itself is agnostic about the theory that is used to reason about the mechanisms that link the micro and the macro level. The diversity of the theoretical approaches used in the chapters of this volume attest to this flexibility of ABM. Yet, if demographic phenomena are phenomena *sui generis*, what kind of theory can we reasonably be looking for to explain them? In this closing section, we describe why we think that evolutionary theory is a particularly attractive candidate for this. Note that this represents our view, which does not necessarily represent the views of the other contributors to this volume.

1.5.1 *The Basic Tenets of Evolutionary Theory*

Inspired by thermodynamics, Lotka (1945) still had a concept of theory in mind consisting of a system of “laws” within which, “by the application of relatively few fundamental principles, the course of events can be rigorously *deduced* for innumerable specific situations” (Lotka 1945, p. 172, italics as in original). However, “demography is neither theoretical physics nor is it mineralogical chemistry”: with this truism, Charbit (2009, p. 48) wants to highlight something he thinks is particular for the human sciences: because demography is a human science, theories are based on factors that are peculiar to a given historical context. Indeed, doing social science is not about finding eternal laws that allow us to predict the future. It might therefore be tempting to dismiss altogether the idea of a general theory of population and to stick with idiosyncratic narratives that might explain in a particular context why things happened the way they did.

Although we agree that historical peculiarities do and should play a role in social scientific research and theory, one could also argue that this epistemological point of view reveals a lack of ambition for the social sciences. Why would this argument hold for the social sciences and not for the biological sciences? Aren't plants and animals, in their phenotypic appearance and behaviour, also peculiar to their historical environment? It is precisely the uniqueness of every plant and animal that is highlighted in the populationist biology inherited from Charles Darwin (Mayr 2002, pp. 90–93). Darwinian evolutionary theory can be considered superior to the earlier, essentialist ways of theorizing about biological diversity because it is able to account for the changing biological diversity and developments that occurred in time not only before, but also after the formulation of the theory (Boyd and Silk 2009; Mayr 2002); it is able to “describe and explain phenomena with considerable precision”, even if it cannot make reliable predictions about the future (Mayr 1961, p. 1504).

While demography and evolutionary biology have followed very different and increasingly divergent pathways after the Second World War, a Darwinian renaissance got started in recent decades, with an increasing number of papers inspired by evolutionary theory being published in mainstream demography journals (Sear 2015a). It would be good to intensify the conversation between demography and evolutionary theory. We concur with Sear (2015b) that the endorsement of evolutionary demography does not at all imply that evolutionary theory would be the only theoretical framework that has value in explaining demographic behaviour, but rather that it can inform, enrich, and stimulate theory development in our field.

The key ideas of evolutionary theory in biology are simple, but nevertheless often poorly understood: in a nutshell, organisms evolve through variation and differential selection. No two living organisms are exactly the same; for both genetic as well as environmental reasons, there is always variation. Not all variants survive and produce offspring in the next generation to the same extent. Those variants that survive and produce a lot of offspring in a given environment have high fitness, which by definition implies that such variants will become more common in the

next generation; variants with fewer offspring will be encountered less frequently in the next generation. This is what is meant by differential selection: in a given environment, some variants will become more common over the generations, others will become less common. Features or variations that lead to high fitness in their environment are called adaptive (Mayr 2002). Of course, environments can and do change, implying that well-adapted organisms at one point in time may turn out to be very badly adapted to the new situation – ‘maladapted’, implying no more nor less than that they will become rarer over the generations.

This basic mechanism is key to explaining how humans and other living organisms evolved (Boyd and Silk 2009). The basic principles have also been applied to the evolution of culture (Richerson and Boyd 2005), although such application of evolutionary theory is still less widely accepted. The same holds for more recent models of gene-culture coevolution. Such models are being developed since it is becoming clear that culture has affected and is affecting the human genome (Laland et al. 2010) through processes such as niche construction (Kendal et al. 2011).

One of the reasons why evolutionary theory seems suitable as a general theoretical framework for human demography (and, more generally, the social sciences) is that it does justice to the fundamental contingency of human populations and societies. Evolutionary theory is not deterministic. Rather to the contrary: it is fundamentally probabilistic and acknowledges the fundamental contingency of life. Evolutionary theory does not allow to predict the substance of the future because it does not have information about the substantive direction. Instead, evolutionary theory contains of “a set of interacting mechanisms resulting in the production of variation and its selection” (Hammel and Howell 1987, p. 142).

Evolutionary theory is not teleological (Mayr 1961, 2002); there is no need to assume that evolution has a direction (in contrast to what has often been claimed, see, e.g., Lotka 1945) It does certainly not claim that evolution leads to perfection (even if we would know what perfection is), nor does it imply that things evolve to always get better – in biological evolution, organisms that may have thrived very well in one environment, may become extinct as the environment changes. Evolutionary theory is also not essentialist. Darwin had a hard time defending his populationist approach against the essentialist claims about the ‘true’ nature of different species (Mayr 1991, 2002).

Demography and populationist thinking is already playing an important role in evolutionary theory. “Human culture and biology jointly and collaboratively drive the evolution of human demography” (Levitis 2015, p. 415). Hammel and Howell (1987) called for an evolutionary theory “in which demographic events are the central mechanism and leading indicators of the coevolution of bodies, minds, and societies” (p.142). Recognizing that birth, marriage, migration, and death have both biological and cultural significance in any human society, and that the subject matter of demography is cutting across the sub-disciplines of the social and biological sciences, they argue that a demographically based formulation of evolutionary theory may integrate important aspects of cultural and biological evolution. More recently, Metcalf and Pavard (2007a) argued that “evolutionary biologists should

be demographers” because evolution depends on fertility, migration, and mortality, as well as on population growth and structure; in other words: “All paths to fitness lead through demography” (Metcalf and Pavard 2007b). Therefore, evolutionary demography aims to cross barriers between social scientific and biological approaches to population processes by combining concepts and tools of demography and evolution, hoping to enhance the scope of both fields (Levitis 2015).

1.5.2 *Agent-Based Modelling and the Evolutionary Approach*

As indicated earlier, human populations are complex adaptive systems. Miller and Page (2007, pp. 78–89) discuss a range of characteristics of ABM that makes the approach particularly well suited to study such systems: the focus on dynamics and processes, the scalability and flexibility, the feasibility to model adaptive rather than optimizing agents, and the enhanced ability to address the role played by heterogeneity and variation. These features also make ABM particularly well suited as a tool to help developing an evolutionary approach in demography.

Agent-based models are inherently *dynamic*: even if one can take snapshots of the system’s situation at discrete points in time, the results of the model change over time and the focus is drawn to the process at least as much as to the outcome. Like evolutionary theory, ABMs are inherently *process oriented*: the focus is on understanding the mechanisms that produce or reduce diversity and change. Evolutionary theory is about mechanisms rather than “laws”, and ABM facilitates the investigation of mechanisms, where mechanisms can be considered halfway “between laws and descriptions” (Billari et al. 2003, p. 13).

Axelrod and Hamilton (1981) powerfully illustrate how a focus on dynamics may be crucial for our understanding. They showed how cooperation in populations may evolve even under conditions that, at any one point in time, imply no cooperation. A criticism by Venturini et al. (2015) on ABM maintains that it “cannot but confirm” individualistic behaviour and that it is unable to understand human cooperation. Indeed, in the first model developed by Axelrod and Hamilton, individual agents face a prisoner’s dilemma that cannot be overcome in a single shot. Yet, when iterated over time, in a second model, cooperation emerges as a viable strategy (Venturini and colleagues seem to have missed this landmark paper). More generally, when developing a theoretical model, one can aim either at reproducing important features of the target system at a given point in time, or at modelling its evolution, i.e., at reproducing the changes that would occur *across generations*. Ideally, however, a good model should be able to reproduce both aspects of the phenomenon, and ABM facilitates such combination (Campenni and Schino 2014). In line with this, evolutionary demography involves investigating both how demographic processes evolve over time and the outcomes of such evolution (including population structure and composition) at given points in time.

Evolutionary demography not only involves integrating the cross-sectional and the longitudinal, it should also integrate insights gained at different levels of

magnitude or scale and in diverging scholarly disciplines (from the molecular micro level of genetics to the macro level of human populations embedded in a globalizing society) (Kaplan and Gurven 2008). Demographic theory “thus faces the same issues raised by Darwinian population thinking: both observed population processes at a local scale and testable models at higher levels of aggregation are necessary, and theoretical formulations confined to one or the other are incomplete” (Kreager 2015a, p. 81). The *scalability* of ABMs and the *flexibility* of specifying agent behaviour and interactions are particularly useful here. The scalability refers to the ability of ABM to explore a system’s behaviour both with a very low and a very high number of agents, and to switch the focus from micro- to macro-level system properties. The flexibility refers to the fact that ABMs can capture a very wide class of behaviours, which is particularly useful for implementing the insights from different study disciplines. Agents may, for example, respond to the constraints imposed by the human metabolic system as well as to the cultural rules implied by human society. Both kinds of rules can be specified in the same ABM, and the emerging properties can be studied at the level of individual agent behaviour, at the neighbourhood level, or at the population level. Mechanisms involved in multiple inheritance models, like the triple inheritance model involving genetic, ecological, and cultural inheritance (see Kendal et al. 2011) can be implemented explicitly in ABM. Change across generations can be simulated over thousands of generations, and snapshots can be taken at each point in time, enabling comparison with real-life data employing standard statistical tools.

Given the dynamic nature and flexibility of ABMs, agents can be designed to be *adaptive*, i.e., as learning from previous experiences within or across generations, or both. This allows moving away from the unrealistic, rationalistic, and atomistic models of well-informed agents who rationally processes all the relevant information to optimize behaviour to maximize utility (Miller and Page 2007, pp. 81–83). With ABM, it is possible to specify agents that learn, build networks, gain or lose power and influence, and inherit knowledge and resources from previous generations. The criticism that ABM is inherently atomistic and apolitical (Venturini et al. 2015) is therefore poorly targeted. For application to human demography, the model of adaptive rather than optimizing agents is much more consistent with evolutionary theory as well as with basic insights from psychology and sociology.

Finally, while conventional models often assume that the underlying agents have a high degree of homogeneity, where differences are typically described in terms of conditional averages, ABM facilitates to focus more on *heterogeneity* – even if it may turn out, empirically, that the aggregate system behaviour does not depend on the details of each agent (Miller and Page 2007, pp. 84–85). ABM does not require making any assumption about the homogeneity of agent populations, which is a key advantage given that heterogeneity is a core aspect of populations and population models (Billari et al. 2003, p. 12). While the focus of conventional statistical approaches is on how averages depend on a set of variables – an approach in the tradition of “the average man” (Quetelet 1835) – this may be insufficient to do justice to the role played by diversity and variation in explaining population patterns

and change. Ernst Mayr even went so far as to imply that statistical methods do not really represent population thinking at all (Kreager 2015a, p. 78).

Enhancing the ability to address the role played by heterogeneity seems important for improving population theory, for example for improving demographic transition theory. In applications of ABM, it has become clear that a given outcome may be produced by different pathways or that a given pathway may lead to very different outcomes, depending on the size and composition of the population. Similarly, ABM has proven to be able to yield both results exemplifying convergent evolution (initially major differences in the population becoming smaller over time) as well as divergent evolution (minor initial differences that magnify over time and generations) (see Axelrod 1997). This matches very well with the observation that, while the transition from (moderately) high to low mortality and fertility in modern populations is a quite general phenomenon, uniform explanations in terms of macro-level factors and processes such as industrialisation, urbanisation, and modernisation have failed the empirical tests to a very large extent (Szreter 1993; Van Bavel 2010b; Watkins 1986).

For example, the secular decline of fertility got started under widely different economic conditions, unexplainable by standard modernisation theories, or failed to kick off when theory would have predicted this. Theories such as those developed by Frank Notestein spoke about interactions between the economy and populations largely at the macro level, without accounting for the heterogeneity within economies and populations. This approach “pushed key aspects of population variation and change to the margins” (Kreager 2015a, p. 79). Thanks to more detailed research in historical demography, often looking at very specific local communities and populations, it became clear that fertility and mortality decline can take place under widely differing conditions. This has stimulated the field to increasingly reconsider the role of local networks of communication in demographic change. In-depth study of local populations, conceived of in the ‘open’ rather than the ‘closed’ way, enabled us to understand more about the role played by distinctive environmental and cultural constraints existing prior to ‘big’ forces such as industrialization and modernization, implying that there is not one universal ‘transition’ pathway. The continuing diversity observed in demographic phenomena like ‘the’ demographic transition highlights that it will be key for demographic theory to understand the mechanisms that continue to renew population heterogeneity (Kreager 2015a, pp. 80–81), and ABM promises to be very helpful in gaining such understanding.

Acknowledgments The authors’ work on this chapter has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement no. 312290 for the GENDERBALL project.

References

- Aparicio Diaz, A., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography*, *48*(2), 559–579.
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, *41*(2), 203–226.
- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, *211*(4489), 1390–1396.
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. D. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the United Kingdom with Gaussian process emulators. *Demographic Research*, *29*, 729–765.
- Billari, F. C. (2015). Integrating macro- and micro-level approaches in the explanation of population change. *Population Studies*, *69*(S1), S11–S20.
- Billari, F. C., Ongaro, F., & Prskawetz, A. (2003). Introduction: Agent-based computational demography. In F. C. Billari & A. Prskawetz (Eds.), *Agent-based computational demography – Using simulation to improve our understanding of demographic behaviour* (pp. 1–17). Heidelberg: Physica Verlag.
- Billari, F. C., Prskawetz, A., Aparicio Diaz, B., & Fent, T. (2007). The “Wedding-Ring”: An agent-based marriage model based on social interaction. *Demographic Research*, *17*, 59–82.
- Boyd, R., & Silk, J. B. (2009). *How humans evolved*. New York: W.W. Norton & Co.
- Burch, T. K. (2003a). Demography in a new key: A theory of population theory. *Demographic Research*, *9*, 263–284.
- Burch, T. K. (2003b). Data, models, theory and reality: The structure of demographic knowledge. In F. C. Billari & A. Prskawetz (Eds.), *Agent-based computational demography – Using simulation to improve our understanding of demographic behaviour* (pp. 19–40). Heidelberg: Physica-Verlag.
- Buss, D. M., Shackelford, T. K., Kirkpatrick, L. A., & Larsen, R. J. (2001). A half century of mate preferences: The cultural evolution of values. *Journal of Marriage and Family*, *63*(2), 491–503.
- Campenni, M., & Schino, G. (2014). Partner choice promotes cooperation: The two faces of testing with agent-based models. *Journal of Theoretical Biology*, *344*, 49–55.
- Charbit, Y. (2009). *Economic, social and demographic thought in the XIXth century: The population debate from Malthus to Marx*. Dordrecht: Springer.
- Ciganda, D., & Villavicencio, F. (2016). Feedback mechanisms in the postponement of fertility in Spain. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 405–435). Springer.
- Cointet, J.-P., & Roth, C. (2007). How realistic should knowledge diffusion models be? *Journal of Artificial Societies and Social Simulation*, *10*(3), 5.
- Coleman, J. S. (1986). Social theory, social research, and a theory of action. *American Journal of Sociology*, *91*(6), 1309–1335.
- Coleman, J. S. (1990). *Foundations of social theory*. Cambridge: Harvard University Press.
- Courgeau, D., Bijak, J., Franck, R., & Silverman, E. (2016). Model-based demography: Towards a research agenda. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 29–51). Springer.
- De Bruijn, B. J. (1999). *Foundations of demographic theory. Choice, process, context*. Amsterdam: Thela-Thesis.
- Deconinck, H., Van Malderen, C., Speybroeck, N., Macq, J., & Chiem, J.-C. (2016). Agent-based modelling to inform health intervention strategies: The case of severe acute malnutrition in children in high-burden low-income countries. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 461–486). Springer.
- Dublin, L. I., & Lotka, A. J. (1925). On the true rate of natural increase: As exemplified by the population of the united states, 1920. *Journal of the American Statistical Association*, *20*(151), 305–339.

- England, P., & McClintock, E. A. (2009). The gendered double standard of aging in US marriage markets. *Population and Development Review*, 35(4), 797–816.
- Epstein, J. M. (2006). Agent-based computational models and generative social science. In J. M. Epstein (Ed.), *Generative social science: Studies in agent-based computational modeling* (pp. 4–46). Princeton: Princeton University Press.
- Gray, J., Bijak, J., & Bullock, S. (2016). Deciding to disclose: A decision theoretic agent model of pregnancy and alcohol misuse. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 301–340). Springer.
- Greitemeyer, T. (2007). What do men and women want in a partner? Are educated partners always more desirable? *Journal of Experimental Social Psychology*, 43(2), 180–194.
- Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., et al. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, 310(5750), 987–991.
- Groeneveld, J., Klabunde, A., O'Brien, M. L., & Grow, A. (2016). How to describe agent-based models in population studies? In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 237–254). Springer.
- Grow, A. (2016). Regression metamodels for sensitivity analysis in agent-based computational demography. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 185–210). Springer.
- Grow, A., & Van Bavel, J. (2015a). Assortative mating and the reversal of gender inequality in education in Europe: An agent-based model. *PLoS ONE*, 10(6), e0127806.
- Grow, A., & Van Bavel, J. (2015b). *Too few alternatives? The reversal of the gender gap in education and its consequences for union stability*. Paper presented at the IUSSP Scientific Panel on Microsimulation and Agent-based Modelling in Demography/Max Planck Institute for Demographic Research workshop on “The Science of Choice – How to Model the Decision Making Process?”
- Hammel, E. A., & Howell, N. (1987). Research in population and culture: An evolutionary framework. *Current Anthropology*, 28(2), 141–160.
- Hedström, P. (2005). *Dissecting the social: On the principles of analytical sociology*. Cambridge: Cambridge University Press.
- Hedström, P., & Swedberg, R. (1998). Social mechanisms: An introductory essay. In P. Hedström & R. Swedberg (Eds.), *Social mechanisms: An analytical approach to social theory* (pp. 1–31). Cambridge: Cambridge University Press.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *The Behavioral and Brain Sciences*, 33(2–3), 61–83.
- Hernes, G. (1972). The process of entry into first marriage. *American Sociological Review*, 37(2), 173–182.
- Hernes, G. (1976). Diffusion and growth: The non-homogenous case. *Scandinavian Journal of Economics*, 78(3), 427–436.
- Hilton, J., & Bijak, J. (2016). Design and analysis of demographic simulations. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 211–235). Springer.
- Kaplan, H., & Gurven, M. (2008). Top-down and bottom-up research in biodemography. *Demographic Research*, 19, 1587–1602.
- Kashyap, R., & Villavicencio, F. (2016). An agent-based model of sex ratio at birth distortions. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 343–367). Springer.
- Kendal, J., Tehrani, J. J., & Odling-Smee, J. (2011). Human niche construction in interdisciplinary focus. *Philosophical Transactions of the Royal Society, B: Biological Sciences*, 366(1566), 785–792.
- Kertzer, D. I., & Hogan, D. P. (1989). *Family, political economy, and demographic change: The transformation of life in Casalecchio, Italy, 1861–1921*. Madison: University of Wisconsin Press.

- Kluge, F. A., & Vogt, T. C. (2016). Linking income, transfers, and social support in an agent-based family exchange model. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 439–459). Springer.
- Klüsener, S., Scalone, F., & Dribe, M. (2016). Exploring the role of communication in shaping fertility transition patterns in space and time. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 369–403). Springer.
- Kreager, P. (2009). Darwin and Lotka: Two concepts of population. *Demographic Research*, *21*, 469–501.
- Kreager, P. (2015a). Population and the making of the human sciences: A historical outline. In P. Kreager, B. Winney, S. Ulijaszek, & C. Capelli (Eds.), *Population in the human sciences. Concepts, models, evidence* (pp. 55–85). Oxford: Oxford University Press.
- Kreager, P. (2015b). Population theory – A long view. *Population Studies*, *69*(Supplement 1), S29–S37.
- Laland, K. N., Odling-Smee, J., & Myles, S. (2010). How culture shaped the human genome: Bringing genetics and the human sciences together. *Nature Reviews Genetics*, *11*(2), 137–148.
- Levitis, D. A. (2015). Evolutionary demography: A synthesis of two population sciences. In J. D. Wright (Ed.), *International encyclopedia of the social & behavioral sciences*, 2nd Ed. (pp. 413–419). Oxford: Elsevier.
- Lotka, A. J. (1945). The law of evolution as a maximal principle. *Human Biology*, *17*(3), 167–194.
- Lutz, W., & Goldstein, J. R. (2004). Introduction: How to deal with uncertainty in population forecasting? *International Statistical Review*, *72*(1), 1–4.
- Mayr, E. (1961). Cause and effect in biology. *Science*, *134*(3489), 1501–1506.
- Mayr, E. (1991). *One long argument. Charles Darwin and the genesis of modern evolutionary thought*. Cambridge: Harvard University Press.
- Mayr, E. (2002). *What evolution is*. London: Phoenix.
- Metcalf, C. J. E., & Pavard, S. (2007a). Why evolutionary biologists should be demographers. *Trends in Ecology and Evolution*, *22*(4), 205–212.
- Metcalf, C. J. E., & Pavard, S. (2007b). All paths to fitness lead through demography. *Trends in Ecology & Evolution*, *22*(11), 563–564.
- Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: An introduction to computational models of social life*. Princeton: Princeton University Press.
- Prskawetz, A. (2016). The role of social interactions in demography: An agent-based modelling approach. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 53–72). Springer.
- Quetelet, A. (1835). *Sur l'homme et le développement de ses facultés, ou essai de physique sociale*. Paris: Bachelier.
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone. How culture transformed human evolution*. Chicago: The University of Chicago Press.
- Richiardi, M., & Richardson, R. E. (2016). Agent-based computational demography and microsimulation using JAS-mine. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 75–112). Springer.
- Schoen, R. (2006). *Dynamic population models*. Dordrecht: Springer.
- Sear, R. (2015a). Evolutionary demography: A Darwinian renaissance in Demography. In J. D. Wright (Ed.), *International encyclopedia of the social & behavioral sciences*, 2nd Ed. (pp. 406–412). Oxford: Elsevier.
- Sear, R. (2015b). Evolutionary contributions to the study of human fertility. *Population Studies*, *69*(Supplement 1), S29–S55.
- Skopek, J., Schulz, F., & Blossfeld, H.-P. (2011). Who contacts whom? Educational homophily in online mate selection. *European Sociological Review*, *27*(2), 180–195.
- Smaldino, P. E., & Schank, J. C. (2011). Human mate choice is a complex system. *Complexity*, *17*(5), 11–22.
- Szreter, S. (1993). The idea of demographic transition and the study of fertility change: A critical intellectual history. *Population and Development Review*, *19*(4), 659–701.

- Tabutin, D. (2007). Whither demography? Strengths and weaknesses of the discipline over fifty years of change. *Population*, 62(1), 15–31.
- Todd, P. M., Billari, F. C., & Simão, J. (2005). Aggregate age-at-marriage patterns from individual mate-search heuristics. *Demography*, 42(3), 559–574.
- Tsuya, N. O., Wang, F., Alter, G., & Lee, J. Z. (Eds.). (2010). *Prudence and pressure. Reproduction and human agency in Europe and Asia*. Cambridge: MIT Press.
- Van Bavel, J. (2010a). Subreplacement fertility in the west before the baby boom: Past and current perspectives. *Population Studies*, 64(1), 1–18.
- Van Bavel, J. (2010b). The decline of Belgian fertility in the nineteenth century. What have we learned since the Princeton project? In T. Eggerickx & J.-P. Sanderson (Eds.), *Chaire Quetelet 2005. Histoire de la population de la Belgique et de ses territoires* (pp. 429–461). Louvain-la-Neuve: Presses Universitaires de Louvain.
- Vance, R. B. (1952). Is theory for demographers? *Social Forces*, 31(1), 9–13.
- Venturini, T., Jensen, P., & Latour, B. (2015). Fill in the gap. A new alliance for social and natural sciences. *Journal of Artificial Societies and Social Simulation*, 18(2), 11.
- Watkins, S. C. (1986). Conclusions. In A. J. Coale & S. C. Watkins (Eds.), *The decline of European fertility* (pp. 420–449). Princeton: Princeton University Press.
- Willekens, F. J. (1990). Demographic forecasting: State-of-the-art and research needs. In C. A. Hazeu & G. A. B. Frinking (Eds.), *Emerging issues in demographic research* (pp. 9–66). Amsterdam: Elsevier Science.
- Willekens, F. J. (2016). The decision to emigrate: A simulation model based on the theory of planned behaviour. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 257–299). Springer.
- Williams, N. E., O'Brien, M. L., & Yao, X. (2016). Using survey data for agent-based modeling: design and challenges in a model of armed conflict and population change. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 159–184). Springer.
- Wilson, T., & Rees, P. (2005). Recent developments in population projection methodology: A review. *Population, Space and Place*, 11(5), 337–360.
- Wolfson, M., Gribble, S., & Beall, R. (2016). Exploring contingent inequalities: Building the theoretical health inequality model. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 487–513). Springer.
- Zinn, S. (2016). Simulating synthetic life courses of individuals and couples, and mate matching. In A. Grow & J. Van Bavel (Eds.), *Agent-based modelling in population studies: Concepts, methods, and applications* (pp. 113–157). Springer.

Chapter 2

Model-Based Demography: Towards a Research Agenda

Daniel Courgeau, Jakub Bijak, Robert Franck, and Eric Silverman

2.1 Introduction

Over its 350-year history, demography has progressed through successive paradigmatic changes, from period analysis (Graunt 1662) to multilevel analysis in the more recent period (Courgeau 2007). Currently, the prominence of agent-based models (ABMs) has indicated an increased focus on individual behaviours and interactions in the study of populations, and also a desire to bolster the theoretical foundations of demography (Burch 2003a, b; Silverman et al. 2011). Here we posit that ABMs have a potential to become a manifestation of a broader, model-based research programme, which would be much more heavily reliant on computer simulations as a tool of analysis. The key advantage of such methods is that they allow examining interactions between various elements of complex population systems. In our view, such model-based approaches, while firmly rooted in the multilevel paradigm, can form the foundation of the next step in the cumulative progression of demographic knowledge.

This chapter proceeds first by detailing the successive paradigmatic changes evident in the history of demography in Sect. 2.2, and then by describing the challenges of studying uncertainty, complexity and interactions in population

D. Courgeau
Institut national d'études démographiques, Paris, France

J. Bijak (✉)
Department of Social Statistics and Demography, University of Southampton, Southampton, UK
e-mail: j.bijak@soton.ac.uk

R. Franck
Université catholique de Louvain, Louvain-la-Neuve, Belgium

E. Silverman
School of Computing, Teesside University, Middlesbrough, UK

systems in Sect. 2.3. In Sect. 2.4 we discuss the ways of conforming computational methods to the classical scientific programme outline, and make the case for model-based demography as a new research programme for the discipline. Finally, in Sect. 2.5 we propose a research agenda to address the challenges ahead.

2.2 Cumulativity in Demography

Since the origin of demography in the seventeenth century, the field has progressed through a series of paradigmatic changes. Here we use the term paradigm in a somewhat different sense from Kuhn (1962), and from its current usage. We want to point out the methods by which the phenomena observed within a population have been related to the set of key parameters (fertility, mortality and migration) used in demography to explain population growth, decline or stabilization. There are four main methods, each implementing a limited scope of notions which we call the paradigm of the selected method (Courgeau and Franck 2007). In this section we point to the differences between these four paradigms, and to their possible complementarity.

It is important to recall the path taken by Bacon in 1620 in his elaboration of an inductive method for scientific thought. He presented it in contrast with the dominant way of thinking in fashion at this time (Bacon 1620, aphorism 19):

There are and can be only two ways of searching into and discovering truth. The one flies from the senses and particulars to the most general axioms, and from these principles, the truth of which it takes for settled and immovable, proceeds to judgement and to the discovery of middle axioms. And this way is now in fashion. The other derives from the senses and particulars, rising by a gradual and unbroken ascent, so that it arrives at the most general axioms at last of all. This is the true way, but as yet untried.

The first way generates what Bacon called the four Idols, where axioms are not grounded on a meticulous observation of the properties of nature to be studied, but rather on prejudices – unverified notions of human understanding. As Bacon said¹, such axioms should not “avail for the discovery of new works, since the subtlety of nature is greater many times over than the subtlety of argument”. As already argued elsewhere (Courgeau et al. 2014) ‘Idols’ may exist in various areas of contemporary population sciences, for example in the form of behavioural genetics, postmodern theory, hereditarianism, or modern hermeneutics.

On the other hand, the Baconian “second way” became the modern scientific way of thinking, rising from a meticulous observation of studied facts to the “formation of ideas and axioms by true induction”. This method of induction² consists of discovering the principles – the ‘first’ axioms, the ‘lesser’, and the ‘middle’ in

¹Citations in this and in the next paragraph come from Bacon (1620), aphorisms 24, 39 and 40.

²*Induction* is not taken in the sense of Mill (1843) and his followers, i.e. generalisation from particular facts. In Bacon’s sense, induction designates the complete research process (Sect. 2.4).

Bacon's terms – of natural or social properties by way of experimentation and observation. The Baconian induction rests on the requirement that without these principles the properties observed would be different (Franck 2002a).

Graunt (1662) was the first to apply this method for the study of human populations. He no longer considered that phenomena such as births, illnesses and deaths were to be seen as God's secret and therefore out of bounds to scientific scrutiny. He studied each event not as a unique one but as one occurring to a statistical individual, with only a few characteristics. These abstract events became fertility, morbidity and mortality, and lost any direct attachment to a given individual. This was the only way to begin a scientific study of population, called by Petty (1690) *political arismetick*, which prevailed for around 200 years. Graunt's research paved the way for demography, epidemiology, political economics, and for population sciences more generally.

Graunt's demonstration of the links between probability and population science was also vital. Probability was first addressed in 1654 by Pascal and Fermat, but their results were published later (Pascal 1665), and it was in fact Huyghens (1657) who first published a treatise on games, with a foundation for objective probability. Graunt used this concept in order to estimate the population of London from the number of deaths, using an estimation of the probability of dying (Courgeau 2012). The probability of an abstract event in a human life was used for the first time, facilitated by the notion of a *statistical individual*.

We can conclude that the population sciences were without a doubt born in England, and subsequently led to a more general school of scientific thought on population problems. From the end of the seventeenth and throughout the eighteenth century, this way of thinking developed through the work of many leading European researchers such as Halley, Süssmilch, Euler, Moheau, and so on (Courgeau 2007).

During the eighteenth century a new concept of epistemic probability was introduced, first by Bayes (1763), and then refined by Laplace (1774, 1812). In fact, the objective probabilistic approach was already showing how through successive trials, the estimated frequency tends towards such a probability, as is the case for fair games for which we can determine an a priori probability. However, as had been already recognised in the preceding century, such a hypothesis was difficult to justify for events in human life. A new approach was necessary for such events, where all we know is the sample observed. Not only is the population from which it is drawn unknown, but its very existence is a hypothesis. By using a *prior* probability in order to estimate a unique *posterior* one, the epistemic approach permitted answering these questions clearly. Laplace applied it to many phenomena, including a number of population science ones.

In 1809, in Germany, Gauss proposed the method of least squares, which was mainly used in astronomy at this time. Following the work of mainly British but also German and French biological and social scientists³, it became widely used:

³These fields were not so clearly defined at this time: scientists were working in different social or biological sciences and in statistics simultaneously.

by Galton and Pearson in population genetics; Lexis and Quetelet in demography; Edgeworth in statistics; Durkheim in social demography; and Yule in economic demography. Finally, at the beginning of the twentieth century Fisher, a statistician and population geneticist, developed the maximum likelihood theory and theory of statistical inference (see Courgeau 2012). Population science was coming back to an objective approach, as the development of censuses through the nineteenth century permitted the use of exhaustive samples.

The problem with many statistical tools from the nineteenth century, such as the least squares method, is often that they assume a particular mathematical structure among a limited set of macro-variables, irrespective of whether that structure exists in the real world. This was already the case for Durkheim's study of suicide in Prussia (1897). This may lead to what is called an *ecological fallacy*, meaning that aggregate data, as a rule, cannot be used to study individual behaviour. The only instance where this is possible is when the probability of experiencing the event is independent of the area studied and when the population is large enough to cancel out any random difference that may appear (Courgeau 2007).

Another issue here is related to the type of observations, which are exclusively *cross-sectional* or *period-based*. After Courgeau (2007), we can conclude that the paradigm of the cross-sectional approach may be defined as follows: the social facts of a period exist independently of the individuals who experience them, and can be explained by various characteristics of the surrounding society, such as economic, political, religious, or social aspects. This cross-sectional paradigm prevailed in demography till the end of World War II.

The next change came from the US, where population scientists set up a new perspective of *cohort analysis*, following the pioneering sociological work by Mannheim (1928), which introduced the individual's lived time; Whelpton (1949) and Ryder (1951) were the first to promote this approach, and Henry (1959) formalised its theoretical underpinnings. The resultant paradigm is defined by the following postulate: "the demographer can study the occurrence of only a single event, during the life of a generation or a cohort, in a population that preserves all its characteristics and the same characteristics for as long as the phenomenon manifests itself" (Courgeau 2007, p. 36). We will not go further into this approach, as "for the analysis to be feasible, the population must be regarded as homogeneous and the interfering phenomena must be independent of the phenomenon studied" (Courgeau, *idem*, discussing Henry 1959; Blayo 1995). These conditions are restrictive, and led to a new approach permitting us to set such hypotheses aside.

To be able to consider heterogeneous cohorts and to introduce dependencies between phenomena, it became necessary to introduce statistical methods able to analyse different processes simultaneously and look at numerous characteristics of the studied individuals. The general theory of stochastic processes was first developed by the US statistician Doob (1953) and was applied to demographic processes by Aalen (1975) in Sweden. In demography, it was incorporated through the introduction of *event-history analysis* (Courgeau and Lelièvre 1992).

In the event-history paradigm, “individuals follow complex, life-long trajectories that depend, at a given instant, on their earlier trajectories and on the information they had acquired in the past” (Courgeau 2007, p. 58). We can identify the factors at work – both demographic and non-demographic – and analyse their effect on individual behaviour in more detail. In order to do that, it is necessary to use surveys that follow individuals along a large part of their life and to collect data on events, and on the characteristics, fixed as well as time-dependent, which may affect these events. However, here we cannot view an individual trajectory as the outcome of a process specific to each person. As we observe only a single outcome (the individual trajectory), the process is not identifiable.

In this case, we must adopt a collective point of view: all individuals are assumed to follow the same random process, the parameters of which we can estimate from the observation of a sample of individuals with their own characteristics. At first glance, this assumption seems quite bold. However, it is important to realize that this is not a hypothesis about observed people, but about the construction of a process underlying a set of trajectories. In this case, two observed individuals do not necessarily follow the same process, whereas two statistical individuals with the same characteristics do so automatically, as random sampling units with identical characteristics (subject to identical selection conditions). Such an approach again may require adopting a Bayesian point of view (Ibrahim et al. 2001), as it looks at many characteristics measured on a sample of restricted size.

However, the use of individual characteristics may lead to an *atomistic fallacy*, as opposed to the *ecological fallacy* of cross-sectional studies. By concentrating on individual characteristics, we disregard the context in which human behaviours occur. As noted by Courgeau (2007), context clearly may influence individual behaviour, and therefore isolating individuals from the constraints imposed by the social networks of the living environment seems misleading.

We must then introduce the different types of groupings of individuals found in all human societies: social groupings, such as the family, networks of contacts (or, more generally, social networks), etc.; economic groupings, such as the firm or the organisation where a person works; educational groupings, health-care groupings; political groupings; etc. In order to consider not only the individual but different groupings we must develop new methods of *contextual* and *multilevel analysis*. These models have been elaborated by American (Mason et al. 1983) and English (Goldstein 1987) researchers.

Multilevel approaches have permitted us to solve the apparent contradiction between aggregate models and the individual, event-history perspective. Thanks to their properties, we can combine the results of the analyses at the aggregate and individual level by clarifying the apparent paradox between them. As observed by Courgeau (2007, pp. 79–80):

The new paradigm will therefore continue to regard a person’s behaviour as dependent on his or her past history, viewed in its full complexity, but . . . this behaviour can also depend on external constraints on the individual, whether he or she is aware of them or not.

This paradigm allows for removing the two fallacies mentioned before (*idem*):

The ecological fallacy is eliminated, since aggregate characteristics are no longer regarded as substitutes for individual characteristics, but as characteristics of the sub-population in which individuals live and as external factors that will affect their behaviour. At the same time, we eliminate the atomistic fallacy provided that we incorporate correctly into the model the context in which individuals live.

As we have demonstrated previously, demography has advanced effectively thanks to the introduction and refinement of successive paradigms. Each paradigm takes the shortcomings of its predecessors as a starting point and offers a method for surmounting them – without, however, erasing all the knowledge attained through earlier paradigms. Indeed, for some questions that a population scientist may wish to ask, cross-sectional analysis can suffice just as any other form of analysis may be sufficient for other issues. The same is true for some questions asked by the physicist that may be answered perfectly by Newtonian physics, without taking into account Einstein’s physics.

However, in demography these developments have not led to a patchwork landscape of competing approaches, but instead to a *cumulativity* of knowledge, despite being far from linear. This is because different paradigms take a different point of view on the studied phenomena, partly preserving some of the results of the previous ones, as the multilevel analysis compared with cross-sectional and event history analysis. As Courgeau (2012, p. 239) has put it:

Cumulativeness of knowledge seems self-evident throughout the history of population sciences ... the shift from regularity of rates to their variation; the shift from independent phenomena and homogeneous populations to interdependent phenomena and heterogeneous populations; the shift from dependence on society to dependence on the individual, ending in a fully multilevel approach. Each new stage incorporates some elements of the previous one and rejects others. The discipline has thus effectively advanced thanks to the introduction of successive paradigms.

Each of the four paradigms frames the relationship between observations and scientific object differently, and in so doing allows for new methodologies that can alleviate difficulties associated with other methods, as summarised in Table 2.1. The scientific objects of enquiry in population sciences, such as mortality, fertility,

Table 2.1 The four paradigms of demography – a summary

No.	Paradigm	Period	Key focus
1	Period (cross-sectional)	1662–	Population-level (macro) phenomena, observed and measured according to the historical time
2	Cohort (longitudinal)	1950s–	Population-level phenomena, observed and measured along the lifetime of individual cohorts
3	Event history	1980s–	Individual-level (micro) phenomena, observed and measured according to the individual time
4	Multilevel	1980s–	Individual, population, and interim-level phenomena, observed and measured from multiple perspectives

nuptiality, migration and so on, are independent of the theory used to treat them. By contrast, the relationships assumed to exist between these objects are strongly dependent on the key theory underpinning each paradigm: independence between them in cohort analysis, heavy dependence between them in event history analysis. Yet, as argued before, each paradigm also occupies a different context, and therefore previous paradigms remain relevant despite the proliferation of new ones.

The evolution of successive paradigms is an ongoing process, and the paradigms themselves are in a constant need of improvement and refinement, in order to be able to answer emerging research questions. Even the multilevel approaches do not address questions related to interactions between various elements of increasingly complex population systems. In particular, micro-level rules may be hardly linked with aggregate-level rules, while macro-level rules cannot be modelled exclusively with an individual approach, since they transcend the behaviour of the component agents (Holland 1995). As Conte et al. (2012, p. 336) said, in their *Manifesto of Computational Social Science*, such a micro-macro link:

... is the loop process by which behaviour at the individual level generates higher-level structures (bottom-up process), which feedback to the lower level (top-down), sometimes reinforcing the producing behaviour either directly or indirectly.

We will add that in some cases it can go in the opposite direction of the producing behaviour, leading to “perverse effects” as shown by Boudon (1977).

We must go further, however, as the effects of aggregation levels are always defined with respect to the individual. For example, a series of individual actions in a community may foster awareness of a problem that concerns the entire community. This may lead to political measures, taken at more aggregated levels. These measures will naturally affect individual behaviours, generating new actions to offset their perverse effect, and so on. The multilevel approach as described above does not allow for inclusion of this two-way flow. More generally it is necessary to identify the different levels as truly different systems of agency, i.e. of collective action with different goals, specific resource interdependencies between members and specific social processes that help members to manage dilemmas at each level. We will see in the following sections how a model-based research programme may answer these challenges.

2.3 From Empirical to Model-Based Demography, and Back: Uncertainty, Complexity and Interactions in Population Systems

The recent evolution of demography and population studies has coincided with shifting perspectives on the epistemological challenges facing the studies of human populations. In particular, demographers are now paying ever more attention not only to different levels of analysis, but also to the uncertainty and complexity of population phenomena, which are discussed in this section.

Demographic phenomena – as all other aspects of social reality – are inherently *uncertain*, but to a slightly lesser degree than is the case in other areas of social sciences, such as sociology or economics. This comparative advantage of population science is largely due to the strength of the underlying relationships, such as population accounts and persistence of demographic patterns in time, and is helped by the strong empirical slant of population science (Xie 2000; Morgan and Lynch 2001). Still, particular areas of demographic interest differ with respect to their uncertainty: out of the three main components of population change, mortality is usually thought to be the least uncertain, while migration is the most (e.g. NRC 2000). The explicit acknowledgement of the uncertainty challenge has led to a renaissance of statistical demography since the 1980s, and to the “return of the variance” to demography – an important methodological perspective for all four paradigms mentioned above (Alho and Spencer 2005; Courgeau 2012)⁴.

Uncertainty is vastly augmented by social reality becoming increasingly *complex*. Hence, appropriate tools are required to analyse the associated complexities in more depth. In demography, the debate on the complexity versus the parsimony of demographic models has been present especially in the context of predictions (Ahlburg 1995; Smith 1997; Lutz 2012; Raftery et al. 2012). However, the evidence regarding the relative performance of models of varying complexity is inconclusive. For predictive applications it may be tempting to apply Occam’s razor and opt for simple models that describe the uncertainty relatively well (Bijak 2010). On the other hand, despite its importance, prediction is not the only goal of enquiry in population science (Xie 2000). If the perspective shifts towards explanation, exploration, or other non-predictive applications, a different approach is required⁵.

From a statistical point of view, *model uncertainty* needs to be acknowledged as well (Raftery 1995). If the models themselves are to be formally recognised as yet another source of uncertainty in population studies, next to the underlying processes, parameters, and inherent randomness, the most natural and coherent way of describing all these sources is via Bayesian statistical inference and epistemic probability (for details, see Bijak and Bryant 2016). Within the Bayesian paradigm there exist several approaches to model error: from a formal *model selection* out of several competing possibilities, and the related *model averaging* (Raftery 1995); to including an additional *model discrepancy (inadequacy)* term in the modelling process (Kennedy and O’Hagan 2001). In addition to the appealing prospect of reconciling quantitative and qualitative information in a formal way, Bayesian statistics allows for the inclusion of subjective opinion in the process of statistical inference.

⁴Similarly, acknowledgement of the role of space in demography has led to the multi-regional perspective within the cohort paradigm (Rogers 1975), later extended to the multi-state case.

⁵See Epstein (2008) for “sixteen reasons other than prediction to build a model”. Conte et al. (2012) highlight the capability of “generative” models to reproduce qualitative regularities observed in the real world (the stylised facts).

On a larger scale, Bayesian statistics also provides a possible way of reconciling the empirical and computational approaches by returning to empiricism, yet at a different level of analysis. All computational models, no matter how complex, have inputs (parameters) – and outputs (quantities of interest). Their mutual mapping enables statistical analysis. There are techniques available for this purpose, chiefly *Bayesian melding* (Poole and Raftery 2000), and approaches based on *Gaussian process emulators*, also Bayesian (Kennedy and O’Hagan 2001; Oakley and O’Hagan 2002). Both have already been prototyped in demographic applications – the former by Alkema et al. (2007) and Clark et al. (2012), and the latter by Bijak et al. (2013), Silverman et al. (2013), and Hilton and Bijak (Chap. 8, this volume). The application of such methods allows for analysing the properties of complex computational models within a formal statistical framework, which would not be possible with more traditional approaches.

As demography has started incorporating insights regarding its own epistemological limits, new approaches to modelling have begun to flourish. The perspective of population science becoming a *model-based science* (Burch 2003b) has become appealing⁶, mirroring similar movements within the study of biological systems and evolution (Levins 1966; Godfrey-Smith 2006). As argued by Xie (2000), there are certainly insights to be gained from examining the successes and failures of modelling efforts in population biology (see also Bullock and Silverman 2008).

Previous efforts have outlined various approaches toward modelling the *complexity* of population processes, amongst which we can identify two broad trends: *social simulation* and *systems sociology* (Silverman and Bryden 2007). The former is concerned with the application of novel modelling techniques, primarily agent-based models, to specific populations and situations. The latter is a primarily theory-driven enterprise, investigating the consequences of various foundational social theoretic positions – along the lines of the ‘opaque thought experiment’ role for simulations proposed by Di Paolo et al. (2000). Within demography and population sciences, the desire to remain empirically relevant – and to strengthen that relevance through more reliable and nuanced predictions – has led to a focus on social simulation more than systems sociology approaches. Micro-simulations, based on empirical transition rates or probabilities for *simulated (virtual) individuals* (Willekens 2005), clearly belong to this class.

Within simulation approaches, we also need to distinguish between *weak simulations* and *strong simulations* (Huneman 2014)⁷. Weak simulations serve to test some theory or hypothesis, when the system studied cannot be easily modelled by mathematics or when data are limited or unavailable. They are top-down models, which start from setting the hypotheses and assumptions. Strong simulations,

⁶Burch (2003b) points to Nathan Keyfitz (1971) as the pioneer of the model-based demography.

⁷Following Huneman (2014), we give these terms slightly different meanings than for example Thagard (1993, p. 6), for whom the weak simulation is “a calculating device drawing out the consequences of mathematical equations that describe the process simulated,” while a strong simulation “itself resembles the process simulated” (see also Brenner and Werker 2007).

on the other hand, aim to “explore the possible outcomes of a simple model” without any reference to a pre-existing theory or hypothesis (*idem*, p. 72). Many existing agent-based models often proceed in this way, where simulations are used with no pre-existing theory to explain the modelled phenomena, but only some intuitive rules. These models are built from the bottom-up: low-level interactions are supposed to produce high-level complex behaviour. As argued by Conte et al. (2012), such ‘generative explanations’ are often arbitrary – they also suggest that simulation models need to become much more empirical, in order to provide solid micro-foundations for the social mechanisms they attempt to model.

The presence of emerging properties and of ‘downward feedback’ or causation (from macro to micro) in complex models means that we cannot obtain the macro-level patterns by simply aggregating the micro-level outcomes. Instead, we need to model both levels jointly. Therefore, from the point of view of the demographic paradigms, we remain firmly within the realm of the multilevel analysis, only using different tools (simulations) to explore multiple layers of population processes at the same time. Conte et al. (2012, p. 342) suggested that:

...simulations must be accompanied by micro-macro-loop theories, i.e., theories of mechanisms at the individual level that affect the global behavior, and theories of loop-closing downward effects or second-order emergence.

A part of the strength of simulations lies in a potentially wide variety of ways to represent the same problems using a relatively simple set of techniques. However, there is a real danger that the models can be constructed in an arbitrary way, not linked to the observations of the properties of the population systems of interest, and thus become manifestations of Baconian ‘Idols’. This problem can be exacerbated if the models lack an explicit documentation of their construction and core assumptions when simulation results are presented⁸. In such situations, even models with well-grounded and well-justified assumptions may seem arbitrary.

Agent-based models are capable of analysing systems of interacting elements through computational modelling. A part of the appeal of such models is their capacity for explanatory power (see Burch 2003a, b; Silverman et al. 2011). As such, agent-based models by their very nature are intended to represent the import and impact of individual actions on the macro-level patterns observed in a complex system, and vice versa, showing a potential promise to transcend different levels of analysis. Such methods can further theoretical understanding of population processes (Burch 2003a; Chattoe 2003), and using these methods to break from the over-reliance of some micro-simulation models on empirical data at the expense of reasonable theoretical explanations and mechanisms⁹ (Silverman et al. 2011). As mentioned earlier, however, to take full advantage of this potential, we need to

⁸For a discussion of the ABM documentation standards, and the ODD framework (“Overview, Design concepts and Details”), see Grimm et al. (2006), as well as Chap. 9 in this volume.

⁹The problem here is not the empirical basis of such models – quite the contrary – but unrealistic mechanisms. Particularly problematic are Markovian assumptions of the lack of memory, where simulations are based on homogenous matrices of transition probabilities. Examples of micro-

look at how these different levels of aggregation interact, in order to better explain social facts. Simple aggregation of individual-level rules to generate and validate macroscopic patterns – as often implicitly done in existing agent-based models – is not sufficient (Conte et al. 2012).

In population sciences, there are many systems comprised of *interacting individuals, groups, or institutions* which are worthy of enquiry. Population sciences can become model-based by making those interactions between different levels in population systems an explicit object of interest. In so doing, our models would become capable of representing complex, interacting behaviours at various levels, and investigating the roles of different elements of population systems in shaping the observed demographic outcomes. Such models of multilevel interacting systems would have clear potential for contributing to theory-building within population sciences, and perhaps even social science more broadly.

Recent years saw an ever-increasing interest amongst population scientists in new modelling methodologies for complex social realities, many of these inspired by agent-based computational approaches (see Billari and Prskawetz 2003; Aparicio Diaz et al. 2011; Kniveton et al. 2011; Willekens 2012; Bijak et al. 2013; Silverman et al. 2013). The movements toward computational complexity have been matched by a shift coming from the other direction, as agent-based modellers have branched out into areas traditionally covered by statistical approaches in population science (see e.g. Axtell et al. 2002; Geard et al. 2013).

Of course, model-based approaches come with their own shortcomings – in particular, models attempting to represent the complexities of particular population systems are naturally dependent on sensible theories regarding these systems, and on their representation. However, such theories are not only many and varied, but can be notoriously difficult to formalise (Klüver et al. 2003), and validate¹⁰, especially in social science realms (see Moss and Edmonds 2005). Without such theories, it may be difficult to build an adequate model of the systems under study. A possible way forward from this conundrum is to reconnect to the classical research programme which promotes some sort of *functional-mechanistic analysis* (Franck 2002a); this will be discussed in the next sections of this chapter.

A clear strength of population science, and one of the keys to its success, is its applied character, responding to the direct needs of policy makers (Xie 2000; Morgan and Lynch 2001; Hirschman 2008). The methodological developments outlined above can only further this practical, utilitarian aspect of demographic enquiries. The Bayesian approach naturally allows formal statistical decision analysis, which can offer practical support to various decisions which require numerical input, for example for planning purposes (Alho and Spencer 2005; Bijak 2010). On the other hand, model-based approaches, especially coupled with statistical analysis, allows

simulation models that allow for heterogeneous transition patterns or mechanisms, e.g. of partnership formation, include SOCSIM (<http://lab.demog.berkeley.edu/socsim/>).

¹⁰After Franck (2002a), we interpret validation as a continuous process, rather than an achievable state.

the decision makers to trial a range of policy “levers” in a simulated environment. Such experimentation *in silico* would consist of generating coherent scenarios, where mechanistic rules governing the behaviour of *simulated individuals* would be coherent with the empirical patterns for *statistical individuals* observed through a scientific lens (Courgeau 2012).

Demography needs more simulations to be able to answer new research questions, but in order to suit the goals of the discipline, such simulations would need to be grounded in the observables, and the models would need to be built inductively (bottom-up), rather than starting from hypotheses and assumptions. To address this challenge for the future of demography and population sciences we propose a *model-based* research programme, firmly rooted in the wider *functional-mechanistic* approach. If agent-based models, as introduced above, are to belong to this programme, they need to be empirically based and scientifically rigorous.

As a part of this research programme, we posit that demography should investigate the *interactions* between various population systems and the functional *mechanisms* behind them. The interactions and mechanisms are best described by formal models based on data and theory-based rules, derived from observations of system properties by following the Baconian inductive method. This approach can augment the capabilities of the multilevel paradigm, whilst broadening the scope of scientific exploration in demography. In particular, it can enable population sciences to enhance the theoretical base of the discipline, whereby theories represent formal conceptual systems rather than necessarily empirical ones (Franck 2002a; Burch 2003b).


2.4 Conforming Model-Based Approaches to the Classical Scientific Programme Following the Baconian Inductive Method

How may the *model-based* approaches we propose for demography conform to the classical scientific research programme? This programme is at present generally ignored by social scientists – as well as overlooked by philosophers of science – because it has been distorted by the empiricist tradition in philosophy, where the empiricism promoted by David Hume and John Stuart Mill has substituted for the classical empiricism. Francis Bacon’s programme was shared by the other pioneers of modern science, Galileo, Descartes, Newton, Huyghens, Graunt, and others. This programme tacitly continues to guide research in the natural sciences today just as in the past, yet in the social sciences it is often abandoned.

To start with an illustrative example, consider the famous model of a neuron (McCulloch and Pitts 1943), the ancestor of the now-ubiquitous artificial neural networks, which is shown in Fig. 2.1.

The model in Fig. 2.1 represents one neuron. Yet what is represented? Not the soma, nor the axon, nor the dendrites, nor the gene nuclei, nor the membrane, nor the shape of the neuron, nor the way that the various parts of the neuron work together.

Fig. 2.1 The McCulloch and Pitts (1943) model of a single neuron (Source: Adapted from Franck (2002a, p. 143), with permission of Springer)



Starting from the observation of some main properties of the neuron, McCulloch and Pitts tried to represent its *functional* architecture, *without which these properties could not come about as they do*.

Five functions were identified: receiving the stimuli x_1, x_2, \dots, x_N ; weighting them by synaptic coefficients w_1, w_2, \dots, w_N ; calculating the sum of weighted stimuli received (p); fixing a threshold of stimulation (σ) below which transmission does not occur; and finally computing the exit signal s . These functions are arranged in a specific order: the weighting of stimuli must precede the calculation of the weighted sum, and so on. Thus, more precisely, the McCulloch and Pitts model represents *the functional structure of the process generating the observed outcomes* of the neuron. Note that such a model ignores the combination of factors or *causes* that fulfil these functions – it is wholly conceptual.

Let us now have a look at *reverse engineering*¹¹. Reverse engineering follows a similar path: inducing the design of a device from its end products. Its original aim is to make a new device that does the same thing as the device studied. At first glance, modelling the *functional* structure of a device from its products in order to make similar products through new procedures is very far from what McCulloch and Pitts achieved, since they were not driven to manufacture some artificial neuron. Yet they have followed the same *method*: they inferred from the neuron’s properties the structure of functions without which these properties could not be as they are (Franck 2002b).

This method is the one which has been conceived by the classical programme of scientific research: from the sustained observation of some property of nature (light, heat, motion . . .) we try to infer – to *induce*¹² – the *functional* structure – in classical terms the *axiom, form, principle, or law* – which rules the process generating this property. We may, at present, qualify this method as *functional-mechanistic* to

¹¹Reverse engineering denotes today diverse research practices varying with the areas of application. We refer to its initial sense.

¹²Bacon’s *induction* is regularly confounded with induction by philosophers in its usual sense of generalisation. Bacon wrote: “In establishing axioms, another form of induction must be devised than has hitherto been employed, and it must be used for proving and discovering not first principles (as they are called) only, but also the lesser axioms, and the middle, and indeed all. For the induction which proceeds by simple enumeration is childish; its conclusions are precarious and exposed to peril from a contradictory instance; and it generally decides on too small a number of facts, and on those only which are at hand” (Bacon 1620; aphorism 105).

underline that it aims to model *the structure of functions* that rules the *mechanism* – the process – generating some property of nature.

For social properties, the method involves modelling the structure of the social functions (the ‘first’ one, the ‘lesser’ and the ‘middle’ in Bacon’s terms) that rule the social process generating these properties, and without which these social properties could not become as they are. For example, regarding variations of population size and structure, demographers uncovered the ‘first’ principle of the generating processes, namely some combination – which remains to be discovered – of three functions: *fertility*, *mortality*, and *migration*.

The ‘law’ of *supply* and *demand*, as another example, is the ‘first’ structure of functions which was inferred (*induced*) by Adam Smith from the observation of markets: it rules the process of social exchanges generating the market. Karl Marx inferred the general structure of *functions* ruling the process that generates industrial production from a thorough historical study of the technical and social organisation: this ‘first’ principle consists of *separating labour and capital*. Finally, Durkheim inferred the *integration theory* from a sustained statistical analysis of the differences in suicide rates between several social milieus: the social process which generates suicides, whichever their causes, is ruled by the *integration* of the individual agents. The application of the classical programme led to these prominent *theoretical* results at the height of *social sciences*.

Next, the functional structure governing the process generating some social property, once established as well as possible, may guide us in identifying and modelling the social factors which – in some singular, historical situation – have contributed to that process. We may restrict our causal investigation to those variables which plausibly contributed to the combination of functions required for generating the property under study. For example, what social factors (events, agents’ behaviour, etc.) led to a weakening of the *integration* of people in some social milieu, and contributed to the increase of suicide? Another example: when we investigate the ups and downs of the market, we no longer ought to interrogate every plausible factor influencing these variations; instead, it may suffice to investigate and model the factors implied by supply and demand. In demography it is the functions of fertility, mortality and migration which actually delimit its parameter space and channel the *empirical* investigation of demographic properties.

Against this background, we propose that the model-based research programme should proceed in accordance with the classical inductive programme, which we qualify as functional-mechanistic. Model building should start with a collection of all relevant empirical information about the social property under study. This would serve as the basis with which to infer the formal functional structure of the social property in question. Once the structure is modelled it can serve to guide the modelling – also simulation modelling – of the interactions between the systems of individuals, groups and institutions, combining the bottom-up and top-down relationships, and feedbacks between them (Franck 2002a).

The key stages of the inductive functional-mechanistic approach are shown in Fig. 2.2. The solid arrows denote the four main stages of the process. Their implementation leads to the execution and analysis of a computational model

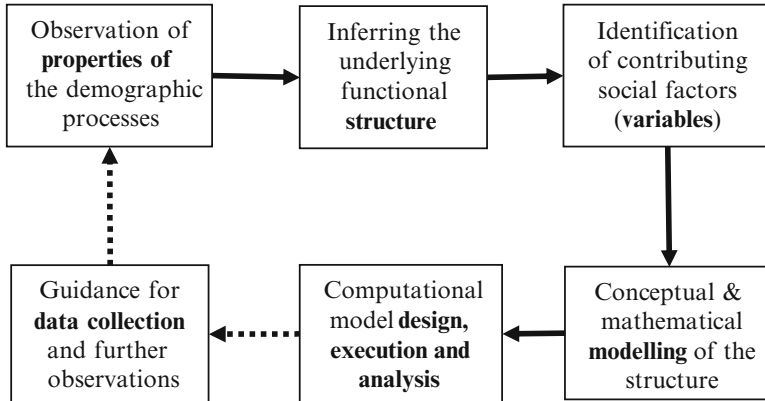


Fig. 2.2 Key stages of the inductive functional-mechanistic approach to model-based demography

designed on the basis of a functional structure of the population processes under study, and derived from empirical observations. The dotted lines depict a possible feedback: the model results can guide the process of further data collection on particular aspects of the process that have not been included in the model. Obtaining more information would enable revisiting the inferred functional structure, as well as the resultant conceptual, mathematical, and computational models. Thus, the proposed approach could be seen as iterative, with successive modelling iterations enabling the researchers to identify gaps in existing knowledge and to fill these gaps by conducting further observations of the processes of interest.

The proposed approach is in line with the suggestion of Conte et al. (2012, p. 342) that "... data can be used to check and validate the results of simulation models and socio-economic theories, but a further step in using them is to take them into account already at the modelling stage." A careful experimental design becomes therefore a crucial part of the modelling process, and statistical methods, such as those discussed earlier in this chapter, need to become embedded in the model construction process, rather than being used only as a validation tool.

2.5 Towards a Research Agenda for Demography

The resulting research agenda we would like to propose for demography is based on three key pillars: (1) adherence to the classical programme of scientific enquiry; (2) enhancement of the ways in which demographic phenomena are measured and interpreted; and (3) the use of formal models, based on the functional-mechanistic principles, as fully-fledged tools of population enquiries.

According to several authors demography should become *interdisciplinary* in order to compensate its perceived shortcomings (e.g. Petit and Charbit 2012).

Demography should borrow *information, methods, and theories* from other social sciences. Although we approve a cautious recourse, when needed, to other disciplines, we have to underline that the solution to the weaknesses of current demographic practices cannot be found in theories and methods borrowed from other social sciences, since the last share similar weaknesses (e.g. Moss and Edmonds 2005). All of these are confronted, just as in demography, with the *complexity* of society; they suffer from *uncertainty* in collecting and treating information often more than in population sciences; and their *theories* are hard to settle. The history of social sciences since the mid of the twentieth century teaches us that many innovative ‘theories’ had a generally short life, or at best remain heavily controversial. Moreover, the flowering of such ‘theories’ nourishes the proliferation of heterogeneous explanations and seem to increase the *complexity* and the *uncertainty* which undermines social sciences.

On the contrary, the first of the pillars of the proposed research agenda – the classical programme of scientific research – helps overcome the *complexity* of society, it reduces the *uncertainty* of the models we are building and of the explanations we are advancing, it establishes the *theoretical* component of research, and it discloses the way to *generalize* social models, something which is reputed to be an inaccessible goal in the social area. This is one of the reasons why we recommend applying this method in particular in demography and population sciences.

The belief that knowledge is something like a copy or an image is widespread. The classical programme conveys a different concept of science: scientific research is not intended to improve or to extend our image of reality. Instead, scientific research consists of discovering the principles governing the processes that generate some properties of nature or of society. We need to collect the best information on some *property* of nature – not about nature as a whole – in order to discover the principles governing the process that generates this property. The same applies for the social sciences: we need to collect the best information about some social *property* of human populations, not about human populations as a whole. Moreover, it is not merely information about this social property that will reinforce our scientific knowledge. Collecting information must be augmented by research on the principles – i.e. the combination (or *structure*) of *functions* – commanding the process generating this property. When selected in this way, the required informative intake varies with the property under study and it is drastically reduced, restricting at the same time the complexity of the task. The classical research programme also restricts the *theoretical* approach to some social property to the modelling of *the structure of functions necessary* to generate this property.

This approach provides a major criterion for selecting information, by restricting our causal investigation to those variables which plausibly contributed to performing the combination of functions required for generating the property under study. This is yet another way to overcome complexity that is delivered by the classical programme of science.

Assumptions are an important source of uncertainty and nourish the proliferation of explanations in the social sciences¹³. The classical research programme recommends setting aside any assumptions in the guiding of scientific investigation (Baconian ‘Idols’). Thus it eliminates the root of any uncertainty arising from assumptions. Managing research without assumptions seems difficult – but can be done if we resist the urge to make hypotheses. Newton argued that “Hypotheses are not to be regarded in experimental Philosophy” (after Ducheyne 2005, p. 124). This way of thinking is not in fashion today in the social sciences, yet it is commonplace in the natural sciences.

The classical research programme substitutes *induction* for the *hypothetical-deductive* approach. As we have seen, *inductive research* in its classical sense consists in inferring, from the sustained observation of and experimentation on some property of nature or population, the *functional* structure – in classical terms the *axiom, form, principle, or law* – which rules the process generating this property, and *without which* this property *could not* come about as it does. This criterion of *necessity* which guides the *inductive* investigation of the *functional structure* ruling the social processes which generate some social property is the best guarantor of the relevance of some *theory*, be it in the natural or in the social sciences.

Besides reducing by far the nagging worries of *complexity* and *uncertainty* characterising the *social* sciences today, the classical programme of science provides another huge advantage. By focusing on the process – or ‘mechanism’ – generating some natural or social property, the functional structure is treated independently of the causal structure. Modelling each of these structures separately allows us to disclose the way to *generalise* social models. Causal structures may never be generalized since populations are diverse and changing; causal structures are at best relatively constant. But a functional structure may be generalized in the sense that, *whenever the same property occurs*, the functional structure of its generative process ought to operate, insofar it has been established that this structure is really required for the property to appear¹⁴. This is the core sense of the *universality* and of the *necessity* of natural *laws*, but it has been regrettably distorted by the Humean empiricist tradition. For the founders of modern science the term *law* was simply a metaphorical synonym of *principle*¹⁵. Thus laws are fully attainable by the social sciences, just as by the natural sciences, insofar we are willing to return to the classical concept of scientific laws.

The second pillar of the tentative research agenda we wish to advance comprises a question: how can we make better use of the measures achieved in demography?

¹³Formulating and testing hypotheses is not wrong, in our opinion, as long as it is based on empirical observations. However, throughout the present chapter we plead for abandoning the hypothetical-deductive approach and for substituting it with the classical induction.

¹⁴The property itself may not be generalized, of course.

¹⁵The principles are traditionally named *theories*; this tradition goes back to Plato’s *theoria*, and reserves to the term *theory* the restricted sense of a corpus of principles. This is far from its present use describing as a ‘theory’ every sort of conceptual hypothesis, or model, or explanatory ‘mechanism’.

One of the main tasks of demography is to measure human populations: their size, density, rate of change, composition, various distributions, as well as the possible causes and consequences of changes in these factors¹⁶. In order to achieve this, demography rests largely on statistical analysis. Yet, *measuring* provokes an increasing dissatisfaction today. This is due, in our opinion, to a distorted view of measurement and quantification. *Measuring* population properties is *judging*, by way of comparisons, the amplitude of these properties regarding their *potentialities*. Measuring also guides the *induction* – in the sense of the classical research programme – of the principles governing these properties.

The term “*potentialities*” refers to what possible effects something might generate in certain circumstances. These *potentialities* are what we have in mind, explicitly or not, when we are measuring – or *judging* – some social property¹⁷. Measuring may not be confounded with – nor reduced to – the mathematical, statistical or other means by which measuring is carried out (for example censuses, surveys or vital registration systems; see Courgeau 2013).

In essence, we ought to multiply the measures of the social properties under study – all sorts of measures which are adequate – and to improve the *quality* of our measures in order to reinforce the quality of our judgments about their *potentialities*. This recommendation is exactly the opposite of what was proposed for the future of demography in recent years by some demographers (e.g. Tabutin 2007; Charbit and Petit 2011; Petit and Charbit 2012, and others): they wish to reduce the importance of *measuring* in demography and to increase our confidence in judgments – assumptions – conceived without measures in other disciplines, and somewhat abusively called *theories*.

Now we reach the third pillar of our tentative research agenda: to promote the *model-based* work programme, based on the functional-mechanistic approach outlined in this chapter. This approach carries with it substantial promise: it complements the four extant paradigms while incorporating insights gained from model-based science. Besides, as we see in other areas of model-based science, the deployment of this kind of approach likely will influence future data collection in demography and other population sciences, not only from surveys and other traditional sources, but also controlled experiments (Conte et al. 2012).

Model-based approaches provide us with the means to expand the range of benefits already provided by multilevel modelling. We gain deeper insights into the interactions between various population systems, and we also gain the capacity to explore the parameter space of the simulations by generating “what-if” scenarios. Simulation parameters – once they result from the functional-mechanistic approach – govern the way in which the complex, interacting social processes in

¹⁶See for example the following definition of demography (IUSSP 1982): “the scientific study of human populations primarily with respect to their size, their structure and their development; it takes into account the quantitative aspects of their general characteristics”.

¹⁷Plato, who was familiar with the concept of *number* developed by the Pythagoreans, developed at length the idea that *measuring is judging*, and that we ought to recourse to measures in order to act wisely in politics as well as in private life (see Bassu 2009, 2011).

the model work, and therefore exploring the parameter space enables us to investigate numerous such scenarios, which could represent policy changes, individual behavioural changes, societal-level changes, and similar (Silverman et al. 2013). Given the construction of these simulations, running them under varied scenarios can illustrate the unforeseen, non-linear impact of changes to these complex processes. This *scenario generation* capability, when coupled with uncertainty quantification, allows us to extend the utility and policy relevance of empirically-grounded population models beyond what is accommodated by the traditional approaches. In addition to addressing the ecological and atomistic fallacies, which is already the case in the current multilevel paradigm, we could now analyse different layers of interactions between population systems.

Such approaches, relying as they do upon inference about systems and interactions between them, are also well-suited to integrating both quantitative and qualitative data into the same simulations, as mentioned before. For example, qualitative information can be gathered from individuals within the population under study, as a means of gaining understanding regarding individual behaviours, intentions, and goals, and these can inform the behavioural rules in the simulated population. Further, qualitative data can even be used to guide the construction or modification of the model itself (e.g., Polhill et al. 2010).

In this chapter, we have discussed what we believe are the key elements of model-based approaches – such as their inductive character – that would be necessary for them to become a real addition to the toolbox of population sciences. If the future demography is to examine complex, multilevel interactions of different elements of population systems seriously, computational approaches are the methodology of choice. However, the models constructed would need to conform to the rigours of scientific enquiry, rather than being based on arbitrary assumptions which often lack empirical basis. The model-based work programme, rooted in the functional-mechanistic approach, offers a general analytical framework to guide this process. Besides, more attention needs to be paid to the role of different levels of analysis, and interactions between them. If this is done correctly, the multilevel paradigm will gain very powerful analytical tools to study new research questions, related to the behaviour of complex population systems.

The next step in developing model-based demographic approaches must consist of proposing some concrete solutions, analytical formalisms and practical guidelines for the modellers. Although this topic remains beyond the scope of the current chapter, in the literature there are already some promising suggestions in that regard. For example, Casini et al. (2011) have proposed using recursive Bayesian networks as an analytical formalism for building “models for prediction, explanation, and control”, which are capable of describing functional mechanisms and causal relations, and of analysing uncertainty in coherent, probabilistic terms. In practice, the process of model-building can be iterative, as shown in Fig. 2.2: we could start with a first approximation of a model that reproduces some well-established qualitative features of the modelled phenomenon (‘stylised facts’), but should not stop there: the model could then be refined by including increasingly more data as they become available. These propositions are clearly worthy of investigating in the demographic context.

Of course, it is unrealistic to expect that every piece of model-based demographic research should contain all the elements discussed above. However, as future studies progress – and as populations under study continue to shift following ever-changing and interacting social processes – model-based approaches to demography will bring about further opportunities for constructing and verifying the models. In this respect the linkage between empirical data on population structures and modelling the social mechanisms and interactions at the root of these structures becomes ever more important – and perhaps more powerful.

Acknowledgments JB and ES acknowledge the Engineering and Physical Sciences Research Council (EPSRC) grant EP/H021698/1 “Care Life Cycle”. We thank Frans Willekens and Anna Klabunde for discussions and to the two anonymous reviewers for helpful suggestions. All the views and interpretations are those of the authors and should not be attributed to any institution with which they are affiliated. All the errors remain exclusively ours.

References

- Aalen, O. O. (1975). *Statistical inference for a family of counting processes*. PhD thesis, University of California, Berkeley.
- Ahlburg, D. A. (1995). Simple versus complex models: Evaluation, accuracy and combining. *Mathematical Population Studies*, 5(3), 281–290.
- Alho, J. M., & Spencer, B. D. (2005). *Statistical demography and forecasting*. Berlin/Heidelberg: Springer.
- Alkema, L., Raftery, A. E., & Clark, S. J. (2007). Probabilistic projections of HIV prevalence using Bayesian melding. *Annals of Applied Statistics*, 1(1), 229–248.
- Aparicio Diaz, B., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography*, 48(2), 559–579.
- Axtell, R., Epstein, J., Dean, J., Gumerman, G., Swedlund, A., Harburger, J., Chakravarty, S., Hammond, R., Parker, J., & Parker, M. (2002). Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciences of the United States of America*, 99(suppl. 3), 7275–7279.
- Bacon, F. (1620). *Novum Organum*. London: J. Bill. English translation: Spedding, J., Ellis, R. L., & Heath, D. D. (1863). *The works* (Vol. VIII). Boston: Taggard and Thompson.
- Bassu, S. (2009). Metretique, éthique et politique: le Protagoras et le Politique de Platon. *Dissertatio*, 29, 85–114.
- Bassu, S. (2011). Ordre et mesure, kosmos et metron de la pensée archaïque à la philosophie platonicienne. In S. Alexandre & E. Rogan (Eds.), *Actes du colloque “Ordres et désordres”*, Université Paris 1 and Université Paris Ouest, Nanterre-La Défense, 4–5 June 2010. Available via: Zetesis, vol. 2 [Online], <https://f.hypotheses.org/wp-content/blogs.dir/3211/files/2015/04/2Bassu.pdf>.
- Bayes, T. R. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53, 370–418.
- Bijak, J. (2010). *Forecasting international migration in Europe: A Bayesian view* (Springer Series on Demographic Methods and Population Analysis, Vol. 24). Dordrecht: Springer.
- Bijak, J., & Bryant, J. (2016). Bayesian demography 250 years after Bayes. *Population Studies*, 70(1), 1–19.
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the United Kingdom with Gaussian process emulators. *Demographic Research*, 29(27), 729–766.

- Billari, F., & Prskawetz, A. (Eds.). (2003). *Agent-based computational demography. Using simulation to improve our understanding of demographic behaviour*. New York: Physica-Verlag.
- Blayo, C. (1995). La condition d'homogeneity en analyse demographique et en analyse statistique des biographies. *Population*, 50(6), 1501–1518.
- Boudon, R. (1977). *Effet pervers et ordre social*. Paris: Presses Universitaires de France.
- Brenner, T., & Werker, C. (2007). A taxonomy of inference in simulation models. *Computational Economics*, 30(3), 227–244.
- Bullock, S., & Silverman, E. (2008). *Levins and the legitimacy of artificial worlds*. A Cross-Disciplinary Workshop “Epistemological Perspectives on Simulation”, Lisbon, 2–3 October 2008.
- Burch, T. (2003a). Data, models, theory and reality: The structure of demographic knowledge. In F. Billari & A. Prskawetz (Eds.), *Agent-based computational demography. Using simulation to improve our understanding of demographic behaviour* (pp. 19–40). Heidelberg/New York: Physica-Verlag.
- Burch, T. (2003b). Demography in a new key: A theory of population theory. *Demographic Research*, 9(11), 263–284.
- Casini, L., Illari, P. M., Russo, F., & Williamson, J. (2011). Models for prediction, explanation and control: Recursive Bayesian networks. *Theoria*, 26(1), 5–33.
- Charbit, Y., & Petit, V. (2011). Towards a comprehensive demography: Rethinking the research agenda on change and response. *Population and Development Review*, 37(2), 219–239.
- Chattoe, E. (2003). The role of agent-based models in demographic explanation. In F. Billari & A. Prskawetz (Eds.), *Agent-based computational demography. Using simulation to improve our understanding of demographic behaviour* (pp. 41–54). Heidelberg/New York: Physica-Verlag.
- Clark, S. J., Thomas, J. R., & Bao, L. (2012). Estimates of age-specific reductions in HIV prevalence in Uganda: Bayesian melding estimation and probabilistic population forecast with an HIV-enabled cohort component projection model. *Demographic Research*, 27(26), 743–774.
- Conte, R., Gilbert, N., Bonelli, G., Cioffi-Revilla, C., Deffuant, G., Kertesz, Loreto, V., Moat, S., Nadal, J.-P., Sanchez, A., Nowak, A., Flache, A., San Miguel, M., & Helbing, D. (2012). Manifesto of computational social science. *European Physical Journal Special Topics*, 214(1), 325–346.
- Courgeau, D. (2007). *Multilevel synthesis. From the group to the individual*. Dordrecht: Springer.
- Courgeau, D. (2012). *Probability and social science. Methodological relationships between the two approaches* (Methodos Series 10). Dordrecht: Springer.
- Courgeau, D. (2013). La mesure dans les sciences de la population. *Cahiers Philosophiques*, 135(4), 51–74.
- Courgeau, D., & Franck, R. (2007). Demography, a fully formed science or a science in the making? An outline programme, *Population-E*, 62 (1), pp. 39–45. (La démographie, science constituée ou en voie de constitution? Esquisse d'un programme. *Population*, 62(1), 39–45).
- Courgeau, D., & Lelièvre, E. (1992). *Event history analysis in demography*. Oxford: Clarendon.
- Courgeau, D., Bijak, J., Franck, R., & Silverman, E. (2014). Are the four Baconian Idols still alive in demography? *Revue Quetelet/Quetelet Journal*, 2(2), 31–59.
- Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In M. Bedau, J. McCaskill, N. Packard, & S. Rasmussen (Eds.), *Proceedings of the 7th international conference on artificial life* (pp. 497–506). Cambridge, MA: MIT Press.
- Doob, J. L. (1953). *Stochastic processes*. New York/Chichester: Wiley.
- Ducheyne, S. (2005). Bacon's idea and Newton's practice of induction. *Philosophica*, 76, 115–128.
- Durkheim, E. (1897). *Le suicide*. Paris: Alcan.
- Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), article 12. <http://jasss.soc.surrey.ac.uk/11/4/12.html>.
- Franck, R. (Ed.). (2002a). *The explanatory power of models. Bridging the gap between empirical and theoretical research in the social sciences* (Methodos series, Vol. 1). Boston/Dordrecht/London: Kluwer Academic Publishers.

- Franck, R. (2002b). Computer simulation and the reverse engineering method. Conclusions of part II. In R. Franck (Ed.), *The explanatory power of models* (Methodos series, Vol. 1, pp. 141–146). Dordrecht/Boston/London: Kluwer Academic Publishers.
- Geard, N., McCaw, J. M., Dorin, A., Korb, K. B., & McVernon, J. (2013). Synthetic population dynamics: A model of household demography. *Journal of Artificial Societies and Social Simulation*, 16(1), article 8. <http://jasss.surrey.ac.uk/16/1/8.html>.
- Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy*, 21(5), 725–740.
- Goldstein, H. (1987). *Multilevel models in educational and social research*. London: Arnold.
- Graunt, J. (1662). *Natural and political observations mentioned in a following index, and made upon the bills of mortality*. London: Tho. Roycroft.
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S., Huse, G., Huth, A., Jepsen, J. U., Jørgensen, C., Mooij, W. M., Müller, B., Pe'er, G., Piou, C., Railsback, S. F., Robbins, A. M., Robbins, M. M., Rossmannith, E., Rügen, N., Strand, E., Souissi, S., Stillman, R. A., Vabø, R., Visser, U., & DeAngelis, D. L. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2), 115–126.
- Henry, L. (1959). D'un problème fondamental de l'analyse démographique. *Population*, 14(1), 9–32.
- Hirschman, C. (2008). The future of demography. *Asian Population Studies*, 4(3), 233–234.
- Holland, J. H. (1995). *Hidden order*. Reading: Addison-Wesley.
- Huneman, P. (2014). Mapping an expanding territory: Computer simulations in evolutionary biology. *History and Philosophy of the Life Sciences*, 36(1), 60–89.
- Huyghens, C. (1657). *De ratiociniis in ludo aleae*. Leyde: Elsevier.
- Ibrahim, J. G., Chen, M.-H., & Sinha, D. (2001). *Bayesian survival analysis*. New York: Springer.
- IUSSP [International Union for the Scientific Study of the Populations]. (1982). *Multilingual demographic dictionary* (2nd ed.). Liège: Ordina.
- Kennedy, M., & O'Hagan, T. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society, Series B*, 63(3), 425–464.
- Keyfitz, N. (1971). Models. *Demography*, 8(4), 571–580.
- Klüver, J., Stoica, C., & Schmidt, J. (2003). Formal models, social theory and computer simulations: Some methodical reflections. *Journal of Artificial Societies and Social Simulation*, 6(2), article 8. <http://jasss.surrey.ac.uk/6/2/8.html>.
- Kniveton, D., Smith, C., & Wood, S. (2011). Agent-based model simulations of future changes in migration flows for Burkina Faso. *Global Environmental Change*, 21(Suppl. 1), S34–S40.
- Kuhn, T. (1962). *The structure of scientific revolutions*. Chicago/London: The University of Chicago Press.
- Laplace, P. S. (1774). Mémoire sur la probabilité des causes par les événements. *Mémoires de l'Académie Royale des Sciences de Paris, Tome, VI*, 621–656.
- Laplace, P. S. (1812). *Théorie analytique des Probabilités* (Vol. 2). Paris: Courcier Imprimeur.
- Levins, R. (1966). The strategy of model building in population biology. *American Scientist*, 54(4), 421–431.
- Lutz, W. (2012). Demographic metabolism: A predictive theory of socio-economic change. *Population and Development Review*, 38(Supplement), 283–301.
- Mannheim, K. (1928). Das Problem der Generationen. *Kölner Vierteljahreshefte für Soziologie*, 7(2), 309–330.
- Mason, W. M., Wong, G. W., & Entwistle, B. (1983). Contextual analysis through the multilevel linear model. In S. Leinhardt (Ed.), *Sociological methodology 1983–1984* (pp. 72–103). San Francisco: Jossey-Bass.
- McCulloch, W. S., & Pitts, W. H. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5(4), 115–133.
- Mill, J. S. (1843). *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence, and the methods of scientific investigation* (Vol. I). London: Harrison.

- Morgan, S. P., & Lynch, S. M. (2001). Success and future of demography. The role of data and methods. *Annals of the New York Academy of Sciences*, 954, 35–51.
- Moss, S., & Edmonds, B. (2005). Towards good social science. *Journal of Artificial Societies and Social Simulation*, 8(4), article 13. <http://jasss.soc.surrey.ac.uk/8/4/13.html>.
- NRC [National Research Council]. (2000). *Beyond six billion: Forecasting the world's population*. Washington, DC: National Academies Press.
- Oakley, J., & O'Hagan, A. (2002). Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*, 89(4), 769–784.
- Pascal, B. (1665). *Traité du triangle arithmétique, avec quelques autres traités sur le même sujet*. Paris: Guillaume Desprez.
- Petit, V., & Charbit, Y. (2012). The French school of demography: Contextualising demographic analysis. *Population and Development Review*, 38(supplement), 322–333.
- Petty, W. (1690). *Political arithmetick*. London: Robert Clavel & Hen. Mortlock.
- Polhill, J. G., Sutherland, L.-A., & Gotts, N. M. (2010). Using qualitative evidence to enhance an agent-based modelling system for studying land use change. *Journal of Artificial Societies and Social Simulation*, 13(2), art. 10. <http://jasss.soc.surrey.ac.uk/13/2/10.html>.
- Poole, D., & Raftery, A. E. (2000). Inference for deterministic simulation models: The Bayesian melding approach. *Journal of the American Statistical Association*, 95(452), 1244–1255.
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology*, 25, 111–163.
- Raftery, A. E., Li, N., Ševčíková, H., Gerland, P., & Heilig, G. K. (2012). Bayesian probabilistic population projections for all countries. *Proceedings of the National Academy of Sciences*, 109, 13915–13921.
- Rogers, A. (1975). *Introduction to multiregional mathematical demography*. New York: Wiley.
- Ryder, N. B. (1951). *The cohort approach. Essays in the measurement of temporal variations in demographic behaviour*. PhD thesis, Princeton University, New York.
- Silverman, E., & Bryden, J. (2007). From artificial societies to new social science theory. In F. Almeida e Costa, L. M. Rocha, E. Costa, I. Harvey, & A. Coutinho (Eds.), *Advances in artificial life, 9th European conference, ECAL 2007 proceedings* (pp. 645–654). Berlin/Heidelberg: Springer.
- Silverman, E., Bijak, J., & Noble, J. (2011). Feeding the beast: Can computational demographic models free us from the tyranny of data? In T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, M. Dorigo, & R. Doursat (Eds.), *Advances in artificial life, ECAL 2011: Proceedings of the eleventh European conference on the synthesis and simulation of living systems* (pp. 747–754). Cambridge, MA: MIT Press.
- Silverman, E., Bijak, J., Hilton, J., Cao, V., & Noble, J. (2013). When demography met social simulation: A tale of two modelling approaches. *Journal of Artificial Societies and Social Simulation*, 16(4), article 9. <http://jasss.soc.surrey.ac.uk/16/4/9.html>.
- Smith, S. K. (1997). Further thoughts on simplicity and complexity in population projection models. *International Journal of Forecasting*, 13(4), 557–565.
- Tabutin, D. (2007). Whither demography? Strengths and weaknesses of the discipline over fifty years of change. Followed by a debate on the future of the discipline, by G. Caselli & V. Egidi, D. Courgeau & R. Franck, J. Hobcraft, & J. Hoem. *Population-E*, 62(1), 13–56.
- Thagard, P. (1993). *Computational philosophy of science*. Cambridge, MA: MIT Press.
- Whelpton, P. (1949). Cohort analysis of fertility. *American Sociological Review*, 14(6), 735–749.
- Willekens, F. (2005). Biographic forecasting: Bridging the micro-macro gap in population forecasting. *New Zealand Population Review*, 31(1), 77–124.
- Willekens, F. (2012). Migration: A perspective from complexity science. Paper for the Complexity Science for the Real World workshop on migration, Chilworth. 16 Feb 2012.
- Xie, Y. (2000). Demography: Past, present and future. *Journal of the American Statistical Association*, 95(450), 670–673.

Chapter 3

The Role of Social Interactions in Demography: An Agent-Based Modelling Approach

Alexia Prskawetz

3.1 Introduction

As we argued 10 years ago (Billari and Prskawetz 2003), we may still conclude that demographers have not been very active in the field of research that uses the computational approach to study human behaviour. This is all the more astonishing since demography offers itself quite naturally for such an approach. Demography looks at human behaviour at the individual level and how this behaviour evolves over the individual life cycle and is shaped by the socio-economic environment.

During the last years consensus has been reached, and could be supported by empirical evidence, that individual behaviour cannot be explained and understood in isolation from the social network one is linked to (e.g. Åberg 2003; Montgomery and Casterline 1996). These networks may consist of family members, friends and other peer groups which will have an impact through social learning and social influence on each other. However, the formalisation of such network effects to explain individual demographic behaviour lags behind the empirical evidence or is often simplified in terms of macro-level diffusion mechanisms that do not allow understanding the mechanisms of social network effects from the bottom up. Agent-based models allow to integrate such network effects into models of individual demographic decision processes and to build up the macro-level demographic patterns (e.g. aggregate fertility rates, marriage rates, etc.) from the bottom up.

A. Prskawetz (✉)

Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology,
Wiedner Hauptstrasse 8-10, E105-3, Vienna, Austria

Wittgenstein Centre for Demography and Global Human Capital (IIASA, VID/ÖAW, WIC),
Welthandelsplatz 2, 1020 Vienna, Austria
e-mail: afp@econ.tuwien.ac.at

© Springer International Publishing Switzerland 2017

A. Grow, J. Van Bavel (eds.), *Agent-Based Modelling in Population Studies*,
The Springer Series on Demographic Methods and Population Analysis 41,
DOI 10.1007/978-3-319-32283-4_3

53

Before we present, in the next section, three examples of our own work relating to social interaction and demographic behaviour, we briefly summarise the main tasks and requisites of an agent-based model (ABM).

We start with a quote from Axelrod (1997, p. 4) “Whereas the purpose of induction is to find patterns in data and that of deduction is to find consequences of assumptions, the purpose of agent-based modeling is to aid intuition.” ABMs constitute computational laboratories that help our understanding of processes underlying the empirical observation of demographic behaviour. As such models are based on individual agents, they are called agent-based models. Prediction is not the purpose of these models, but the emphasis is rather on explanation.

It is important to confine the task of agent-based modelling. The aim is not to understand why specific rules are applied by individuals but to presuppose specific behavioural rules and study whether these rules at the micro level can explain the macro-level pattern we empirically observe. For example, the famous model of Schelling (1971) (see also Schelling 1978) aims to explain the spatial segregation pattern of individuals observed at the macro level based on specific decision rules at the individual level.

As outlined in Axelrod and Tesfatsion (2006), ABM focuses also on how people interact and not just how they behave. The interaction of agents leads to emergent properties that could not be deduced from simple aggregation of individual behaviour. ABMs help in “testing, refinement and extension of existing theories that have proved to be difficult to formulate and evaluate using standard statistical and mathematical tools” (Axelrod and Tesfatsion 2006, p.1651). But also the individual heterogeneity is an important aspect of an agent-based model as well as the possibility that individuals can adapt and learn through time (Gilbert 2008). This deviates clearly from the representative-agent model that assumes static or dynamic predefined rules.

The fact that an ABM needs to be implemented as a computer program requires precision. On the other hand, the mathematical tractability is less of a limitation compared to formalised theoretical constructions. Indeed, there is often a wide gap between theory and techniques in demography and ABMs may help to close this gap. For example, demographers may present interesting theories of behaviour and good statistical models but frequently the link is missing. Hence, many statistical models suffer from an insufficient theoretical basis. Moreover, ABMs may be regarded as a tool to test theories for which data are not easily obtained or not available at all. Examples are subjective aspects of demography such as values, norms, psychological aspects, cognition or emotions where we often lack concrete data but argue in theories about their importance for explaining human behaviour.

The three examples we present in Sect. 3.2 should convey main properties of an ABM as outlined above. Based on rules at the individual level we aim to explain the macro level demographic pattern of marriage (Sect. 3.2.1) and fertility (Sect. 3.2.2). The case of ABMs as computational laboratories is best presented by our third example (Sect. 3.2.3) where we study the role of family policies for fertility. Common to all three examples is the fact that individuals are heterogeneous and interact within social networks.

3.2 Three Examples

Based on three of our papers we present how agent-based models can be applied to investigate the role of social interactions and social learning to explain macro-demographic phenomena like the age-at-marriage curve (Billari et al. 2007), age-specific fertility rates (Aparicio Diaz et al. 2011) and the role of family policies for fertility (Fent et al. 2013). In particular we focus on the various steps that need to be followed when building up an agent-based model. These include the discussion of the macro-demographic phenomena to be explained and the underlying micro-demographic mechanisms, the implementation of behavioural rules of the model in a mathematical representation, the setup of the simulations and finally a verification of the simulation results with the macro-demographic phenomena to be explained.

3.2.1 *The Wedding Ring: Mate Search and Marriage*

3.2.1.1 Theory and Assumptions

The marriage market constitutes an intuitive case study to apply agent-based modelling as it is based on individual agents that interact and may follow specific rules how to search for partners. The aim of our model is to explain the typical shape of the aggregate age pattern of marriage as it emerges from the micro dynamics of individual agents. The benchmark against which we test our model is the shape of the age-at-marriage hazard function (cf. Billari et al. 2007, Fig. 1) which has a skewed unimodal shape where the rise of age-specific probabilities is faster than its decrease.

While demographers have mainly applied statistical and mathematical models at the macro level to explain and model the age pattern of marriage, psychologists and economists have studied and modelled the process of partner search at the micro level. Applying agent-based models allows us to combine both approaches. Such models account for the macro-level marriage pattern starting from plausible micro-level assumptions and allow for the interaction between potential partners.

To model the social diffusion of marriage at the micro level we assume that each agent is embedded in a social network. Members of the agent's social network (relevant others) who are already married may influence the agent's willingness to marry, and the chance of actually marrying will depend on the availability of partners. These mechanisms are also underlying the macro-level diffusion marriage model by Hernes (1972). Marriage rates are therefore high within social networks that have a high share of married and unmarried agents. To allow for the fact that the set of relevant others may change during the life course we assume that individual characteristics such as the age of the individual agent will determine the size and characteristics of the set of relevant others. These assumptions are based on stylised empirical facts that show that the number of relevant others increases during youth

and adulthood and thereafter is reduced again. Based on empirical facts that show a strong homogamy of marriage within socioeconomic groups, we assume that the social network is determined by individual characteristics such as age, kinship, spatial location, education, etc. To yield a parsimonious representation of the social network we restrict the set of characteristics to the two-dimensional space only, with age and spatial location as the two key characteristics. A further assumption is that we neglect divorce, modelling marriage as an irreversible process.

3.2.1.2 Implementation

One of the most difficult steps is to formalise the various theories and assumptions in a way that they can be implemented in a computer simulation.

We start by defining the world in which agents move. For this we locate agents along a torus and establish each agent with two characteristics: a spatial location $\phi \in [0, 2\pi]$ along a circular line on the torus and a second characteristic, the vertical location on the torus, which may represent the age of the individual (cf. Fig. 3.1). The geometry of a circular line gave rise to naming it the “wedding ring” and it has the advantage that the neighbourhood for each agent is contained within the circular line. We next define the set of relevant others as a two-dimensional neighbourhood that is symmetric w.r.t. the location and asymmetric w.r.t. the age of the agent. The time scale of our simulations is therefore a calendar year that corresponds to the age of the agents. The share of married persons in the network determines the social pressure to get married. We allow that also in case of zero married couples in the set of relevant others, social pressure is positive, and we assume that social pressure is increasing with the number of married individuals in the set of relevant others and follows an S-shaped function. The social pressure itself together with the age of the agent determines another set of relevant others within which an agent looks




Fig. 3.1 Implementation of the agent-based model (Billari et al. 2007, p. 65)

for partners. Based on stylised empirical facts we assume that age influences the extent of this region in a non-monotonic way, being highest between ages 21 and 38. We only allow marriages between different sexes and assume a mutual search, i.e. two agents will marry only if both agents are part of the set of relevant others of the other agent. Once married, agents will have children according to a set of predetermined age-specific fertility rates where we adjust fertility to keep the total population constant. Children born to married couples are located near parents and their characteristics are initialised in a similar way as for the initial population.

Next we need to define the characteristics of each agent, formalise the specific rules of agents and define the various simulation steps. At this stage a flow diagram may be helpful to follow the working of the model (Fig. 3.2). Each agent is characterised by a numerical identifier, year of birth, sex, age, spatial location, length of symmetric interval in which the agent searches for potential partners, social pressure, marital status, identifier of partner if married, marriage duration, set of relevant others and set of potential partners. Note that all characteristics are time-dependent except the birth year, sex and the numerical identifier.

To run the simulations we initialise the starting population by the age distribution of the US in 1995 and based on the same population we assign the sex and marital status randomly. For the initial population we also assign the marriage duration randomly to each agent. We run the model for 150 years. To define the set of relevant others we assume five kinds of agents that differ in their preference as to whether they prefer others in the same, younger or older ages and combinations of these.

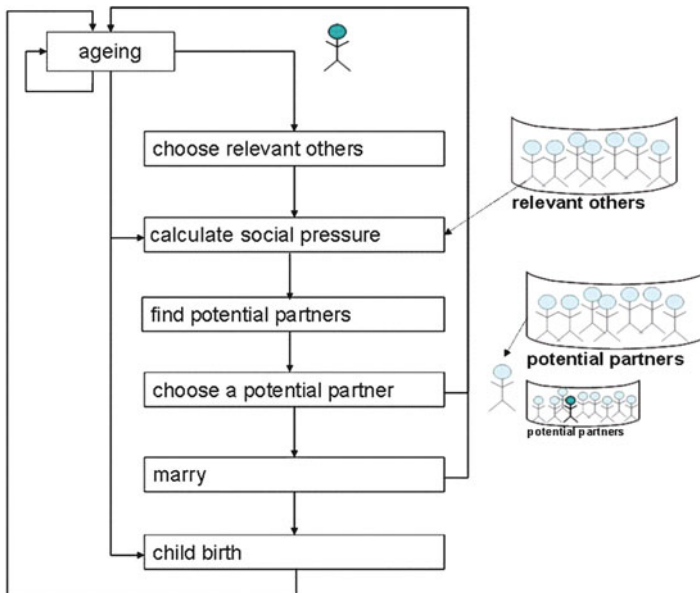


Fig. 3.2 Flow diagram (Billari et al. 2007, p. 70)

Next we randomly assign agents a type, and for each type we randomly choose the midpoint of each age interval, the width being again a random variable. Among this set of agents we choose a random number of agents to be their relevant others. The social pressure is then determined by the number of relevant others that are already married. To define the space of potential partners we transform the number of the social pressure into a distance: $d = sp(pom) * m(N) * ai(x)$, where the social pressure sp increases with the share of relevant others being married as denoted by pom , $m(N)$ denotes a factor that depends on the initial size of the population N and $ai(x)$ reflects the fact that age x determines the network size with $ai(.)$ presenting a non-monotonic function first increasing in age and then decreasing again at older ages. The functions $sp(.)$ and $ai(.)$ are both bounded in the interval $[0, 1]$. Potential partners are therefore in the spatial range of $[\phi - d, \phi + d]$ and the age range $[x - sp(pom) * ai(x) * c, x + sp(pom) * ai(x) * c]$ where we assume the positive constant c to be equal to 25. Hence, the maximum age difference for potential partners will be 25 years.

In each simulation step, the agent ages by one year and the final age at which agents die is set to 100. Agents start to search for partners at age 16. In every simulation step, agents choose the set of relevant others which then determines the social pressure. The arrow from the aging box to social pressure indicates that the specific value of the social pressure is age dependent. Next, the agent determines the set of potential partners which will depend on the social pressure. Within the set of potential partners the agent looks for a partner. If the agent finds a partner it is checked whether the agent herself is in the set of potential partners of her partner in which case the two agents get married. Once married the agent gives birth to new agents according to an exogenously fixed age-specific fertility rate (we applied the US age-specific fertility rate of 1995). If the agent is not married in one simulation step it ages and starts the search for relevant others and potential partners again. Otherwise, when the agent is married it just follows the ageing process over its remaining life cycle.

3.2.1.3 Simulation

The software used for the simulations is NetLogo. To obtain smooth simulation results we set the initial population to 800 and take the average over 75 consecutive cohorts and 100 simulation runs. To validate our simulations we test whether our model can replicate the qualitative shape of the age-at-marriage curve observed at the aggregate level. In a next step we can then investigate which of our assumptions are most important to capture the qualitative shape of the age-at-marriage hazard curve. Indeed, agent-based models represent a toolkit to perform such counterfactual experiments. Our simulation results (cf. Fig. 3.3) demonstrate that our model is able to reproduce the empirically observed right-skewed bell-shaped distribution of the age-at-marriage hazard. However, if we ignored the age dependency of the set of

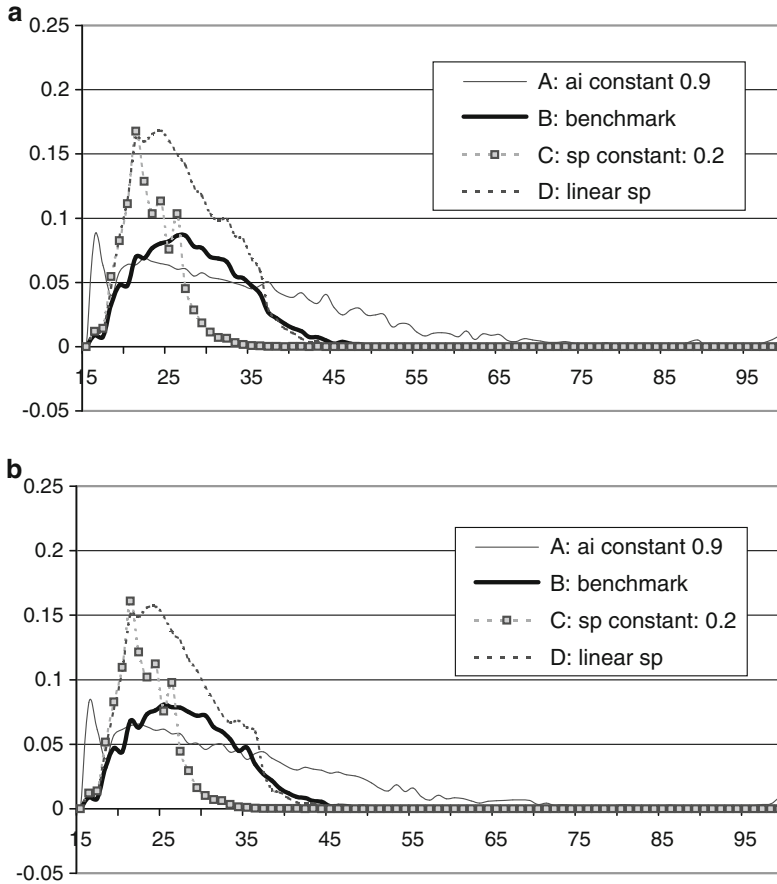


Fig. 3.3 Hazard of marriage in a population of simulated agents with alternative settings for social pressure. (a) Women (b) Men (Billari et al. 2007, p. 72)

relevant others or/and used a linear or constant functional form of the social pressure as it depends on the married couples in the set of relevant others, the shape of the age-at-marriage hazard rate would be very different. Further sensitivity analyses w.r.t. the asymmetry of the age interval that determines the set of relevant others show that either extreme—no asymmetry or a much higher asymmetry—implies an increase in the proportion of married agents within the set of relevant others compared to the benchmark simulations. Moreover, the exact form of the S-shaped social pressure function will also determine the quantitative shape of the age-at-marriage hazard.

3.2.1.4 Conclusion

Our model of the marriage market nicely demonstrates how the shape of the age-at-marriage pattern emerges as the outcome of individual behaviour and social interactions. Not only can we replicate the qualitative shape of the age-at-marriage hazard but our model allows us to discern the most important characteristics of the network structure that may explain the qualitative shape of the marriage process. For a quantitative analysis, however, we would need an empirical validation of our parameters for different societies.

The mechanisms we present are in accordance with the work by Dixon (1971), who argues that three key variables are important in determining the marriage pattern. These include the *availability of mates* (represented by the set of potential partners), the *feasibility of marriage* (represented indirectly through the initial distribution of ages at marriage) and the *desirability of marriage* (represented by the influence of the social pressure).

3.2.2 *Transition to Parenthood: Social Interactions and Endogenous Networks*

3.2.2.1 Theory and Assumptions

In fertility theories, diffusion processes have long been shown to underlie the observed fertility decline (e.g. Palloni 1998; Watkins 1987; Kohler 2001). Fertility behaviour not only depends on family background variables and the life course path, but also on the behaviour and characteristics of other individuals which are linked through social networks. Within such networks, beliefs, norms, services etc. are exchanged, learning from others takes place and agents may feel induced to conform to prevailing norms. However, the formalisation of social interactions to explain fertility behaviour is missing in theoretical demographic models. Agent-based models allow testing the importance of social interaction effects and the prevailing micro and macro feedback to explain actual fertility patterns. In Aparicio Diaz et al. (2011) we have set up an agent-based model to study the role of social interaction for the transition to parenthood and in particular we focused on endogenous network formation according to Watts and Strogatz (1998) and Watts et al. (2002). More specifically we assumed that age and intended education determine the affiliation to a social network. The share of mothers within the social network positively impacts the desire to have a child for the agent. We then also applied our model to project age-specific fertility rates for Austria.

3.2.2.2 Implementation

To formalise our model we need to first define the characteristics of our agents. We assume a one-sex model with only female agents and distinguish agents by their age, intended education (which we assume to be already known at childhood) and parity. We distinguish six stages of parity ranging from 0 to 5+ and three groups of education: primary and lower secondary, upper secondary, and tertiary. Each agent is furthermore characterised by a numerical identifier and her social network. Age and intended education determine the affiliation to a social group. Agents choose \bar{s} members in their social network. Her mother and her siblings are also part of her social network. These selected members of the social network will then influence the birth probabilities of the agent. We assume the reproductive period to start at age 15 and end at age 49 years. The maximum age of agents is set to 100 years.

To initialise our model we assume an age distribution that reflects the Austrian female age distribution and set the intended education similar to the age-specific educational distribution of Austrian females at age 30. To set the parity level we refer to the age and educational parity distribution of Austrian females and assign an age at first birth based on the Austrian female population in case the agent is of parity greater than one.

Next we need to define the various simulation steps, among them the formation of the endogenous social network and the feedback rules between the micro and macro behaviour.

In every simulation step (that equals one calendar year) agents age and are at risk of dying. At age 15 we assume that a social network is formed. Based on the social influence of members of the social network si_i , where the subscript i denotes the i -th agent, the empirically observed age and parity-specific birth probabilities $\overline{bpr}_t(x, p)$ are altered, with x denoting age and p denoting parity. In case of parity greater than zero we also postulate that birth probabilities are related to the age of the youngest child xc_i where the relation is represented by the functional form $g(xc_i)$ that is decreasing in xc_i . The average of the birth probabilities at step t at the individual level determines the new updated birth probabilities at the macro level for step $t + 1$. To sum up, at stage t individual birth probabilities are determined as follows: $\overline{bpr}_t(x, 0) = \overline{bpr}_t(x, 0)si_i$ for parity zero and $\overline{bpr}_t(x, p) = \overline{bpr}_t(x, p)si_i g(xc_i)$ for parity greater than zero. In stage $t + 1$ the new macro-level birth probabilities are given by $\overline{bpr}_{t+1}(x, p) = \overline{bpr}_t(x, p)\overline{si}_t(x, p)$.

To endogenously set up the social network we define a distance d_{ij} between any two agents i and j that is determined by their difference in education e and age x : $d_{ij} = \epsilon|e_i - e_j| + |x_i - x_j|$ where ϵ is a constant to adjust for the fact that the maximum age difference is much higher than the maximum education difference. The agent then chooses a distance d with probability $pr_1(d) = c \exp(-\alpha d)$, with the parameter α denoting the degree of homophily and c representing a normalisation parameter. High values of α imply that agents chosen are more similar. Agents are searched for until the social network size is equal to a desired number s which itself is drawn from a log-normal distribution with a given mean of \bar{s} . In addition to friends chosen




Fig. 3.4 Social influence function (Aparicio Diaz et al. 2011, p. 569)

by this procedure, each social network also contains the mother of the agent and her siblings. Furthermore we assume a mutual friendship relation in the sense that the selecting agent is also added to the social network of the selected agents.

We assume that the share of mothers within the social network will induce a positive influence on the agent's parity transition and implement the social influence function as a threshold function similar to Granovetter (1978) and Schelling (1978). To account for different education-specific network effects we model the social interaction so as to vary by educational level of the agent assuming different intensities of social influence by education. Formally we model the social influence for an agent of parity p as an S-shaped function (cf. Fig. 3.4) which increases in the difference between the share of mothers at higher parity compared to the agent, in her social network π and in the entire population π^* : $si = 1 + \xi \left[\frac{\exp(\beta(\pi - \pi^*))}{1 + \exp(\beta(\pi - \pi^*))} - \frac{1}{2} \right]$. The parameters β and ξ measure the intensity of the social influence. For ξ we use different values for different educational groups and for different parities. We assume higher values for higher educational groups and lower values for higher order births representing the fact that higher educated women conform more to social pressure and social pressure decreases with higher order births. We assume that agents are only influenced by members in their social network with higher parity. Only when the share of mothers with higher parity in the social network differs from the corresponding share in the whole population will there be a positive social influence.

3.2.2.3 Simulation

Our simulations are based on the Austrian age-, education- and parity-specific distributions for 1981, 1991 and 2001, and we rely on the Austrian life table for those years to account for age-specific mortality. Setting our initial population equal to $N = 50,000$, we take the average over 25 simulation runs. We start our first set of

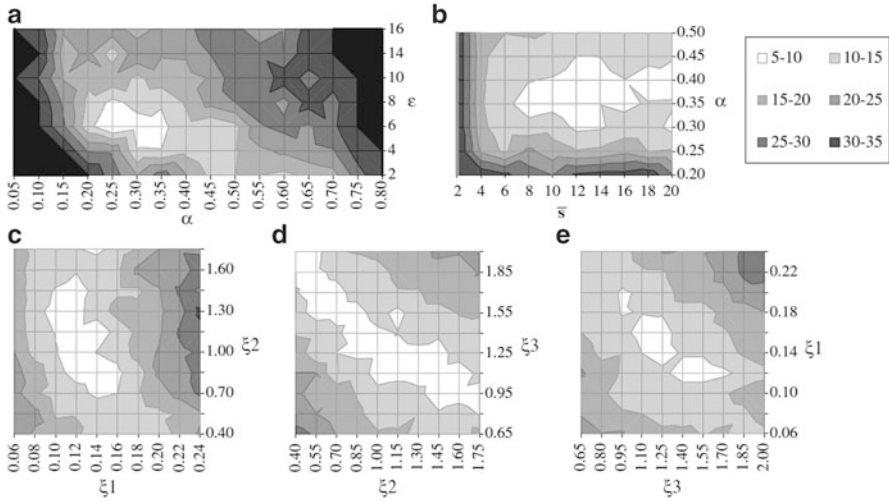


Fig. 3.5 Fit of simulated age-specific fertility rates for combinations of model parameters. In panel **a**, the level of homophily α (the higher the value the more similarities among network members) and the weight of education difference ϵ (the higher the value the more important becomes the education when choosing network members) are varied; in panel **b**, the mean social network size \bar{s} and α are varied; in panels **c–e**, the education-specific intensities of network influence are pairwise compared. Fixed parameters are set at $\alpha = 0.35$, $\epsilon = 5$, $\bar{s} = 10$, $\xi_1 = 0.14$, $\xi_2 = 0.85$, and $\xi_3 = 1.25$ (Aparicio Diaz et al. 2011, p. 572)

simulations with the Austrian population in 1981 and birth probabilities as of 1984¹ and simulate the model forward in time for 20 years. As a performance measure for our simulations we take the sum of absolute differences between simulated and observed age-specific fertility rates in 2004. We choose two-dimensional contour plots of this difference where we vary two of the model parameters simultaneously (cf. Fig. 3.5). In this way we can investigate the sensitivity of our model to model parameters. Or alternatively, the exercise allows us to calibrate the parameters such that our model can replicate the observed path of fertility trends. Results indicate that our model performs best for medium ranges of the homophily parameter α , a minimum network size \bar{s} of six and when agents of lowest education are influenced the least by their peers.

Similar as to the marriage model we perform counterfactual simulations where we assume no social interaction. In this case our simulation results can neither replicate the observed increase in the mean age at birth nor the decrease in fertility rates observed in Austria since 1981 (cf. Fig. 4 in Aparicio Diaz et al. (2011), p. 573). However, when we include social interaction, our model is capable of replicating the development of the Austrian fertility pattern between 1984 and 2004 very well (cf. Fig. 3.6). The trend of the simulated TFR in panel (a) fits the trend

¹Age-specific birth probabilities are not available for the time before 1984.

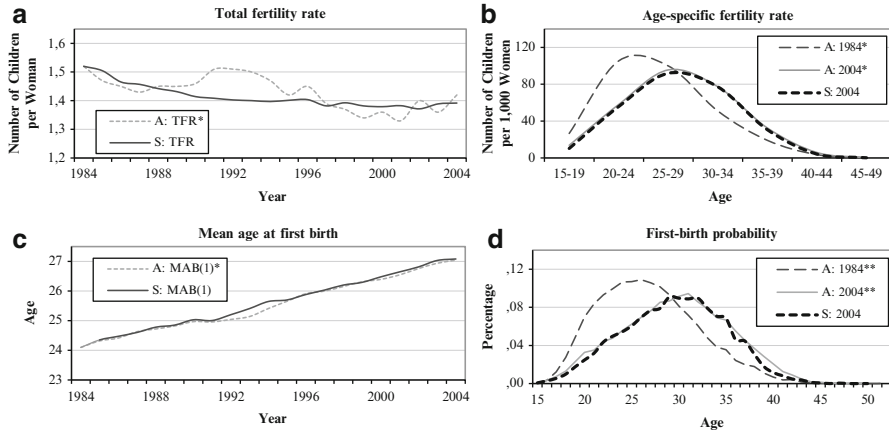


Fig. 3.6 Simulation results for simulating 20 years starting from 1984. “A” stands for empirical data of Austria, and “S” denotes simulated results. An asterisk denotes that data are from Statistik Austria (2007), double asterisks indicate computations by T. Sobotka based on data provided by Statistics Austria (Aparicio Diaz et al. 2011, p. 574)

in empirical TFR quite well. Panels (b), (c), and (d) show that the simulated curve of age-specific fertility, mean age at first birth and first birth probabilities constitute a remarkably good fit to the observed ones. Similarly, if we run the model from 1991 onwards for 15 years our results indicate a good fit with the actual Austrian fertility development. In particular, our model can reproduce the shift in the timing of fertility that occurred during the last decades.

We have also investigated forecasts of our model and compared these with the fertility assumptions postulated by the Austrian population projections. Such an exercise highlights the potential of using agent-based models as a causal model to explain trends in fertility rather than applying only projections based on trend extrapolation of past fertility development.

3.2.2.4 Conclusion

In the application summarized in this section we have applied an agent-based model to suggest a more realistic model of how social interaction (through social learning and social influence) may help to explain observed family formation patterns in contemporary Europe. By applying an agent-based modelling framework we can conduct “thought experiments that explore plausible mechanisms that may underlie observed patterns” (Macy and Willer 2002, p.147). Such an approach is particularly helpful when trying to explain the trends in fertility timing and quantum over the last decades. Our numerical results and further sensitivity analyses have clearly identified the importance of the characteristics of the social network in explaining the observed fertility patterns. Most importantly, our model can not only explain

the decrease in fertility at younger ages but also its increase at higher ages. We are not arguing that other factors such as socioeconomic conditions (employment, education, etc.) and shifting values of younger women towards less traditional female roles in the family and in society do not play a role in explaining family formation. The aim of our model is to highlight that social interaction may be an important mechanism as well. As a next step we need to test whether our model can also explain fertility patterns in other European countries. In particular, the framework of agent-based modelling allows us to experiment with alternative mechanisms that may underlie the timing and quantum of fertility in different social environments. Our model can therefore be regarded as an attempt in the exploration of identifying mechanisms that may underlie observed family formation patterns. Such an approach is indispensable in order to understand past trends and to project future developments of demographic behaviour.

3.2.3 Family Policies and Social Structure

3.2.3.1 Theory and Assumptions

The social structure within a society may not only influence demographic behaviour at the individual level, but also mediate the role of policies targeted at the individual level. As we have demonstrated in Fent et al. (2013), the effectiveness of family policies may be closely correlated with the prevailing social structure. So far, the empirical evidence on the importance of family policies to explain fertility is rather mixed. Gauthier (2007) shows in a comprehensive survey of the literature that family policies at the micro level have an effect on completed fertility while studies at the macro level indicate an effect of family policies on the timing of fertility. However, studies on the role of family policies for fertility—whether at micro or macro level—commonly ignore the prevailing social structure in a society though the role of social networks is undisputed in models of fertility behaviour. In Fent et al. (2013) we build on these ideas and assume that fertility preferences are subject to a diffusion process among individuals. Family policies, by altering fertility preference in some females, may therefore act as an effective device to induce a snowball process with the changing preferences spreading from person to person. We distinguish between a *direct effect* of family policies as captured by an alleviation of resource constraints (through e.g. institutional childcare or financial benefits) at the individual level and an *indirect effect* of family policies that captures the diffusion of fertility intentions via social ties as described above. The latter indirect effect captures the argument that any additional child resulting from family policies may induce others in the same peer group to increase their own fertility intentions. Such multiplier effects that operate through peer group effects form the basis of the work of Feyrer et al. (2008). Within an agent-based model we can combine these various levels of analysis: family policies at the macro level, their diffusion within social networks and their ultimate effectiveness at the individual

level. Our model can be regarded as an attempt to show that family policies can only be effective if they account for the characteristics of the society in which they are implemented.

3.2.3.2 Implementation

To build up an agent-based model we start to define the characteristics of the agents. We assume a one-sex population of females only. Each agent is characterised by a numerical identifier, age, her household budget, parity, number of dependent children who do not yet have their own income, intended fertility, her social network and her assignment to a specific income quantile. The household budget is composed of income as well as the monetary value of non-working time. Household income is allocated in order to satisfy the agent's own and her children's needs. We assume that an agent stays in the same income quantile over all her life but may progress to higher income levels as the agent ages. We then define necessary conditions to have a child. First, intended fertility has to exceed actual parity and secondly, the disposable income (i.e. income less the consumption needs of the agent and her children) has to exceed the costs of a child. If both conditions are fulfilled, an agent is exposed to the biological probability (fecundity) of having another child. Every new-born child is then linked to its mother and sisters. Agents age by one year in every time step and depending on their age and labour force participation they become adults. Once an agent is an adult, she gets her own income, own social network and own fertility intention.

Next we have to formalise the working of family policies and how they intervene in the fertility decision. We distinguish between a fixed family policy b^f and a variable family policy b^v , the latter being proportional to the household budget $w_{i,t}$ where i denotes the i -th agent. Both kinds of family policies reduce the costs of $n_{i,t}$ dependent children: $c_{i,t}^{n_{i,t}} = n_{i,t}(\tau\sqrt{w_{i,t}} - b^f - b^v w_{i,t})$ where the costs of children $\tau\sqrt{w_{i,t}}$ are measured in terms of their consumption which is a concave function of the household budget $w_{i,t}$ and τ denotes a parameter.

The social network is built up in a similar way as in Aparicio Diaz et al. (2011), but we are assuming that not only two but three characteristics (age x , income z and intended fertility f) determine the closeness between agents i and j as presented by the distance $d_{ij} = |x_i - x_j| + \epsilon_1 |z_i - z_j| + \epsilon_2 |f_i - f_j|$. Agents choose a distance d with probability $pr_1(d) = c \exp(-\alpha d)$ and pick an agent with this distance as a new friend. We define another probability pr_2 , which determines whether this new friend is chosen among those individuals who are not linked to any of the agent's peers or only among those individuals who are linked to at least one of the agent's friends. We therefore assume a specific level of network transitivity, i.e. two agents being connected to the same agent can build a mutual relationship. The constant c is a normalisation parameter to ensure that the probabilities of all of the feasible distances sum up to one, and the parameter α determines the agents' level of homophily. The selecting agent is also added to the network of the selected agent. Thus, we assume a mutual friendship relation. We repeat this procedure until the

desired number of peers, s , is found. This desired network size is drawn from a lognormal distribution with mean $\bar{s} = 10$ and rounded to the nearest integer.

We next define a diffusion mechanism that is based on local ties and operates on the intended fertility. The specific social effects are modelled as in Goldenberg et al. (2007). We assume that intended fertility $f_{i,t}$ of agent i increases (decreases) by one with probability pr_3 (pr_4) due to the social effects exerted by a peer with a parity greater (less) than the agent's intended fertility. Then, we compute π_i^+ (π_i^-), the number of agents j who are linked to i and have a parity greater (less) than the intended fertility of agent i , i.e. $p_{j,t} > f_{i,t}$ ($p_{j,t} < f_{i,t}$). We next compute the probabilities for an agent to be positively or negatively influenced by at least one agent from the peer group,² $p_{i,t}^+ = 1 - (1 - pr_3)^{\pi_i^+}$ and $p_{i,t}^- = 1 - (1 - pr_4)^{\pi_i^-}$. Individuals may be exposed to positive influence, negative influence, both positive and negative influence, or neither. Hence, the probability of being only positively (negatively) influenced becomes $(1 - p_{i,t}^-)p_{i,t}^+$ (respectively $(1 - p_{i,t}^+)p_{i,t}^-$) and the probability of being positively and negatively influenced is $p_{i,t}^+p_{i,t}^-$. We use the parameter κ (or $(1 - \kappa)$) to determine the fraction of individuals who increase (decrease) their intended fertility in the case of mixed influence. Then, the probabilities of increasing, decreasing or keeping the intended fertility constant are

$$\begin{aligned} p_i(f_{i,t+1} = f_{i,t} + 1) &= (1 - p_{i,t}^-)p_{i,t}^+ + \kappa p_{i,t}^+p_{i,t}^- \\ p_i(f_{i,t+1} = f_{i,t} - 1) &= (1 - p_{i,t}^+)p_{i,t}^- + (1 - \kappa)p_{i,t}^+p_{i,t}^- \\ p_i(f_{i,t+1} = f_{i,t}) &= (1 - p_{i,t}^+)(1 - p_{i,t}^-). \end{aligned}$$

3.2.3.3 Simulation

We start with six distinct populations of agents. Each agent is characterised by her age, parity, number of dependent children, intended fertility and household budget. The distribution the individual characteristics are drawn out of is the same for all six populations, i.e. the populations only differ with respect to the realisation of the specific values. Each population consists of 5,000 agents. Our interest is in the role of social interaction for the effectiveness of family policies. We therefore vary the following set of parameters: the level of fixed and proportional family policy, the homophily parameter that describes the network structure, the degree of transitivity of the network structure, the importance of the level of intended fertility as a characteristic that determines the social network, and the strength of positive and

²If pr_3 is the probability of increasing intended fertility due to meeting one peer with a higher parity, then $(1 - pr_3)$ is the probability of not increasing intended fertility despite this one peer, $(1 - pr_3)^{\pi_i^+}$ is the probability of not increasing intended fertility despite π_i^+ peers with higher parities, and $1 - (1 - pr_3)^{\pi_i^+}$ is the probability of increasing intended fertility when being exposed to π_i^+ peers with higher parities.

negative social influence. We choose quite an extended set of parameter variations resulting in 88 different sets of family policies run on 8,424 different societies where each simulation is run for 100 time steps. The outcome variables against which we test our model are the aggregate values of cohort fertility, intended fertility and the fertility gap. While intended fertility allows us to measure the indirect effect of family policies, the fertility gap, i.e. the difference between intended and realised fertility, allows to measure the direct effect of fertility policies.

In Fig. 3.7 we present results of our simulations on cohort fertility, intended fertility and the fertility gap as a function of variable family policies (left column) and fixed family policies (right column). Within each figure we distinguish between presenting the results as average for all simulations and alternatively as average over simulations of a specific variable family policy and alternatively over simulations of a specific fixed family policy. Since the effect of family policies on cohort fertility and the fertility gap is stronger as compared to the effect on intended fertility, the direct effect of family policies seems to be more important according to our simulations. We also run OLS regressions on our simulation results with the various fertility measures as dependent variable and the monetary values of fixed and variable family policies as explanatory variables. Our regressions confirm the graphical representations and indicate that fixed family policies have a stronger impact.

Further results of our simulations reveal that the degree of homophily has a strong impact on the indirect effect of family policies (since intended fertility is very sensitive to this parameter). Similarly, the difference between being positively or negatively influenced by peers also has a very pronounced effect on intended fertility, even exceeding the effect of family policies. To better quantify the role of family policies within societies being characterised by different social structures we conducted another extended regression analysis on our simulations where we included further explanatory variables. In particular we added variables that characterise the social network and allowed for interactions between them and the variables of fixed and variable family policies. Results of these regressions indicate that variable family policies contribute more to the indirect effects while fixed family policies contribute nearly the same to indirect and direct effects. While the indirect effect is more sensitive to social effects for variable as compared to fixed family policies, the reverse holds for the direct effect of family policies. Our simulation results therefore clearly indicate that neglecting the social structure in which family policies operate will yield a wrong assessment of family policies.

3.2.3.4 Conclusion

In the application summarized in this section we have applied an agent-based modelling approach to combine the literature on social interactions and fertility behaviour with the literature on the role of family policies for fertility intentions and fertility realisations. Such an approach allows us to experiment with different family policies and their relation to the prevailing social structure in a society. Most

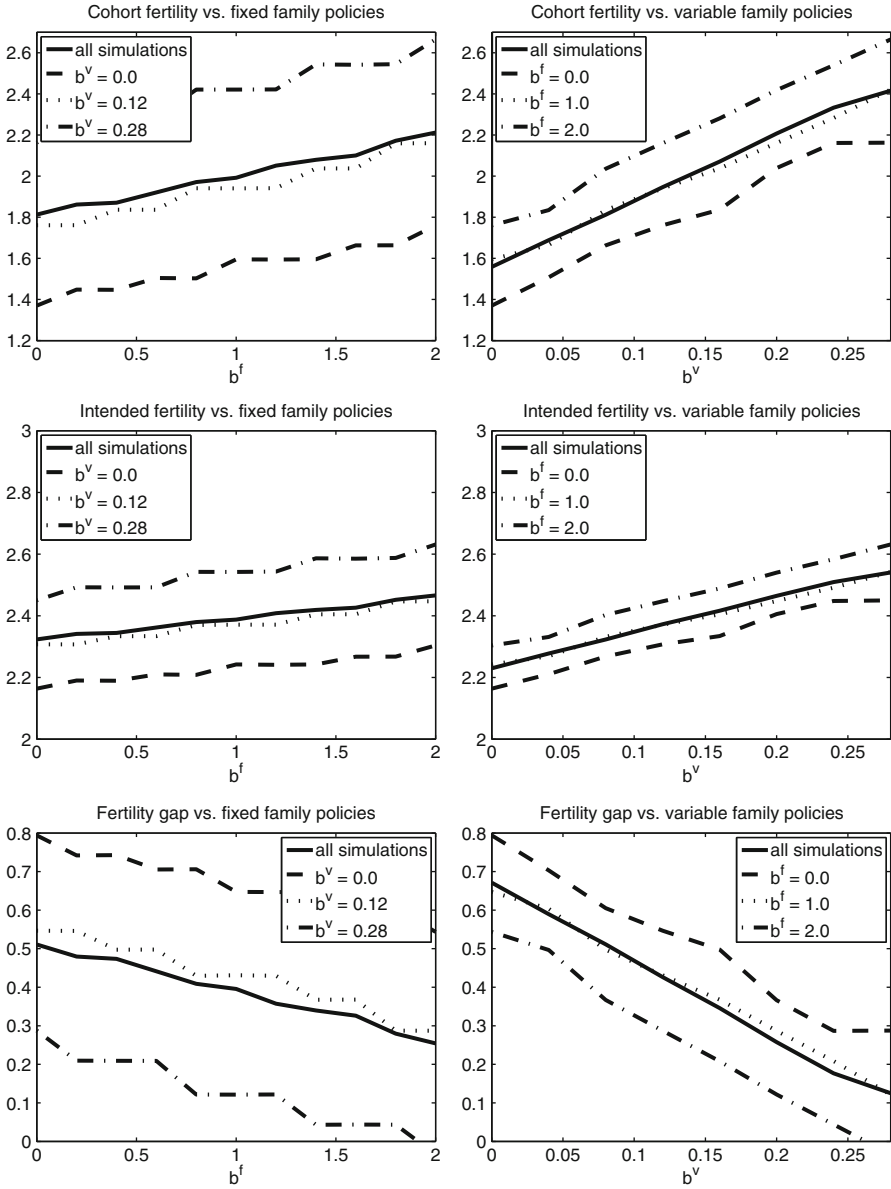


Fig. 3.7 Completed cohort fertility, intended fertility and fertility gap by fixed, b^f , and variable, b^v , family policies (Fent et al. 2013, p. 972)

interestingly, we can study the social multiplier effect on fertility, i.e. to which extent are family policies mitigated or reinforced by social effects. Furthermore we could also show that while a higher level of homophily among network partners induces a positive impact on fertility intentions and realisations, it may at the same time reduce the effectiveness of family policies. Our model setup also allowed us to differentiate between an indirect and direct effect of family policies and will help to better understand the fertility gap. In summary, our framework highlights the fact that the success of a family policy will depend on whether it takes into account the societal structure to which it is applied. Since our framework so far only constitutes a selection of variables that may influence fertility, further extensions are obvious. At the micro level partnership formation and employment uncertainties are important determinants. At the meso level social capital and place of residence are further important variables. At the macro level economic trends, advancement in reproductive technologies, and changes in attitudes and norms are important further determinants.

3.3 Summary

As the previous three examples on demographic behaviour and social interaction have shown, agent-based modelling can be regarded as a tool of *theorising via simulation*, complementary to *theorising via statistical modelling*. Agent-based models offer a tool to unify the typical rate-based approach of micro-simulation with a new rule-based approach. Similar to the life course approach, agents play a central role, however they may also interact with other agents.

According to the typology presented in Gilbert (2008), the set of agent-based models reviewed in this paper are rather toy models or mid-range models.³ We have combined well-known and partly empirically observed micro mechanisms into a larger but still abstract model to produce familiar macro mechanisms in a new way.

In the Wedding Ring model our framework allows for an endogenous explanation of the age-at-marriage hazard. By modelling the fact that social networks vary with age, we could—compared to previous models that relied on exogenous explanations to prevent the age-at-marriage hazard to increase too fast at younger ages—replicate the age-at-marriage hazard in younger ages in a more realistic way. In our second example on the transition to parenthood we could replicate the age-specific changes in fertility rates as they have been observed in Austria over the last decades. The last example on family policies goes beyond a mere description of the macro behaviour.⁴ It aims to test the effectiveness of family policies on fertility decisions at the micro level in the presence of social interaction and studies how macro outcomes such

³In contrast to the class of toy and mid-range models presented in this paper, there exist highly data-driven ABMs that model real situations, e.g. the study by Axtell et al. (2002) on a historical population and studies on household dynamics and land use change by Entwisle et al. (2008).

⁴See also Baroni et al. (2009) for the integration of policies to explain fertility in an agent-based modelling framework.

as cohort fertility and the fertility gap (difference between intended and realised fertility) may differ depending on the interaction between family policies and the prevailing social structure.

The challenge in setting up an agent-based model is to select the characteristics and rules of the agents and to define how agents may interact and how macro-economic behaviour may feed back on the micro-level decisions processes. These choices together with the specific functional forms and parameters chosen need to be based on empirical evidence as far as possible. Once an agent-based model has been implemented an extensive sensitivity analysis is imperative. Obviously this is quite a computational intense task. In a first step the model should be calibrated to match the macro behavior to be explained. In a second step the set of specific characteristics, rules and interactions of agents that have been postulated should be reassessed in terms of their relevance. E.g. in the Wedding Ring model we could show that the age dependence of the set of relevant others and the specific nonlinear shape of the social interaction function are important to explain the qualitative shape of the age-at-marriage hazard. In our second model on the transition to parenthood we have tested the specific model parameters to identify which combination of parameters may best replicate the observed fertility behaviour. In the third example on the role of social structure for family policies we have investigated the importance of fixed versus variable family policies by applying regression analysis on our simulation results.

Agent-based modelling requires a good knowledge of tools in computer simulations, but also in statistics and probability theory. Developing the formal behavioural rules and interactions of agents also requires some skills in mathematical formalization. From our experience so far the acceptance of agent-based models in demography is closely related to how far one can convincingly demonstrate that the model is able to explain and yield new insight into quantitative or at least qualitative shapes of important macro-demographic behaviour. For this, an extensive sensitivity analysis that highlights the key ingredients of the model is imperative.

We have presented our models following the same sequence of steps to be conducted. The development of a standard protocol to describe ABMs is definitely very important and needs to be improved. Our approach in this survey chapter is only a first attempt in this direction.

Acknowledgements I would like to thank Bernhard Rengs and Thomas Fent for their comments and suggestions and Werner Richter for proof reading.

References

- Åberg, Y. (2003). *Social interaction. Studies of contextual effects and endogenous processes. Stockholm studies on social mechanisms*. Department of Sociology, Stockholm University.
- Aparicio Diaz, B., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous network. *Demography*, 48(2), 559–579.

- Axelrod, R. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton: Princeton University Press.
- Axelrod, R., & Tesfatsion, L. (2006). A guide for newcomers to agent-based modeling in the social sciences. In *Handbook of computational economics* (vol. 2, pp. 1649–1659). Amsterdam: Elsevier.
- Axtell, R. L., Epstein, J. M., Dean, J. S., Gumerman, G. J., Swedlund, A. C., Harburger, J., Chakravarty, S., Hammond, R., Parker, J., & Parker, M. (2002). Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciences of the United States of America*, 99(suppl. 3), 7275–7279.
- Baroni, E., Hallberg, D., Eklöf, M., Lindh, T., & Zamac, J. (2009). Fertility decisions – Simulation in an agent-based model (IFSIM). In A. Zaidi, A. Harding, & P. Williamson (Eds.), *New frontiers in microsimulation modelling*, part III, (pp. 265–286). Farnham: Ashgate.
- Billari, F. C., & Prskawetz, A. (2003). *Agent-based computational demography: Using simulation to improve our understanding of demographic behaviour*. Heidelberg: Physica Verlag.
- Billari, F. C., Prskawetz, A., Aparicio Diaz, B., & Fent, T. (2007). The “Wedding Ring”: An agent-based marriage model based on social interaction. *Demographic Research*, 17, 59–82.
- Dixon, R. B. (1971). Explaining cross-cultural variations in age at marriage and proportions never marrying. *Population Studies*, 25(2), 215–233.
- Entwisle, B., Malanson, G., Rindfuss, R. R., & Walsh, J. (2008). An agent-based model of household dynamics and land use change. *Journal of Land Use Science*, 3(1), 73–93.
- Fent, T., Aparicio Diaz, B., & Prskawetz, A. (2013). Family policies in the context of low fertility and social structure. *Demographic Research*, 29, 963–998.
- Feyrer, J., Sacerdote, B., & Stern, A. D. (2008). Will the stork return to Europe and Japan? Understanding fertility within developed nations. *Journal of Economic Perspectives*, 22(3), 3–22.
- Gauthier, A. H. (2007). The impact of family policies on fertility in industrialized countries: A review of the literature. *Population Research and Policy Review*, 26(3), 323–346.
- Gilbert, N. (2008). *Agent-based models* (Quantitative applications in the social sciences). Thousand Oaks: Sage Publications.
- Goldenberg, J., Libai, B., Moldovan, S., & Muller, E. (2007). The NPV of bad news. *International Journal of Research in Marketing*, 24(3), 186–200.
- Granovetter, M. (1978). Threshold models of collective behavior. *The American Journal of Sociology*, 83(6), 1420–1443.
- Hernes, G. (1972). The process of entry into first marriage. *American Sociological Review*, 37, 47–82.
- Kohler, H. P. (2001). *Fertility and social interaction: An economic perspective*. Oxford: Oxford University Press.
- Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28, 143–166.
- Montgomery, M. R., & Casterline, J. B. (1996). Social learning, social influence, and new models of fertility. *Population and Development Review*, 22, 151–175.
- Palloni, A. (1998). Theories and models of diffusion in sociology. CDE Working Paper No. 98–11, Madison, Center for Demography and Ecology, University of Wisconsin.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2), 143–186.
- Schelling, T. C. (1978). *Micromotives and macrobehavior*. New York: Norton.
- Statistik Austria. (2007). Statistisches Jahrbuch Österreich 2007.
- Watkins, S. (1987). The fertility transition: Europe and the third world compared. *Sociological Forum*, 2(4), 645–673.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393, 440–442.
- Watts, D. J., Dodds, P. S., & Newman, M. E. J. (2002). Identity and search in social networks. *Science*, 296(5571), 1302–1305.

Part II
Designing, Analysing, and Reporting
Agent-Based Models

Chapter 4

Agent-Based Computational Demography and Microsimulation Using JAS-mine

Matteo Richiardi and Ross E. Richardson

4.1 Introduction

In this chapter, we present the implementation of a dynamic microsimulation with a rich set of demographic processes (birth, death, household formation and dissolution) and other life course events (educational choices, labour market participation and employment outcomes), using the recently upgraded JAS-mine simulation platform (www.jas-mine.net).

The chapter is meant to provide a step-by-step guide to the development of dynamic microsimulations/agent-based models. From a practical perspective, the model presented here is highly reusable and can be easily modified in order to develop other microsimulation/agent-based models.¹ This is thanks to the JAS-mine architecture, which envisages a neat separation between data (parameters

¹The model and the supporting documentation can be downloaded from the demo section of the JAS-mine website (www.jas-mine.net/demo/demo07).

M. Richiardi (✉)

Institute for New Economic Thinking at the Oxford Martin School, Mathematical Institute,
University of Oxford, Oxford, UK

Nuffield College, Oxford, UK

Collegio Carlo Alberto, Moncalieri, Italy

Department of Economics and Statistics, University of Torino, Turin, Italy

e-mail: matteo.richiardi@maths.ox.ac.uk; matteo.richiardi@unito.it

R.E. Richardson

Institute for New Economic Thinking at the Oxford Martin School, Mathematical Institute,
University of Oxford, Oxford, UK

e-mail: ross.richardson@maths.ox.ac.uk

and coefficients) and code, and a clear distinction between *modelling objects*, which specify the structure of a model and should be the primary concern of a researcher, and *auxiliary objects*, which perform useful tasks such as input-output communication, real-time visualization, etc.

The chapter is structured as follows. Section 4.2 motivates the need for a unique platform for agent-based models and dynamic microsimulations, integrating tools used by both modelling approaches. The section also lists other requirements that were specified for the platform. Section 4.3 briefly describes the technical solutions that were adopted to meet these requirements. Section 4.4 presents the general structure of a JAS-mine project. Section 4.5 describes the specific simulation model implemented. Section 4.6 illustrates the JAS-mine implementation and Sect. 4.7 concludes.

4.2 Convergence Between Agent-Based and Microsimulation Models

Microsimulation is a technique used in a large variety of fields to simulate the states and behaviours of different units (individuals, households, populations, etc.) as they evolve in a given environment (a market, a region, an institution). The word “dynamic” refers to the fact that the population being simulated is also changing, as opposed to “static” microsimulations (such as tax and benefit simulators) which examine the impact of a policy change on a given population (Li and O’Donoghue 2012). The modelling of demographic processes is therefore the salient characteristic of dynamic microsimulations.

Agent-based models are also computational models with individuals as the primary object of analysis. They mainly differ from microsimulations for their emphasis on the role of interaction and for explicit departures from the standard assumptions of economic models: rational expectations, perfect knowledge about the environment, infinite computational ability, absence of centralised “top down” coordination devices (Richiardi 2012).

Agent-based (AB) and microsimulation (MS) models share many features and can be described as belonging to the same class of discrete-event simulations. Indeed, from a mathematical and computational perspective the two approaches are identical. Both AB and MS models are recursive models, where the number and individual states of the agents in the system are evolved by applying a sequence of algorithms to an initial population (Gilbert and Troitzsch 2005). As computer-based simulations, they face the problem of reproducing real-life phenomena, many of which are temporally continuous processes, using discrete microprocessors. The abstract representation of a continuous phenomenon in a simulation model requires that all events be presented in discrete terms, hence the label discrete-event simulation.

However, in their historical development AB models and microsimulations have followed different trajectories (Richiardi 2013): AB models have focused more on

theory, while MS models have evolved as more data oriented, with the processes generally specified as probabilistic regression models. As a generalisation, AB models are structural models with a primary concern on *understanding*, while microsimulations are reduced-form models (as such, they often focus on one side of a market only), with a primary concern on *forecasting*.² However, a trend towards a convergence of the two approaches is currently underway, with AB models becoming increasingly empirically oriented, and MS models including more feedback effects (see again Richiardi (2013)). An example of this fruitful integrated approach is the recent field of agent-based computational demography (Billari and Prskawetz 2003).

The differences in scope and perspective between the two approaches have, however, impinged on the structure of the computer models used within each community. AB models lead naturally to an explicit object-oriented representation, while MS models are generally built around a database which is evolved forward in time. This has led to the development of simulation toolkits which are specific to each field, as for instance NetLogo (Wilenski 1999), RePast (North et al. 2013) and Mason (Luke et al. 2005) for AB modelling, and Modgen (Statistics Canada 2009), LIAM2 (De Menten et al. 2014) and JAMSIM (Mannion et al. 2012) for MS modelling – to name just a few.

JAS-mine was created to make the development of “hybrid” AB-MS models easier, and to allow researchers to use the same tools for both approaches, to exploit economies of scale in learning and coding. Its combination of features distinguish it from all the above platforms.

4.3 The JAS-mine Architecture

JAS-mine is an object-oriented Java-based platform for discrete-event simulations. The philosophy of JAS-mine is to foster *clarity*, *transparency* and *flexibility*. The rationale behind this is the belief that a major bottleneck in agent-based and dynamic microsimulation modelling comes from humans, rather than machines: minimizing modelling time then becomes even more important than minimizing computing time.³ To this aim, a strict adherence to the open source paradigm was enforced

²Structural models often include unobservable parameters that help describe individual behaviour at a deep level (say, in terms of utility maximisation); reduced-form models aim more simply at identifying statistical relationships between observable characteristics.

³The performances of JAS-mine with respect to speed of execution, though, are noteworthy. An exercise aimed at testing the performance of the simulation platform with respect to scaling involved the implementation in JAS-mine of a complex mixed AB-microsimulation model of the two-way relationship between health and economic inequality, calibrated on both US and Canadian cities. The JAS-mine implementation can run five million agents with a time-step equivalent to 1 day for 500 years (182,500 time-steps) in 50 min on a standard laptop (using less than 4GB of RAM).

in the design of the platform, which makes it less of a black-box with respect to proprietary software and encourages cooperative development of the platform by the community of users: all functions can be inspected and, if necessary, modified or extended. Also, it was decided not to develop an ad-hoc grammar and syntax – as in NetLogo and LIAM2 – but to allow the user to choose from a wide range of classes and interfaces which extend the standard Java language. The JAS-mine libraries therefore provide open tools to “manufacture” a simulation model, making use whenever possible of solutions already available in the software development community. This ensures efficiency and a maximum amount of flexibility in model building.

In the platform architecture, a clear distinction is made between objects with a modelling content, which specify the structure of the simulation, and objects which perform useful but auxiliary tasks, from enumerating categorical variables to building graphical widgets, from creating filters for the collection of agents to computing aggregate statistics to be saved in the output database. Moreover, a separation is made between code and data, with all parameters and input tables stored either in an input database or in specific MS Excel files. For instance, the *regression* package provides tools for simulating outcomes from standard regression models (OLS, probit/logit, multinomial probit/logit): in particular, there is no need to specify the variables that enter a regression model, as they are directly read from the data files. This greatly facilitates exploration of the parameter space, testing different econometric specifications, and scenario analysis.

From a modelling viewpoint, JAS-mine extends the *Model-Observer* paradigm introduced by the Swarm experience (Minar et al. 1996; Luna and Stefansson 2000) and introduces a new layer in simulation modelling, the *Collector*. The Model deals mainly with specification issues, creating objects, relations between objects, and defining the order of events that take place in the simulation. The Observer allows the user to inspect the simulation in real time and monitor some pre-defined outcome variables as the simulation unfolds. The Collector collects the data and computes the statistics needed both by the simulation objects and for post-mortem analysis of the model outcome, after the simulation has completed. This three-layer methodological protocol allows for extensive re-use of code and facilitates model building, debugging and communication.

As for input/output (I/O) communication, building on the vast number of software solutions available, JAS-mine allows the user to separate data representation and management from the implementation of processes and behavioural algorithms. The management of input data persistence layers and simulation results is performed using standard database management tools, and the platform takes care of the automatic translation of the relational model of the database into the object-oriented simulation framework, through object-relational mapping (ORM).⁴ This also allows

⁴ORM is a programming technique for converting data between incompatible type systems in object-oriented programming languages. This creates, in effect, a “virtual object database” that can be used from within the programming language.

to separate data creation from data analysis. As the statistical analysis of the model output is possibly intensive in computing time, performing it in real time might be an issue in large-scale applications. A common solution is to limit it to a selected subset of output variables. This, however, requires identifying the output of interest before the simulation is run. If additional computations are required to better understand how the model behaves, the model has to be run again: the bigger the model, the more impractical this solution is. On the other hand, the power of modern relational database management systems (RDBMS) makes it feasible to keep track of a much larger set of variables, for later analysis. Also, the statistical techniques envisaged, and the specific modeller's skills, might suggest the use of external software solutions, without the need to integrate them in the simulation machine. Finally, keeping data analysis conceptually distinct from data production enhances brevity, transparency and clarity of the code.

The architectural characteristics of JAS-mine are discussed in detail on the JAS-mine website (www.jas-mine.net). To summarise, the main features of the platform are:

- a discrete-event simulation engine, allowing for both discrete-time and continuous-time simulation modeling;⁵
- a *Model-Collector-Observer* structure (see Sect. 4.4);
- interactive (GUI based), batch and multi-run execution modes, the latter allowing for detailed design of experiments (DOE);
- a library implementing a number of different matching methods, to match different lists of agents;
- a library implementing a number of different alignment methods, to force the microsimulation outcomes to meet some exogenous aggregate targets (Li and O'Donoghue 2014);
- a library implementing a number of common econometric models, from continuous response linear regression models to binomial and multinomial logit and probit models;
- a statistical package based on the *cern.jet* package;
- an embedded H2 database;
- MS Excel I/O communication tools;

⁵Discrete-event simulations can be organized into two categories, depending on how time is treated. *Discrete-time* simulations break up time into regular time slices (Δt), while the simulator calculates the variation of state variables for all the elements of the simulated model between one point in time and the next. Nothing is known about the order of the events that happen within each time period: discrete events (marriage, job loss, etc.) could have happened at any moment in Δt while inherently continuous events (aging, wealth accumulation, etc.) are often thought to progress smoothly between one point in time and the next. By contrast, *continuous-time* simulations are characterized by irregular timeframes that are punctuated by the occurrence of the events. What is modelled is not whether an event occurs or not in the reference period, but rather the time elapsed before its occurrence (duration models). Between consecutive events, no change in the system is assumed to occur; thus the simulation can directly jump in time from one event to the next. Inherently continuous events must therefore be discretised (Keller et al. 1993).

- automatic GUI creation for parameters by using Java annotation;
- automatic output database creation;
- automatic agents' sampling and recording;
- powerful probes for real-time statistical analysis and data collection;
- a rich graphical library for real-time plotting of simulation outcomes;
- Eclipse plugin, which enables the creation of a JAS-mine project in just a few clicks, with template classes organised in the JAS-mine standard package and folder structure;⁶
- Maven version control.

4.4 The Structure of a JAS-mine Project

In the JAS-mine architecture, agents are organized and managed by components called *managers*. As already mentioned, there are three types of managers in this architecture: *Model*, *Collector* and *Observer*. Models serve to build artificial agents and objects, and to plan the time structure of events. Collectors are managers that build data structures and routines to calculate (aggregate) statistics dynamically, and that build the objects used for data persistence. The definition of a Collector's schedule specifies the frequency of statistics updating and agent sampling, and consequent storage in the output database. Observers are managers that serve to build graphical widget objects that indicate the state of the simulation in real time, and define the frequency with which to update these objects.

JAS-mine allows multiple Models (and multiple Collectors and Observers) to run simultaneously, since they share the same scheduler.⁷ This allows for the creation of complex structures where agents of different Models can interact. Each Model is implemented in a separate Java class that creates the objects and plans the schedule of events for that Model. Model classes require the implementation of the *SimulationManager* interface, which implies the specification of a *buildObjects* method to build objects and agents, and a *buildSchedule* method for planning the simulation events. Analogously, Collector classes must implement the *CollectorManager* interface, and Observer classes must implement the *ObserverManager* interface.

When a new JAS-mine project is created using the JAS-mine Eclipse plugin, several packages are created:

- **data**: package containing the classes that describe the structure of coefficients, parameters and agent population tables contained in the database to be loaded by

⁶Eclipse Integrated Development Environment is a software application that provides support to aid software development. A description of how to start using Eclipse and the JAS-mine plugin can be found at <http://www.jas-mine.net/how-to-create-and-run-a-new-jas-project-using-eclipse>.

⁷Technically, the scheduler is a "singleton". In software engineering, the singleton pattern is a design pattern that restricts the instantiation of a class to one object.

the ORM. When using Excel files to specify input data, no specific classes need to be included in this package.

- **model:** package containing the classes that specify the model structure; in particular, it contains the Model manager class(es) and the class(es) of agents that populate the simulation.
- **model.enums:** subpackage containing the definition of the enumerations used (if any).⁸
- **experiment:** package containing the classes that deal with running the simulation experiment(s); it contains, in particular, the Start class where the main method and the type of the experiment (interactive vs. batch mode, single run vs. multiple runs) are defined. The package might also contain one or more Collectors, who compute statistics and persist the output in the database, one or more Observers for online statistics collection and display, and a MultiRun class that manages repeated runs for parameter exploration.
- **algorithms:** package containing classes that implement algorithms for determining events and applying processes to the agents. These implementations, in a cooperative effort of users, are potential candidates to extend the set of standard functions included in the JAS-mine libraries.

In addition to sources, the project also contains two folders for data input-output. The input folder contains input data in excel or H2 embedded formats. The output folder contains the output of different simulation experiments. At the beginning of each run, JAS-mine creates a sub-folder that is labelled automatically⁹ with a copy of the input files plus an empty output database, with the same structure of the input database as defined by the annotations added to the model classes. Coherence between the input database (if any), the output database and the classes representing the agents in the simulation (known as *entity* classes) is guaranteed by the ORM.

By default, JAS-mine executes the simulations in embedded mode: the databases are modified directly by the JDBC driver included in JAS-mine.¹⁰ The standard database uses a H2 database engine. Other databases supporting embedding can be used, such as Microsoft Access, Hypersonic SQL, Apache Derby, etc.

⁸Enumerations specify a set of predefined values that a property can assume. These values might be categorical (strings, e.g. gender), quantitative (discrete numbers, e.g. age) or even objects with their set of characteristics and properties (e.g. a predefined set of banks to which a firm can be linked). The ORM detects that a value is an enumeration when the property is declared with the annotation *@Enumerated* (see Sect. 4.6.3.1). Through enumerations the ORM automatically manages reading/writing operations in both text and numerical format.

⁹The folder name can be modified dynamically through labels set by the user.

¹⁰A JDBC driver is a software component enabling a Java application to interact with a database.

4.5 The Dynamic Microsimulation Model

The model that we implement is inspired by the *Demo07* sample model included in LIAM2.¹¹ It features a population of 20,200 persons grouped in 14,700 households undergoing a number of demographic and other life course events on an annual basis between the years 2002–2061:

- **Birth:** all women aged between 15 and 50 (inclusive) in any simulation year can give birth to a child, with a probability which is year- and age-specific and is reported in the file *p_birth.xls*.
- **Education:** education (lower secondary, upper secondary or tertiary) is predetermined at birth. Individuals exit lower secondary education at age 16, upper secondary education at age 19, and tertiary education at age 24.
- **Exit from parental home:** individuals aged 24 or over who are not yet married leave their parental home to set up a new household.
- **Marriage:** all individuals aged 18 or over whose civil state is either single, divorced or widowed, are eligible for getting married in any given simulation year. The probability of marriage depends on age, gender and civil state, and is stored in the *p_marriage.xls* file. Given these probabilities, a subset of the unmarried population is sampled and those chosen are entered into the matching algorithm. Actual matching involves ordering all the females first; then starting with the top ranked female, all males are ordered and the best available male is matched. Then for the second ranked female, the remaining males are ordered and the best available male is matched, and so on until no more matches can be made (because there are either no more males or females to match). Females are ordered according to their age difference (in absolute value) with respect to the average age in the pool of females to be matched, $|age - \text{mean}(age)|$: females with an age closer to average ‘choose’ first, while older or younger females ‘choose’ later. For each female, males are ranked by looking at how their age and work status compares with the female’s age and work status: regression coefficients are stored in the *reg_marriage.xls* file. Matched couples form new households.
- **Divorce:** divorce probability depends on age difference between the partners, elapsed marriage duration, number of children and work status of both partners: regression coefficients are stored in the *reg_divorce.xls* file.

¹¹The model differs from the LIAM2 version in that it collapses the work states of unemployment and inactivity into a single non-employment state. This is done by removing the unemployment module from the corresponding LIAM2 simulation, with everything else staying the same. The change is motivated by the fact that the distinction between unemployment and inactivity was implemented in a very unnatural way in LIAM2, and did not affect any subsequent choice on the part of the agents.

- **Employment:** all individuals who are of working age (males: between 15 and 65; females: between 15 and 61) and whose previous work state was neither student nor retired are considered to be available to work. Conditional on this, individuals are employed with a probability which depends on age, lagged work state (either employed, unemployed or inactive), gender and marital status: regression coefficients are stored in the *reg_inwork.xls* file. The model does not distinguish between unemployment and non-employment.
- **Death:** death is also a probabilistic event, with year- and age-specific death probabilities contained in the files *p_death_m.xls* and *p_death_f.xls*, for males and females respectively.

The divorce and employment processes are subject to alignment. This is a technique widely used in (dynamic) microsimulation modelling to ensure that the simulated totals conform to some exogenously specified targets, or aggregate projections (Baekgaard 2002; Klevmarken 2002; Li and O’Donoghue 2014). Alignment targets (aggregate frequencies) for divorce and employment are stored in the *p_divorce.xls* and *p_inwork.xls* files respectively.

One important thing to note is that the processes to be aligned are executed at an individual level, while alignment always takes place at the population level. That is, individual outcomes or probabilities are determined for each individual based on the chosen econometric specification and the estimated coefficients. This in general leads to a mismatch between the simulated (provisional) totals and the aggregate targets, which can of course be assessed only at the population level. The alignment algorithm then directly modifies the individual outcomes or probabilities of occurrence.

The specific algorithm used in the LIAM2 implementation is called “Sorting by the difference between the predicted probability and a random number” (SBD), see Li and O’Donoghue (2014), and – though quite common in the microsimulation literature – can be criticised on many theoretical grounds, see Stephensen (2014). The JAS-mine *alignment* library implements it for completeness, though its use is deprecated. Here we use it to remain as close as possible to the original LIAM2 version (the reader does not need to understand precisely how it works).

Finally, note that, though agents’ interactions are limited to matching in the marriage market, the model contains all the basic ingredients of a standard agent-based model. Further “agentisation” could entail introducing more interaction in the labour market, or introducing competition for instance in residential locations.¹²

¹²The interested reader will find a JAS-mine implementation of the Schelling Segregation Model, with added microsimulation features for illustrative purposes (a dynamic population, with birth, ageing and death processes) in the demo section of the JAS-mine website (www.jas-mine.net/demo/extended-schelling).

4.6 The JAS-mine Implementation

The JAS-mine class structure of the *Demo07* model is organised as in Table 4.1.

The core of the simulation lies in the *model* package, which contains the classes *PersonsModel*, *Person* and *Household*. The *experiment* package contains the *StartPersons* class that specifies to run the simulation in interactive mode, the *PersonsCollector* class that collects all the statistics needed to monitor the simulation and updates the output database, and the *PersonsObserver* class that creates and manages the graphical object for runtime monitoring. Parameters and coefficients are loaded into the *Parameters* class in the *data* package. All filters used to filter collections are grouped in the *data.filters* subpackage. The categories used for gender, educational levels, civil state and work state are stored as Enums in the *model.enums* subpackage. Finally the *algorithms* package contains classes that perform technical tasks (in the example, *MapAgeSearch* searches age- and gender-specific values in a map of coefficients, with age and gender as keys). Classes in the *algorithms* package are meant to be of general use beyond the specific model being developed, and are candidates for inclusion in the core libraries in future releases of the platform.

The project is organised in the standard JAS-mine folder structure, as described in Table 4.2.

The Java classes are contained in the *src* folder. The *input* folder contains the MS Excel parameter files and the initial population, stored as an h2 database (*input.h2*). The *output* folder is initially empty. For each new simulation experiment, a new subfolder is created with the appropriate time stamp, so as to uniquely identify the experiment (e.g. 20141218151116, for experiments initiated on the 18th of

Table 4.1 Class structure

Package	Class
Algorithms	MapAgeSearch
Data	Parameters
data.filters	ActiveMultiFilter
	FemaleFilter
	FemaleToCoupleFilter
	FemaleToDivorce
	MaleFilter
	MaleToCoupleFilter
Experiment	PersonsCollector
	PersonsObserver
	StartPersons
Model	Household
	Person
	PersonsModel
model.enums	CivilState
	Education
	Gender
	WorkState

Table 4.2 File structure

Folder	Files	Notes
Input	p_birth.xls	Birth probabilities, by age and (simulated) year
	p_death_f.xls	Death probabilities, by age and (simulated) year, for females
	p_death_m.xls	Death probabilities, by age and (simulated) year, for males
	p_marriage.xls	Marriage probability, by age group, gender and civil state
	p_divorce.xls	Alignment target for the divorce probability, by age group and (simulated) year
	p_inwork.xls	Alignment target for the employment probability, by age group, gender and (simulated) year
	reg_marriage.xls	Marriage score coefficients: determine how well a specific male fits a given female
	reg_divorce.xls	Divorce coefficients: determine the (unaligned) probability of divorcing
	reg_inwork.xls	Employment coefficients: determine the (unaligned) probability of being employed
		input.h2
Output	(Empty)	
Src	(All java classes)	See Table 4.1
Target	(Compiled classes)	
Libs	(External libraries and JARs, empty)	
(Root)	pom.xml	Maven version control

December 2014, at 16 s after 3.11 pm). The subfolder contains a copy of all the input files (in the *input* directory) and an output database (out.h2, in the *database* directory).

The *target* and *libs* folders contain technical content of no immediate interest to the modeller. The root folder also contains a *pom* (project object model) file, which contains information on the JAS-mine version used for the project. Apache Maven, an open source software project management and comprehension tool uses this information to manage all the project's build, reporting and documentation. In particular, by specifying in the *pom* file the desired release for each library used (including the JAS-mine libraries), Maven automatically downloads the relevant libraries from the appropriate repositories.¹³

4.6.1 Parameters

As JAS-mine supports a clear distinction between modelling classes and data structures, parameters are loaded and stored in a specific class, called Parameters.

¹³This implies that each JAS-mine project has its own copy of all the libraries used, ensuring that the project is self-contained and that it keeps working exactly as intended even when new versions of the libraries are released (and even if backward compatibility is not respected).

The class makes use of the `ExcelAssistant.loadCoefficientMap()` method to read the parameters from MS Excel files: this requires to specify a .xls file, a worksheet name, the number of key columns and the number of value columns.¹⁴ Parameters are then stored in `MultiKeyCoefficientMap` objects, which are basically standard Java maps with multiple keys (Box 4.1).

Box 4.1 The `Parameters.loadParameters()` Method

```
public static void loadParameters() {

    // probabilities
    pBirth = ExcelAssistant.loadCoefficientMap("input/
    p_birth.xls", "Sheet1", 1, 59);

    pDeathM = ExcelAssistant.loadCoefficientMap("input/
    p_death_m.xls", "Sheet1", 1, 59);

    pDeathF = ExcelAssistant.loadCoefficientMap("input/
    p_death_f.xls", "Sheet1", 1, 59);

    pMarriage = ExcelAssistant.loadCoefficientMap(
    ("input/p_marriage.xls", "Sheet1", 3, 4);

    pDivorce = ExcelAssistant.loadCoefficientMap(
    ("input/p_divorce.xls", "Sheet1", 2, 59);

    pInWork = ExcelAssistant.loadCoefficientMap("input/
    p_inwork.xls", "Sheet1", 3, 59);

    // regression coefficients
    coeffMarriageFit = ExcelAssistant.loadCoefficientMap(
        "input/reg_marriage.xls", "Sheet1", 1, 1);

    coeffDivorce = ExcelAssistant.loadCoefficientMap(
        "input/reg_divorce.xls", "Sheet1", 1, 1);

    coeffInWork = ExcelAssistant.loadCoefficientMap(
        "input/reg_inwork.xls", "Sheet1", 3, 1);

    // definition of regression models
    regMarriageFit = new LinearRegression(coeffMarriageFit);

    regDivorce = new LogitRegression(coeffDivorce);

    regInWork = new LogitRegression(coeffInWork);

}
```

¹⁴It is also possible to load the parameters from a table in the input database. See the online documentation for further details.

There are two types of parameters in *Demo07*: probabilities and regression coefficients.

Birth (p_{Birth}) and death (p_{DeathM} and p_{DeathF}) probabilities have one key (age), while the value columns refer to different simulation years: birth and death probabilities are therefore age- and year-specific.

Divorce probabilities ($p_{Divorce}$) have two keys (the lower and upper bounds defining age groups), while value columns refer again to different simulation years: divorce probabilities are therefore age group- and year-specific.

Marriage ($p_{Marriage}$) and employment (p_{InWork}) probabilities have three keys (the lower and upper bounds defining age groups and gender). Value columns in $p_{Marriage}$ refer to different civil states: marriage probabilities are therefore age group-, gender- and civil state-specific. Value columns in p_{InWork} refer to different simulation years: employment probabilities are therefore age group-, gender- and year-specific.

Table 4.3 shows how the $p_{birth.xls}$ file looks like.

Regression coefficients can have one key ($coeff_{MarriageFit}$ and $coeff_{Divorce}$) which is the regressor variable name, and a corresponding value with the estimated coefficient. They might have additional keys, as in $coeff_{InWork}$, if the coefficients are differentiated by some other variables (gender and employment state, in this example). Table 4.4 shows what the corresponding $reg_{inwork.xls}$ file looks like.

Table 4.3 Extract from the $p_{birth.xls}$ file

	Simulation year		
Age	2002	...	2060
15	0.00068	...	0.00075
16	0.00186	...	0.00181
...			
50	0.00010	...	0.00021

Table 4.4 Extract from the $reg_{inwork.xls}$ file

Regressors	Gender	workState	Coefficients
Age	Male	Employed	-0.19660
isMarried	Male	Employed	0.18928
workIntercept	Male	Employed	3.55461
...			
Age	Male	NotEmployed	0.97809
workIntercept	Male	NotEmployed	-12.39108
...			
Age	Female	Employed	-0.27405
isMarried	Female	Employed	-0.09068
workIntercept	Female	Employed	3.64871
...			
Age	Female	NotEmployed	0.82176
isMarried	Female	NotEmployed	-0.55910
workIntercept	Female	NotEmployed	-10.48043
...			

Note that the name of the regressor variable must appear in the first column, as the regression classes expect it to be the first key in the `MultiKeyCoefficientMap` instance. The name of the headings for the additional key columns must match the name of a field in the relevant class, in this case, the `Person` class.

The appropriate regression models are then defined based on the regression coefficients.

4.6.2 The `PersonsModel` Class

4.6.2.1 Objects

The `Model` extends the `AbstractSimulationManager` class. This requires implementing the `buildObjects()` and the `buildSchedule()` methods. The `buildObjects()` method contains the instructions to create all the agents and the objects that represent the virtual environment for model execution (see Box 4.2).¹⁵ In *Demo07*, this involves loading the parameters for the simulation and the initial population, made of persons and households. Three other methods complete the simulation setup: `initializeNonDatabaseAttributes()` initializes attributes that do not appear in the input database, such as the education level; `addPersonsToHouseholds()` registers household members, and `cleanInitialPopulation()` checks the internal consistency of the initial population and removes errors, making sure that all marriage partnerships are bilateral, that all partners belong to the same household, and that no empty households exist.¹⁶

Box 4.2 The `PersonsModel.buildObjects()` Method

```
@Override
public void buildObjects() {

    Parameters.loadParameters();

    persons = (List<Person>) DatabaseUtils.loadTable
        (Person.class);

    households = (List<Household>)
        DatabaseUtils.loadTable(Household.class);

    initializeNonDatabaseAttributes();
}
```

(continued)

¹⁵The `@Override` annotation is used by the Java interpreter to ensure that the programmer is aware that the method declared is overriding the same method declared in the superclass.

¹⁶This method is absent in the `LIAM2` implementation, which does not get rid of all the errors in the initial database.


```

addPersonsToHouseholds();

cleanInitialPopulation();

}

```

As we have seen, the general rule is that parameters should not be hard-coded in the simulation. The only exception is with *control parameters* that can be changed from the GUI before the simulation starts or while the simulation is running in order to experiment with the model behaviour in interactive mode. Control parameters are properties of a simulation, they are annotated with *GUIparameter*, are automatically loaded into the JAS-mine GUI, and are automatically saved in a separate table of the output database. In *Demo07* there are just three such parameters, as described in Box 4.3.

Box 4.3 PersonsModel: Control Parameters

```

GUIparameter(description="Simulation begins at
year [valid range 2002-2060]")
private Integer startYear = 2002;

GUIparameter(description="Simulation ends at year
[valid range 2003-2061]")
private Integer endYear = 2061;

GUIparameter(description="Retirement age for women")
private Integer wemra = 61;

```

4.6.2.2 Schedule

The *buildSchedule()* method contains the plan of events for the simulation. Events are planned based on a discrete event simulation paradigm. This means that events can be scheduled dynamically at specific points in time. The frequency of repetition of an event can be specified in the case of periodic events. An event can be created and managed by the simulation engine (a system event e.g. terminating the simulation), it can be sent to all the components of a collection or list of agents or it can be sent to a specific object/instance. Events can be grouped together if they share the same schedule.

In *Demo07*, all events are scheduled right from the beginning of the simulation (there is no dynamic scheduling), and occur on a yearly basis. They are grouped in an EventGroup called *modelSchedule*, which is scheduled at every simulation

period starting at `startYear` with the `scheduleRepeat(Event event, double atTime, int withOrdering, double timeBetweenEvents)` method:

```
getEngine().getEventList().scheduleRepeat(modelSchedule,
startYear, 0, 1.);
```

The events of *Demo07* are typically directed to a collection of objects – persons or households – and are inserted into an `EventGroup` with the instruction

```
modelSchedule.addCollectionEvent(Object object, [some action the
object must perform]);
```

The actions to be performed can be specified in two ways. The simplest is to use Java reflection and simply specify the object’s method name to be invoked.¹⁷ For instance, asking all persons to perform the `ageing()` method would require the instruction:

```
modelSchedule.addCollectionEvent(persons, Person.class,
"ageing");
```

Java reflection, however, generally has a reputation for being quite slow. A better approach is to use the `EventListener` interface. When an object implements this interface, it must define an `onEvent()` method that will receive specific enumerations to be interpreted. We will describe how the `Person` class implements the `onEvent()` method in Sect. 4.6.3.3. For now, we simply note that by using the `EventListener` interface, the scheduling of the `ageing()` method becomes:

```
modelSchedule.addCollectionEvent(persons, Person.Processes.
Ageing);
```

By default, the broadcasting of an event to a collection of objects is performed in *safe mode* (read only), and does not allow the concurrent modification of the collection itself. This is not a problem with the `ageing()` process, as ageing per se does not entail any modification in the list of persons, that is, it does not add or remove anyone. This is not true with other processes, like `birth()` or `death()`. In order to allow the collection to be changed while iterated by the simulation engine, this feature has to be switched off, as in

```
modelSchedule.addCollectionEvent(persons, Person.Processes.Death,
false);
```

The last argument specifies that the collection is subject to changes while being iterated, and the JAS-mine engine treats it accordingly.

The order of the events in the simulation follows the original LIAM2 implementation and is specified in Box 4.4: there is first a set of demographic events (ageing, death, birth, marriage, exit from parental home, divorce, household composition) and then a set of events that define the work status (whether in education, retired, other non-employed, or employed).

¹⁷This requires the method – `ageing()` in this case – to be declared public.

Box 4.4 The PersonsModel.buildSchedule() Method

```

@Override
public void buildSchedule() {

    EventGroup modelSchedule = new EventGroup();
    // 1: Ageing
    modelSchedule.addCollectionEvent(persons, Person.Processes.
    Ageing);

    // 2: Death
    modelSchedule.addCollectionEvent(persons, Person.Processes.
    Death, false);

    // 3: Birth
    modelSchedule.addCollectionEvent(persons, Person.Processes.
    Birth, false);

    // 4: Marriage
    modelSchedule.addCollectionEvent(persons, Person.Processes.
    ToCouple);

    modelSchedule.addEvent(this, Processes.MarriageMatching);

    // 5: Exit from parental home
    modelSchedule.addCollectionEvent(persons, Person.Processes.
    GetALife);

    // 6: Divorce
    modelSchedule.addEvent(this, Processes.DivorceAlignment);

    modelSchedule.addCollectionEvent(persons, Person.
    Processes.Divorce);

    // 7: Household composition
    // (for reporting only: household composition is
    // updated whenever
    // needed throughout the simulation)
    modelSchedule.addCollectionEvent(households,
    Household.Processes.HouseholdComposition);

    // 8: Education
    modelSchedule.addCollectionEvent(persons, Person.Processes.
    InEducation);

    // 9: Work
    modelSchedule.addEvent(this, Processes.InWorkAlignment);

    getEngine().getEventList().schedule(modelSchedule, 0, 1);

    getEngine().getEventList().schedule(

```

(continued)

```

new SingleTargetEvent(this, Processes.Stop),
endYear - startYear);

}

```

Marriage is performed in two steps. First, a subset of suitable males and females are selected for matching by invoking the method `Person.toCouple()`¹⁸; then, matching takes place. As we have seen in Sect. 4.5, matching uses a “centralised” algorithm and is therefore performed by the Model itself. Consequently, this event is a single target event, rather than a collection event, and is inserted into our EventGroup `modelSchedule` with the instruction

```
modelSchedule.addEvent(this, Processes.MarriageMatching);
```

Similarly, the divorce and work events are subject to alignment and are managed directly by the Model, with the methods `divorceAlignment()` and `inWorkAlignment()`, though divorce also requires some actions taken by the individuals themselves – in the `divorce()` method in the Person class – after they have been selected to divorce. `householdComposition()` is the only method which is directed to the collection of households. It simply updates the number of adults and children in each household for reporting purposes. A final single target event is scheduled for the last year of the simulation with the method `scheduleOnce(Event event, double atTime, int withOrdering)`: its target is the Model itself and brings the simulation to a halt:

```

getEngine().getEventList().scheduleOnce
(new SingleTargetEvent(this, Processes.Stop), endYear,
Order.AFTER_ALL.getOrdering());

```

4.6.2.3 The EventListener Interface

Since the Model performs actions during the simulation, as with the Person and Household classes, it implements the EventListener interface. This requires first to enumerate all the actions that the Model is supposed to perform (this is done by defining the specific Enum Processes), and then to specify the method `onEvent()` – see Box 4.5.

¹⁸See Sect. 4.6.3.3.

Box 4.5 Implementation of the EventListener Interface in PersonsModel

```

public enum Processes {
    MarriageMatching,
    DivorceAlignment,
    InWorkAlignment,
    Stop;
}

@Override
public void onEvent(Enum<?> type) {
    switch ((Processes) type) {
        case DivorceAlignment:
            divorceAlignment();
            break;
        case InWorkAlignment:
            inWorkAlignment();
            break;
        case MarriageMatching:
            marriageMatching();
            break;
        case Stop:
            log.info("Model completed");
            getEngine().pause();
            break;
    }
}

```

We now dig into the matching and alignment methods performed by the Model.

4.6.2.4 The Matching Algorithm

Prior to matching, a sample of the population to marry at this time is determined randomly using the `Person.toCouple()` method. Subsequently, matching involves first ordering all the females; then, for each female starting from the top of the ranking, all males are ordered and the most suitable male is matched. This continues until there are either no more females or males to match. Females are ordered according to their age difference (in absolute value) with respect to the average age in the pool of females to be matched, $|age - \text{mean}(age)|$; the female whose age is closest to the average is ranked first. To compute this ranking, the average age of the subset of females selected for matching is required. There are a number of ways to perform this computation, which is preliminary to the application of the

matching algorithm. The one that is implemented in *Demo07* makes use of Java closures (Box 4.6).¹⁹

Box 4.6 Computing the Average Age for the Eligible Females in Persons-Model.marriageMatching()

```
final AverageClosure averageAge =
new AverageClosure() {

    @Override
    public void execute(Object input) {
        add(((Person) input).getAge());
    }
};

Aggregate.applyToFilter(getPersons(),
new FemaleToCoupleFilter(), averageAge);
```

The JAS-mine *collection* package defines an `AverageClosure` as a closure that receives values from objects as an input and returns the mean of these values as an output. Here, it is used to compute the average age of a given set of persons. The set is defined by applying the `FemaleToCouple` filter to the list of all persons, with the instruction

```
Aggregate.applyToFilter(getPersons(), new FemaleToCoupleFilter(),
averageAge);
```

The `averageAge` closure now contains the average age of all filtered females. In turn, the `FemaleToCouple` filter simply selects the female persons who have the `toCouple` flag switched on (Box 4.7).

Box 4.7 The FemaleToCouple Filter

```
public class FemaleToCoupleFilter
implements Predicate {

    @Override
    public boolean evaluate(Object object) {
        Person agent = (Person) object;
```

(continued)

¹⁹Technically, a *closure* is a function that refers to free variables in their lexical context. A free variable is an identifier (the identity of the person which is included in the evaluation set, in our example) that has a definition outside the closure: it is not defined by the closure, but it is used by the closure. In other words, these free variables inside the closure have the same meaning they would have had outside the closure.

```

        return (agent.getGender().equals(Gender.Female)
&& agent.getToCouple());
    }
}

```

Having the filters specified as separate classes, grouped in the separate package *data.filters*, might look cumbersome at first (and there are other ways to do this, see the online documentation) but allows to keep the core code clean while using the standard Apache Predicate approach to filtering – remember that the JAS-mine approach supports the use of existing software solutions whenever possible, and envisages to keep the specificities of the JAS-mine libraries to a minimum in order to minimise the “black box” feeling of many simulation platforms.

Matching is then performed, following the LIAM2 implementation, by making use of a simple one-way matching procedure (the agents in one collection – females in our example – choose, while the agents in the other collection – males – remain passive) implemented in the SimpleMatching class:

```

matching(collection1, filter1, comparator1, collection2,
filter2, matchingScoreClosure, matchingClosure);

```

and it is invoked as

```

SimpleMatching.getInstance().matching(...);

```

The matching method requires seven arguments:

1. **collection1**: the first collection (e.g. all individuals in the population);
2. **filter1**: a filter to be applied to the first collection (e.g. all females with the *toCouple* flag on);
3. **comparator1**: a comparator to sort the filtered collection, which determines the order that the agents in the filtered collection will be matched.
4. **collection2**: the second collection, which can be the same as *collection1* (e.g. all individuals in the population) or a different one; the two collections do not need to have the same size;
5. **filter2**: a filter to be applied to the second collection (e.g. all males with the *toCouple* flag on);
6. **matchingScoreClosure**: a piece of code that assigns, for every element of the filtered *collection1*, a double value to each element of the filtered *collection2*, as a measure of the quality of the match between every pair;
7. **matchingClosure**: a piece of code that determines what to do upon matching.

As in the computation of the average age, the use of closures – which are relatively new to the Java language – allows a great simplification of the code. While it is not required that the user knows about closures, it is interesting to understand why they are so useful. In the example, suppose that the females in the population are sorted according to some criterion, for example beauty: the prettiest woman is the first to choose a partner, the second prettiest is the second to choose etc. The

matchingScoreClosure sorts all possible mates according to some other criterion, for example wealth. Hence, the prettiest woman gets the richest man, the second prettiest gets the second richest, etc. In such a case, a comparator would suffice to order the males in the population, as the ranking is the same irrespective of the female who is evaluating them. But suppose now that the attractiveness of a man depends on the age differential between himself and the potential partner: in such a case, the ranking is specific to each woman in the population. A simple comparator would still do the job, but the comparator should be able to access the identity of the woman who is making the evaluation as an argument, which requires a lot of not-so-straightforward coding. Closures allow to bypass this technical requirement because they can pass a functionality as an argument to another method; in other words, they treat functionality as method argument, or code as data.

Closures in the *matching()* method are easier to understand when illustrated by an example: the seven arguments are listed in Box 4.8.

Box 4.8 The Matching Algorithm in *PersonsModel.marriageMatching()*

```
SimpleMatching.getInstance().matching(
    // collection1: the whole population
    persons,

    // filter1:
    new FemaleToCoupleFilter(),

    // comparator1: a comparator that assigns priority to the
    // individual that has a lower difficulty in matching
    // (this is determined by an individual's age in relation
    // to the average)
    new Comparator<Person>() {

        @Override
        public int compare(Person female1, Person female2) {
            return (int) Math.signum(
                Math.abs(female1.getAge() -
                    averageAge.getAverage()) -
                Math.abs(female2.getAge() -
                    averageAge.getAverage()));
        }
    },

    // collection2: same as collection1
    persons,

    // filter2:
    new MaleToCoupleFilter(),

    // MatchingScoreClosure: a closure that, given a specific
```

(continued)


```

// female,
// computes for every male in the population a matching score
new MatchingScoreClosure<Person>() {

    @Override
    public Double getValue(Person female, Person
male) {
        return female.getMarriageScore(male);
    }
},
// matchingClosure: a closure that creates a link between a
// specific female and a specific male, and sets up a new
// household.
new MatchingClosure<Person>() {

    @Override
    public void match(Person female, Person male) {

        female.marry(male);
        male.marry(female);
    }
}
);

```

4.6.2.5 Alignment

Alignment involves comparing the *provisional* outcomes of the simulation with some external aggregate targets, and then modifying the simulation outcomes in order to match the external totals. We show how this is implemented in *Demo07* by looking at the *divorceAlignment()* method; the *inWork()* alignment method works similarly. When it comes to divorce, as in *marriageMatching()*, the focus is on females: males are passive recipients of their partners' choices. Different targets are specified for different age groups and simulated years; as we have seen in Sect. 4.6.1, these are read from the file *p_divorce.xls* and stored in the `MultiKeyCoefficientMap` *pDivorce* in the `Parameters` class. The *divorceAlignment()* method works cell by cell, that is, it aligns each age group of the population to its year-specific target: this means that the alignment algorithm is applied once for every age group (as defined in the *p_divorce.xls* parameter file). The structure of the method is therefore as follows:

- For each age group: do alignment:
 - Read target from *pDivorce*.
 - Select the relevant subgroup of married females.

- Compute, for each of the selected females, a probability to divorce that depends on the age group to which they belong.
- Select the couples that divorce by applying the SBD algorithm: each female is ranked according to the signed difference between their divorce probability and a random number uniformly distributed between 0 and 1; then, the number of couples equal to the target are selected to divorce by starting with the top ranked female and going down the ranks until the target number is reached.²⁰

The `MultiKeyCoefficientMap` *pDivorce*, which contains the targets, has a three dimensional key: the lower and upper bounds for the age group, and the year of the simulation. The age group-specific and year-specific targets are read with the instruction reported in Box 4.9.

Box 4.9 `PersonsModel.divorceAlignment()`: Reading the Targets

```
MultiKeyCoefficientMap pDivorceMap =
Parameters.getpDivorce();

for (MapIterator iterator =
pDivorceMap.mapIterator(); iterator.hasNext();) {

    iterator.next();
    MultiKey mk = (MultiKey) iterator.getKey();
    int ageFrom = (Integer) mk.getKey(0);
    int ageTo = (Integer) mk.getKey(1);
    double divorceTarget = ((Number)
pDivorceMap.getValue(
    ageFrom,
    ageTo,
    getStartYear() + SimulationEngine.
getInstance().getTime())) .doubleValue();
    [...]
}
```

The alignment methods require four arguments:

1. **collection**: a collection of individuals whose outcome or probability of an event has to be aligned (e.g. all the population);
2. **filter**: a filter to be applied to the collection (e.g. all females selected to divorce);
3. **alignmentProbabilityClosure** or **alignmentOutcomeClosure**: a piece of code that i) computes for each element of the filtered collection a probability for the

²⁰The ranking involves a stochastic component (the random number that is subtracted from the divorce probability score) in order to give individuals with a low predicted probability some chance to experience the event. As we have already noted, the SBD algorithm is quite distortive and its use is deprecated in JAS-mine; it is employed here only for consistency with the LIAM2 implementation.

event (in the case that the alignment method is aligning probabilities, as in the SBD algorithm) or an outcome (in the case that the alignment method is aligning outcomes), and ii) applies to each element of the filtered collection the specific instructions coming from the alignment method used;

4. **targetShare** or **targetNumber**: the share or number of elements in the filtered collection that are expected to experience the transition; the SBD algorithm uses *targetShare*.

Box 4.10 shows how the alignment method is implemented in *Demo07*.

Box 4.10 *PersonsModel.divorceAlignment()*: Applying the SBD Alignment Algorithm

```

new SBDAlignment<Person>().align(

    // collection:
    persons,

    // filter:
    new FemaleToDivorce(ageFrom, ageTo),

    // alignmentProbabilityClosure:
    new AlignmentProbabilityClosure<Person>() {

        // i) compute the probability of divorce
        @Override
        public double getProbability(Person agent) {
            return agent.computeDivorceProb();
        }

        // ii) determine what to do with the aligned
        probabilities
        @Override
        public void align(Person agent,
            double alignedProbability)
        {
            boolean divorce = RegressionUtils.event(
                alignedProbability,
                SimulationEngine.getRnd()
            );

            agent.setToDivorce(divorce);
        }
    },

    // targetShare:
    divorceTarget

);

```

4.6.3 The Person Class

4.6.3.1 Entities

The Person class is an Entity class, as specified by the `@Entity` annotation:

```
@Entity
public class Person implements Comparable<Person>,
EventListener, IDoubleSource {
    [...]
}
```

This implies that the class is linked to a table in the database with the same name, and that all properties which are not annotated as `@Transient` are persisted in the database, when the simulation output is saved.

Entity classes must specify a `PanelEntityKey` (annotated as `@Id`), which is a three-dimensional object which identifies the agent id, the simulation time and the simulation run. These three keys uniquely identify each record in the database:

```
@Id
private PanelEntityKey id;
```

The ORM expects that the field names in the database are the same as the property names in the Java class, except when a different name is specified as in

```
@Column(name="dur_in_couple")
private Integer durationInCouple;
```

Enumerations can be interpreted by the ORM both as a string and as ordinal values (0 for the first *enum*, 1 for the second, etc.), depending on how they are annotated:

```
@Enumerated(EnumType.STRING)
private WorkState workState;
```

4.6.3.2 The IDoubleSource Interface

The Person class implements the `IDoubleSource` interface. This interface provides a simple way of asking a class to return a specific value.²¹ Similarly to the `EventListener` interface, it requires to declare an `Enum` which lists all the variables that can be queried, and the `getDoubleValue()` method for returning their value (Box 4.11). It is used by the Regression classes as a way of decoupling the regression model specification from the code: as long as a variable is enumerated in the specific

²¹As such, it is also used by JAS-mine distribution plots, see Sect. 4.6.6.

Enum called `Regressors`, it can be used (or removed) as a covariate in a regression model without the need to modify the code.²²

Box 4.11 Implementation of the `IDoubleSource` Interface in `Person`

```
public enum Regressors {

    // For marriage regression, check with potential
    // partner's properties
    potentialPartnerAge,
    potentialPartnerAgeSq,
    potentialAgeDiff,
    inWorkAndPotentialPartnerInWork,
    notInWorkAndPotentialPartnerInWork,
    ...

    // For in work regression
    age,
    ageSq,
    ageCub,
    isMarried,
    workIntercept;

}

public double getDoubleValue(Enum<?> variableID) {

    switch ((Regressors) variableID) {

        //For marriage regression
        case potentialPartnerAge:
            return getPotentialPartnerAge();
        case potentialPartnerAgeSq:
            return getPotentialPartnerAge() *
                getPotentialPartnerAge();

            ...

        //For work regression
        case age:
            return (double) age;
        case ageSq:
            return (double) age * age;
        case ageCub:

```

(continued)

²²Regression classes also have a method to read directly the values of the variables from the agent class, without the need of implementing the `IDoubleSource` interface. However, this requires that all the variables used by a regression model are defined as (possibly transient) properties in the class. This is particularly tedious when the covariates refer to another agent (such as a potential partner, or the spouse), as is common in our case.

```

        return (double) age * age * age;
    case isMarried:
        return civilState.equals(CivilState.Married) ?
            1.0 : 0.0;
    case workIntercept:
        return 1.0; //Constant intercept, multiply
// regression coefficient by 1
    default:
        throw new IllegalArgumentException(
            "Unsupported regressor " +
            variableID.name() + " in
            Person#getDoubleValue");
    }
}

```

4.6.3.3 Methods

The Person class implements the EventListener interface and is therefore able to be activated by the scheduler with the *onEvent()* method. The calls that a Person is able to respond to – enumerated in a specific Enum called Processes (Box 4.12) – are:

- **Ageing:** age and marriage duration are increased; work status is set to retired if retirement age is reached.
- **Death:** an age-, gender- and year-specific death probability is read from the MultiKeyCoefficientMaps *pDeathM* and *pDeathF* stored in the Parameters class; this probability is then compared with a uniformly distributed random number between 0 and 1 to determine the occurrence of the event:

```
RegressionUtils.event(deathProbability);
```

If death occurs, the partner's status is updated to widow and the person is removed from all the lists (that is, from his/her household and from the model).

- **Birth** (applied to all females aged between 15 and 50 inclusive): an age- and year-specific probability of having a baby is read from the MultiKeyCoefficientMap *pBirth* stored in the Parameters class; then the occurrence of the event is determined in a similar fashion to the *death()* process. No multiple births such as twins can occur. Newborns are given a potential educational level that will be reached with certainty. Following the LIAM2 implementation, the person is assumed to be a student until completion of their studies (at age 16 for lower secondary education, 19 for upper secondary education, and 24 for tertiary education).

- **ToCouple** (applied to all unmarried individuals aged between 18 and 90 inclusive): this method reads an age-, gender- and civil state-specific probability of forming a partnership from the `MultiKeyCoefficientMap pMarriage` and determines whether the Boolean flag `toCouple` is set to true, to be used by the `marriageMatching()` algorithm in the `PersonsModel` class.
- **GetALife** (leave parental home): a new household is created if the individual is aged 24 or over, unmarried and still living in the parental household.
- **Divorce**: after divorce is decided by the `Model`'s alignment method, partner links are broken, civil states are updated, females retain their household and males move to a newly created household.
- **InEducation**: this method examines the person's age and education level to determine whether an individual is still in education, or must exit education and enter the labour market as unemployed.

Box 4.12 The `Person.Processes` Enum, Defining the Processes a Person Undertakes When Activated by the Scheduler

```
public enum Processes {
    Ageing,
    Death,
    Birth,
    ToCouple,
    Divorce,
    GetALife,
    InEducation;
}
```

Other significant methods of the `Person` class include:

- **getMarriageScore()**: computes the score of each male in the marriage pool, for a given female, based on a linear regression model specified by the `MultiKeyCoefficientMap regMarriageFit`; it is used by the `marriageMatching()` method in the `PersonsModel` class.
- **marry()**: creates a link between the two partners and sets up a new household where they move to; it is used by the `marriageMatching()` method in the `PersonsModel` class.
- **computeDivorceProb()**: computes the divorce probability, based on a logit regression model specified by the `MultiKeyCoefficientMap regDivorce`; it is used by the `divorceAlignment()` method in the `PersonsModel` class.
- **computeWorkProb()**: computes the employment probability, based on a logit regression model specified by the `MultiKeyCoefficientMap regInWork`; it is used by the `inWorkAlignment()` method in the `PersonsModel` class.

Given that the regression coefficients have already been loaded from Excel files into the Parameters class, and the IDoubleSource interface method *getDoubleValue()* takes care of reading the values of the regressor variables, the simulation of outcomes or probabilities based on regression models is straightforward:

```

    marriageScore = Parameters.getRegMarriageFit().
getScore(this, Person.Regressors.class);
    divorceProb = Parameters.getRegDivorce().
getProbability(this, Person.Regressors.class);
    workProb = Parameters.getRegInWork().
getProbability(this, Person.Regressors.class);

```

Again, if the specification of the model is changed by adding or removing covariates, or if new coefficient estimates become available, nothing has to be changed in the code, except for adding any new covariate to the Person.Regressors Enum and providing a method for the new case in the *getDoubleValue()* method.²³

4.6.4 The Household Class

This class contains a list of all household members and is able to count the number of adults and children in the household. It is defined as an Entity class and is therefore linked to a table with the same name in the database. It implements the EventListener interface because it responds to calls by the scheduler requesting that the household composition is updated.

4.6.5 The PersonsCollector Class

The Collector collects statistics and manages the persistence of the simulation outputs on the database. It extends the AbstractSimulationCollectorManager class and requires, similarly to the Model, the implementation of a *buildObjects()* method and a *buildSchedule()* method.

The *buildObjects()* method creates several CrossSection objects, which collect specific values from each individual in the population (Box 4.13).

²³The change in specification is instead achieved by updating the regression coefficient input files (e.g. *reg_inwork.xls*).

Box 4.13 The PersonsCollector.buildObjects() Method

```

@Override
public void buildObjects() {

    final PersonsModel model =
        (PersonsModel) getManager();

    ageCS = new CrossSection.Integer(model.getPersons(),
        Person.class, "age", false);

    nonEmploymentCS = new CrossSection.Integer(
        model.getPersons(), Person.class, "getNonEmployed",
        true);

    employmentCS = new CrossSection.Integer(
        model.getPersons(), Person.class, "getEmployed", true);

    retiredCS = new CrossSection.Integer(
        model.getPersons(), Person.class, "getRetired", true);

    inEducationCS = new CrossSection.Integer(
        model.getPersons(), Person.class, "getStudent", true);

    lowEducationCS = new CrossSection.Integer(model.
        getPersons(), Person.class, "getLowEducation", true);

    midEducationCS = new CrossSection.Integer(model.
        getPersons(), Person.class, "getMidEducation", true);

    highEducationCS = new CrossSection.Integer(model.
        getPersons(), Person.class, "getHighEducation", true);

}

```

The Collector's schedule is made up of two processes only, which take place at every simulation period: the CrossSections are updated (*Processes.Update*), and the persons and households are persisted in the database (*Processes.DumpInfo*) (see Box 4.14).

Box 4.14 The PersonsCollector.buildSchedule() Method

```

@Override
public void buildSchedule() {

    EventGroup collectorSchedule = new EventGroup();

```

(continued)

```

collectorSchedule.addEvent(this, Processes.Update);
collectorSchedule.addEvent(this, Processes.DumpInfo);

getEngine().getEventList().schedule(collectorSchedule, 0, 1);
}

```

The Collector also implements the `EventListener` interface, featuring the Enum `Processes` and `onEvent()` method (see Box 4.15).

Box 4.15 Implementation of the EventListener Interface in PersonsCollector

```

public enum Processes {
    Update,
    DumpInfo;
}

@Override
public void onEvent(Enum<?> type) {
    switch ((Processes) type) {
        case Update:
            ageCS.updateSource();
            nonEmploymentCS.updateSource();
            employmentCS.updateSource();
            retiredCS.updateSource();
            inEducationCS.updateSource();
            lowEducationCS.updateSource();
            midEducationCS.updateSource();
            highEducationCS.updateSource();
            break;

        case DumpInfo:
            try {
                DatabaseUtils.snap(((PersonsModel) getManager()).
                    getPersons());

                DatabaseUtils.snap(((PersonsModel) getManager()).
                    getHouseholds());

            } catch (Exception e) {
                log.error(e.getMessage());
            }
            break;
    }
}

```

As we have seen in Sect. 4.6.2, implementing the `EventListener` interface is not necessary, as the class can be activated by the scheduler using Java reflection. However, grouping all the updating in one single *Update* process improves on clarity.²⁴

Updating the `CrossSection` objects only involves simple instructions such as

```
ageCS.updateSource();
```

Similarly, dumping the simulation outputs is done by the *DumpInfo* process and only requires

```
DatabaseUtils.snap(((PersonsModel) getManager()).getPersons());
DatabaseUtils.snap(((PersonsModel) getManager()).getHouseholds());
```

4.6.6 The *PersonsObserver* Class

The `PersonsObserver` builds graphical objects that allow monitoring and inspection of the simulation outcome in real time. It extends the `AbstractSimulationObserverManager` interface and, similarly to the other simulation managers (the `Model` and the `Collector`), requires the implementation of a *buildObjects()* method and a *buildSchedule()* method.

The *buildObjects()* method creates three plots. The first one (*agePlotter*) depicts the evolution of the average age of the simulated population: it takes the *ageCS* `CrossSection` from the `Collector`, with information on the age of each individual, and computes its mean (by creating a `MeanArrayFunction` object). Similarly, the *workPlotter* plots the frequency of students, retired, other non-employed and employed individuals in the population, and the *eduPlotter* plots the share of individuals with each educational level (Box 4.16).

Box 4.16 The `PersonsObserver.buildObjects()` Method

```
@Override
public void buildObjects() {

    final PersonsCollector collector = (PersonsCollector)
        getCollectorManager();

    agePlotter = new TimeSeriesSimulationPlotter
        ("Age", "age(years)");
```

(continued)

²⁴Because updating is a common activity, it is also defined as a *CommonEventType* Enum in the JAS-mine *event* library (together with saving). Passing the scheduler this Enum does not require implementing the `EventListener` interface. An example of this approach is implemented in the `Observer`.

```

agePlotter.addSeries("avg",
    new MeanArrayFunction(collector.getAgeCS()));
GuiUtils.addWindow(agePlotter, 250, 50, 500, 500);

workPlotter = new TimeSeriesSimulationPlotter
("Work status", "proportion");
workPlotter.addSeries("employed",
    new MeanArrayFunction(collector.getEmploymentCS()));
workPlotter.addSeries("non-employed",
    new MeanArrayFunction(collector.getNonEmploymentCS()));
workPlotter.addSeries("retired",
    new MeanArrayFunction(collector.getRetiredCS()));
workPlotter.addSeries("students",
    new MeanArrayFunction(collector.getInEducationCS()));
GuiUtils.addWindow(workPlotter, 750, 50, 500, 500);

eduPlotter = new TimeSeriesSimulationPlotter("Education
level", "proportion");
eduPlotter.addSeries("low",
    new MeanArrayFunction(collector.getLowEducationCS()));
eduPlotter.addSeries("mid",
    new MeanArrayFunction(collector.getMidEducationCS()));
eduPlotter.addSeries("high",
    new MeanArrayFunction(collector.getHighEducationCS()));
GuiUtils.addWindow(eduPlotter, 1250, 50, 500, 500);
}

```

Other plots can be easily added. In particular, by building on the JFreeChart library, the `CollectionBarSimulationPlotter` class in JAS-mine allows to create histograms for representing distributions of given variables in the simulated population, at any given simulation period.

The schedule of the `PersonsObserver` class manages the updating of these three plots (Box 4.17). Here, the built-in JAS-mine Enum `CommonEventType.Update` is used, rather than a class-specific implementation of the `EventListener` interface as in the `Collector`. This requires scheduling the update of each graph separately, but allows for a better control of the display frequency. The latter is obtained by means of an extra parameter which is loaded into the GUI:

Box 4.17 The `PersonsObserver.buildSchedule()` Method

```

@GUIparameter
private Integer displayFrequency = 1;

@Override
public void buildSchedule() {

```

(continued)

```

getEngine().getEventList().schedule(new
    SingleTargetEvent(agePlotter,
        CommonEventType.Update), 0,
    displayFrequency);
getEngine().getEventList().schedule(new
    SingleTargetEvent(workPlotter,
        CommonEventType.Update), 0,
    displayFrequency);
getEngine().getEventList().schedule(new
    SingleTargetEvent(eduPlotter,
        CommonEventType.Update), 0,
    displayFrequency);

}

```

4.6.7 The StartPersons Class

The Start class initialises the JAS-mine simulation engine and defines the list of models to be used. In general, the Start class is designed to handle two types of experiments:

- performing a single run of the simulation in **interactive mode**, through the creation of a Model and related Collectors and Observers, with their GUIs;
- performing a single run of the simulation in **batch mode**, through the creation of the Model and possibly the Collectors; this involves managing parameter setup, model creation and execution directly, and is aimed at capturing only the simulation's numerical output;

The Start class is ignored when performing a **multi-run session** (whose structure is defined in a class extending the MultiRun interface) where the simulation is run repeatedly using different parameter values, so as to explore the space of solutions and produce sensitivity analyses on the specified parameters.

The Start class implements the ExperimentBuilder interface, which defines the *buildExperiment()* method. This method creates managers and adds them to the JAS-mine engine. In *Demo07*, the simulation is run in interactive mode (Box 4.18).

Box 4.18 The StartPerson Class

```

public class StartPersons implements
ExperimentBuilder {

public static void main(String[] args) {

```

(continued)

```

        boolean showGui = true;

        StartPersons experimentBuilder = new StartPersons();

        final SimulationEngine engine =
            SimulationEngine.getInstance();
        MicrosimShell gui = null;
        if (showGui) {
            gui = new MicrosimShell(engine);
            gui.setVisible(true);
        }

        engine.setExperimentBuilder(experimentBuilder);

        engine.setup();

    }

    @Override
    public void buildExperiment(SimulationEngine engine)
    {

        PersonsModel model = new PersonsModel();
        PersonsCollector collector = new PersonsCollector(model);
        PersonsObserver observer = new
            PersonsObserver(model, collector);

        engine.addSimulationManager(model);
        engine.addSimulationManager(collector);
        engine.addSimulationManager(observer);

    }

}

```

The `Start` class contains the standard `main()` method which allows a Java application to run. By selecting the “Run As Java Application” option from the Eclipse menu, this procedure launches the JAS-mine GUI, creates a model instance and gives it to the simulation engine. It then creates a Collector and an Observer and calls the `setup()` method of the simulation engine, which has the task of loading the experiment into memory.

The JAS-mine GUI contains a mask for setting the specific Model parameters, another mask for defining the specific Observer parameters and the three dynamic graphs defined in the Observer class. Figure 4.1 depicts the graphical output of one simulation run.

The Tools > ‘Database explorer’ tab in the JAS-mine GUI allows to browse the input and output databases. By selecting a specific database and pressing the ‘Show database’ button, the data can be explored in the default web browser using SQL

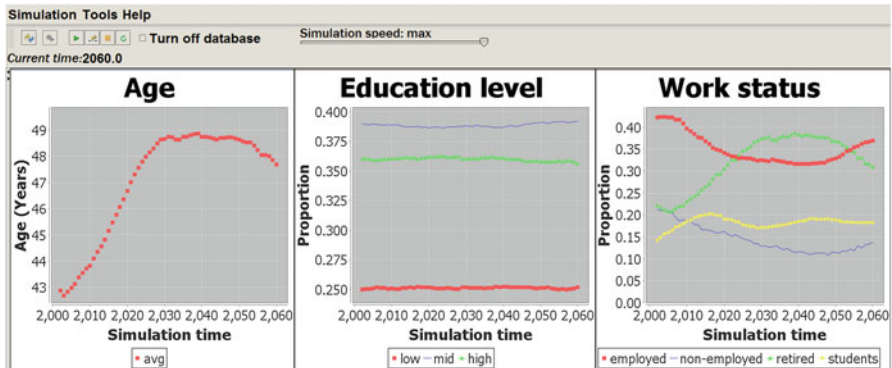


Fig. 4.1 The graphical output of one simulation run

commands. The output tables can also be exported in CSV format for subsequent analysis using specific statistical tools by typing:

```
CALL CSVWRITE('person.csv', 'SELECT * FROM PERSON');
```

4.7 Conclusions

The JAS-mine simulation platform achieves a convergence between agent-based and microsimulation tools. Its main goal is to speed up model development, facilitate model documentation, and foster model testing and sharing. The rationale behind this choice lies in the observation that computer power is always increasing, while researchers' time is not. Also, large-scale simulation projects are generally beyond the reach of a single scientist. Even when they remain under the control of a restricted group of people, they generally require a prolonged effort, often on a stop-and-go basis. The possibility of building on previous work done by the same authors or by other researchers is crucial. Simulation modelling needs cooperative development. The choice of an entirely open-source tool should be evaluated in this light. Moreover, JAS-mine does not force the user to adopt predefined solutions to the problems faced in simulation modelling. By offering a set of libraries that extend the capability of the standard Java classes, JAS-mine leaves entirely open the possibility of using other libraries and tools, either as an alternative or on top of the JAS-mine toolkit.

Acknowledgments Matteo Richiardi has benefited from support by Piedmont Region (research project: "From work to health and back: The right to a healthy working life in a changing society") and Collegio Carlo Alberto (research project: "Causes, Processes and Consequences of Flexsecurity Reform in the EU: Lessons from Bismarkian Countries"), as well as from support by a Marie Curie Intra European Fellowship within the 7th European Community Framework Programme.

We are grateful to Michele Sonnessa, who developed an earlier version of the platform and has contributed to the coding of the microsimulation model.

References

- Baekgaard, H. (2002). *Micro-macro linkage and the alignment of transition processes: Some issues, techniques and examples*. National Centre for Social and Economic Modelling (NAT-SEM) Technical paper No. 25.
- Billari, F., & Prskawetz, A. (2003). *Agent-based computational demography*. Berlin: Springer.
- De Menten, G., Dekkers, G., Bryon, G., Liègeois, P., & O'Donoghue, C. (2014). LIAM2: A new open source development tool for discrete-time dynamic microsimulation models. *Journal of Artificial Societies and Social Simulation*, 17(3), art. 9.
- Gilbert, N., & Troitzsch, K. (2005). *Simulation for the social scientist* (2nd ed.). Maidenhead: Open University Press.
- Keller, A.M., Agarwal, S., & Jensen, R. (1993). Enabling the integration of object applications with relational databases. *ACM SIGMOD Proceedings*.
- Klevmarcken, A. (2002). Statistical inference in micro-simulation models: Incorporating external information. *Mathematics and Computers in Simulation*, 59(1–3), 255–265.
- Li, J., & O'Donoghue, C. (2012). Simulating histories within dynamic microsimulation model. *International Journal of Microsimulation*, 5(1), 52–76.
- Li, J., & O'Donoghue, C. (2014). Evaluating binary alignment methods in microsimulation models. *Journal of Artificial Societies and Social Simulation*, 17(1), art. 15.
- Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., & Balan, G. (2005). MASON: A multiagent simulation environment. *Simulation*, 81(7), 517–527.
- Luna, F., & Stefansson, B. (2000). *Economic simulations in Swarm: Agent-based modelling and object oriented programming*. Dordrecht: Kluwer.
- Mannion, O., Lay-Yee, R., Wrapson, W., Davis, P., & Pearson, J. (2012). JAMSIM: A microsimulation modelling policy tool. *Journal of Artificial Societies and Social Simulation*, 15(1), art. 8.
- Minar, N., Burkhart, R., Langton, C., & Askenazi, M. (1996). The Swarm simulation system: A toolkit for building multi-agent simulations. *Santa Fe Institute Working Paper*, 96-06-042, Santa Fe.
- North, M. J., Collier, N. T., Ozik, J., Tataru, E., Altaweel, M., Macal, C. M., Bragen, M., & Sydelko, P. (2013). Complex adaptive systems modeling with repast simphony. *Complex Adaptive Systems Modeling*, 1, 1–26.
- Richiardi, M. (2012). Agent-based computational economics. A short introduction. *The Knowledge Engineering Review*, 27(2), 137–149.
- Richiardi, M. (2013). The missing link: AB models and dynamic microsimulation. In S. Leitner & F. Wall (Eds.), *Artificial economics and self organization* (Lecture notes in economics and mathematical systems, Vol. 669). Berlin: Springer.
- Statistics Canada (2009). *Modgen version 10.1.0 developer's guide*. <http://www.statcan.gc.ca/eng/microsimulation/modgen/doc/devguide/dev>. Accessed 17 Dec 2014.
- Stephensen, P. (2014). An information-loss-minimizing approach to multinomial alignment in microsimulation models. *DREAM Working Paper*, 2014:4.
- Wilensky, U. (1999). *NetLogo user manual*. <https://ccl.northwestern.edu/netlogo/docs>. Accessed 10 Apr 2015.

Chapter 5

Simulating Synthetic Life Courses of Individuals and Couples, and Mate Matching

Sabine Zinn

5.1 Introduction

This paper presents a novel modeling and simulation approach for fine-grained population projection, which merges demographic microsimulation and agent-based modeling. The main idea behind this approach is to model and simulate life course dynamics of individuals and couples by means of traditional demographic microsimulation and to use agent-based modeling for mate matching.


Demographic microsimulation is well suited to conducting fine-grained population projection if only independent entities such as individuals or couples are concerned (van Imhoff and Post 1998). However, as soon as kinship and/or inter-individual interactions are to be considered as well, microsimulation is likely to encounter specification problems (see, e.g., Ruggles 1993; van Imhoff and Post 1998; Murphy 2006). The approach presented here offers an opportunity to tackle this problem: it facilitates the specification of life courses of individuals and couples, and also the establishment of a partnership market. In this way, it allows us to test, for example, how different policies might affect demographic events which depend on the mutual decision of both spouses—for instance, whether implementing parental leave benefit for both men and women can significantly affect childbearing decisions. In detail, population dynamics in our model are driven by synthetic life courses of individuals and couples, which are defined by the sequences of states that individuals and couples enter over time, and by the waiting times between these state transitions, see Fig. 5.1. These states, which individuals and couples enter

S. Zinn (✉)

National Educational Panel Study, Leibniz Institute for Educational Trajectories,
Bamberg, Germany

Max Planck Institute for Demographic Research, Rostock, Germany
e-mail: sabine.zinn@lifbi.de

TWO INDIVIDUAL LIFE COURSES



LIFE COURSE OF COUPLE



Fig. 5.1 The upper part of the figure shows the life courses of two opposite-sex individuals. Once they get married, a couple life course is created. Through divorce the couple life course dissolves and the individual life courses of both (divorced) spouses are renewed

over time, summarize the demographically relevant categories that an individual or a couple can belong to. Generally, the state space is determined by the problem to be studied, but commonly it will at least comprise the elementary demographic characteristics of sex, fertility, and marital status. In demographic microsimulation, life courses usually evolve along two time scales: individual age and calendar time. A possible third time scale is the time that an individual or a couple has already spent in the current demographic state, for example, the time that has elapsed since getting married. In our setting, life courses are specified in continuous time, that is, transitions between states (i.e., events) can occur at any instant of time. No changes take place in between subsequent events and only a finite number (i.e., a discrete sequence) of events can occur in a finite time span.¹ In other words, our model is a discrete event model (Cellier 1991, Section 1.9). In a microsimulation model like ours, the underlying stochastic model is parameterized with transition rates that are assumed to vary with age and also with calendar time. Commonly, for their estimation statistical methods of event history analysis are applied to retrospective or prospective life histories that are reconstructed from longitudinal data and/or vital statistics. Assumptions about future rates then define the projection scenarios.

Unfortunately, as soon as a conventional demographic microsimulation is confronted with inter-individual interaction and couple behavior, problems of parametrization arise: on the one hand, inter-individual interaction processes such as mate choice mechanisms are largely unobservable, and only the outcome of these processes can be seen (de Vos and Palloni 1989). That is, it is hard—if not impossible—to estimate accordant transition rates. On the other hand, surveys

¹This is opposed to continuous simulation (not to be confused with continuous time microsimulation) where system dynamics are continuously tracked.

mostly gather only very limited information on partnership relations (Huinink and Feldhaus 2009). Thus, a realistic microsimulation model of couple behavior is usually hindered by data limitations. However, researchers are currently noticing considerable improvements regarding data availability (Mayer 2009).

Concerning the first problem described, agent-based modeling poses a solution: as soon as we face hidden processes driving certain life course decisions, we use components of agent-based modeling to generate life-course events. The basic idea is simple: within the scopes of microsimulation, let individuals interact such that in sum the microsimulation output resembles observed aggregates. To make it more concrete: for a certain year and population, we might know in advance how many couples will have formed. We might also know the within-couple age distribution and the distribution of the spouse's educational attainment. This information will suffice to set up, for example, a partnership market that produces couples whose attributes closely resemble the observed distributions (Zinn 2012).

In order to accurately specify this novel modeling approach, we suggest using the ml-DEVS formalism as described by Uhrmacher et al. (2007). This formalism is a variant of the discrete event specification (DEVS) model language developed by Zeigler et al. (2000), extended by explicitly supporting multi-level modeling. We deem this formalism adequate for our purposes as it allows us, on the one hand, to specify population dynamics in the requested way and, on the other hand, to benefit from well-proven and efficient up-to-date simulation methodology.

The remainder of this paper is structured as follows: in Sect. 5.2 we detail the stochastic model applied to describe life course dynamics of individuals and couples, and we describe the simulation processing used to construct synthetic life courses. Then, in Sect. 5.3, we present the mate matching procedure applied for building couples over simulation time. Section 5.4 outlines the implementation of the simulation model: we describe the ml-DEVS model which we have designed to specify the population model at hand as well as its execution semantics. In Sect. 5.5, we illustrate the capability of the novel simulation approach using a simplified example that aims to forecast the contemporary Dutch population. We analyze partnership and smoking behavior. We conclude our work in Sect. 5.6 by summarizing its key lines, and by presenting problems that remain for future work.

5.2 Simulating Individual and Couple Behavior in Continuous Time

The model of a demographic microsimulation comprises a virtual population and a stochastic model of individual and couple behavior. The virtual population includes all individuals and couples for whom life courses are simulated over simulation time, that is, individuals and couples who are part of the base population (i.e., the set of individuals and couples with whom we start our simulation), individuals who enter the population by birth, couples that are built during simulation, and individuals and couples who immigrate into the population. To describe individual and couple behavior we use continuous-time multi-state models. Continuous-time

multi-state models are stochastic processes—commonly Markovian processes—that, at any point in time, occupy one out of a set of discrete states (Hougaard 1999). The state space summarizes all discrete states considered. Generally, the state space is determined by the problem under study but, commonly, it will at least comprise the elementary demographic characteristics of sex, fertility status, and marital status. In our terminology used here, an individual’s and a couple’s state is a combined characteristic given by the combination of several attributes. To simplify the notation, we define state variables. These are the demographic categories considered, such as sex, fertility status, and marital status. All unique combinations of these state variables’ values thus form the state space.

We define individuals and couples to be the atomic components of our model. Hence, the state space of our simulation model can be decomposed into two sub-state spaces: one for individuals and one for couples. Table 5.1 gives an example of four state variables: sex, fertility status, marital status, and mortality. In the example, a potential state of an individual is (*female, childless, never married, alive*), and a potential state of a couple is (*one child, married, both alive*).

Over simulation time, individuals might experience events, that is, transitions between states. In principle, transitions may occur between all states of the state space. However, the problem under investigation restrains the set of possible events. Table 5.2 shows a list of feasible events for the example state space given in Table 5.1. Two individuals form a couple after a successful mating period. The corresponding process is described in Sect. 5.3. Once a couple is formed the individual life courses of both spouses are combined into one couple life course. If a couple experiences a dissolution event, such as divorce or widowhood, the

Table 5.1 Example of state variables that individuals and couples might occupy, inclusively possible values, separated by commas

State variable	Individual values	Couple values
<i>Sex</i>	Female, male	Opposite-sex couple
<i>Fertility status</i>	Childless, child(ren)	Childless, child(ren)
<i>Marital status</i>	Never married, divorced, widowed	Married, divorced, widowed (she is dead), widowed (he is dead)
<i>Mortality</i>	Alive, dead	Both dead, she is dead (he is alive), he is dead (she is alive), both alive

Table 5.2 Example of possible events that individuals and couples might experience

Event type	Individual events	Couple events
<i>Fertility event</i>	Giving birth (for females), becoming father (for males)	Becoming parents
<i>Partnership events</i>	Marrying	Getting divorced
<i>Mortality events</i>	Dying	Female dies and male is widowed, male dies and female is widowed, both die

couple life course decomposes again into two individual life courses (in case of a dissolution event) or into one individual life course (in case of a widowhood event), cf. Fig. 5.1.

In a continuous-time simulation model—such as the one considered here—the propensity of an individual or couple to experience a transition is determined by transition rates. Depending on the process class used to describe individual or couple behavior these rates might depend on different time scales (e.g., age, calendar time, time elapsed since a specific event) and sequences of states already passed. For reasons of convenience, we use transition rates that depend on age, calendar time, the state of occurrence, and the state of destination. In other words, we use nonhomogeneous continuous-time Markov chains $Z(t)$, $t \in \mathbb{R}_0^+$, to describe individual and couple behavior. The process time t maps the time span over which we observe the life course of a unit. The time t is set to zero when an individual or couple is created and evolves throughout their life time.

The process $Z(t)$ is fully defined by the two-dimensional process $(J_n, T_n)_{n \in \mathbb{N}_0}$ where $(J_n)_{n \in \mathbb{N}_0}$ is a Markov Chain mapping all states that an individual or couple occupies, and $(T_n)_{n \in \mathbb{N}_0}$ is the sequence of the corresponding transition times along process time t . Thus, the transition rate of an individual or couple to undergo a transition from a state s_j to a state s_k is

$$\lambda_{s_j, s_k}(t) = \lim_{h \downarrow 0} \frac{1}{h} P[J_{n+1} = s_k, T_{n+1} \in (t, t+h] \mid J_n = s_j, T_{n+1} \geq t].$$

By way of translation, the process time t can easily be mapped onto age and calendar time.² Thus, modeling age and calendar time dependence is straightforward. Once the transition rates of a Markovian process are known, the distribution functions of the waiting times in the distinct states of the state space can be derived and synthetic life courses can be constructed. More precisely, the distribution function of leaving state s_j at time t after waiting time w for moving on to state s_k is

$$F(w, t) = 1 - \exp \left\{ - \Lambda_{s_j, s_k}(w, t) \right\}, \quad (5.1)$$

where

$$\Lambda_{s_j, s_k}(w, t) = \int_t^{t+w} \lambda_{s_j, s_k}(u) du$$

is the corresponding integrated hazard rate. By means of the distribution function (Eq. 5.1), we can derive a random waiting time in state s_j until moving

²Specifically, the function $C(T_n)$ maps the calendar time at T_n and the function $A(T_n)$ maps the age of an individual at T_n . Similarly, $A^f(T_n)$ maps the age of a female spouse at T_n and $A^m(T_n)$ maps the age of a male spouse at T_n . At individual birth time the functions $A(T_n)$, $A^f(T_n)$, and $A^m(T_n)$ take the value zero, and $C(\cdot)$ is zero at, for example, 01-01-1970 00:00:00.

on to state s_k . For this purpose, we use the inversion theorem (Rubinstein and Kroese 2008, p. 51f.). Hence, we yield a random waiting time w from the correct distribution by

$$w = \Lambda_{s_j, s_k}^{-1}(w, t)[- \ln(1 - u)],$$

with u being a standard uniformly distributed random number. By generating sequences of random waiting times until the next event, synthetic life courses can be constructed. For example, first a waiting time until school enrolment is derived, then a waiting time until school graduation, followed by a waiting time until first birth, and so on. Here, we have to consider that an individual might not only be exposed to one single event but to a set of possible events (i.e., to competing risks). For example, a 16-year old female might experience the event school graduation or teenage pregnancy as a next event. We account for competing risks such as these by computing for each possible destination state a random waiting time and by selecting the shortest one as the one to be realized. Over simulation time, this computation of the shortest waiting times is repeated for each individual of the virtual population. Once simulation ending time has been reached, the event times of all surviving individuals are censored.

The computation of random waiting times requires the inversion of the integrated hazard rate Λ_{s_j, s_k} . In demography, assuming constant transition rates over yearly intervals is a suitable and widespread approximation to Λ_{s_j, s_k}^{-1} , which clearly eases its derivation (Gill and Keilman 1990). Integrated hazards become piecewise linear, and waiting time distributions piecewise exponential.

Standard approaches to estimate transition rates are occurrence-exposure rates or the Poisson generalized additive model. Both approaches are discussed in Zinn (2011). Commonly, longitudinal survey data, vital statistics, or census data are used to estimate transition rates. For individual-based questions such data are available. However, the current data situation could hamper the estimation of transition rates for all types of events to which couples might be exposed—though, momentarily, we observe the buildup of a number of more complex surveys dealing (among many other things) with partnership issues.³ An idea to describe couple behavior, even if couple data are not available is to combine models of individual behavior. For this purpose we have to make assumptions of how individual behavior must be interlinked to yield couple behavior. This means, we have to model how (and for which transitions) the transition rates of the couple need to be modified relative to the rates of the individual spouses. An example: we may have age-specific rates of quitting smoking for men and women. If, within one couple, one partner quits smoking, this rate for the other partner suddenly will be much higher than before. In principle, interlinking individual behavior this way is a task in statistical

³Examples are the German National Educational Panel Study NEPS (<https://www.neps-data.de/en-us/home.aspx>) and the German Family Panel pairfam (<http://www.pairfam.de/en/study.html>).

modeling, and many general approaches are available. Mostly we will have to base the analysis on external knowledge about the studied phenomenon and the data available to model it.

5.3 Matching of Individuals

So far, we have detailed how, in a continuous-time microsimulation, synthetic life courses of individuals and couples can be constructed. However, we have not made clear how couples are formed, that is, how mate matching is performed. As mate choice mechanisms are largely unobservable and the starting time of mate search activities is usually not known exactly, transition rates for matching individuals are hard (if not impossible) to estimate. Without such rates, pure microsimulation is not capable of performing mate matching. This is in contrast to agent-based modeling, where individuals interact according to rules being usually based on societal or behavioral theories. Such theories might be substantiated by empirical and hypothesized data. In other words, agent-based models allow us to replicate hidden processes by combining behavioral theories and observed data. Thus, even though some decision processes (such as mate choice decisions) are to a large extent unobservable, agent-based models facilitate the replication of observed aggregates (for example, the number of homosexual couples within one year). In a continuous-time microsimulation model, we compute waiting times until the next event. As a consequence, we know in advance when couples have to be formed. Subsequently, we interpret the occurrence of partnership formation as the outcome of a hidden mating process which we describe by an agent-based model. In other words we simulate partnership formation events via microsimulation and we decide who mates whom via agent-based modeling. This way, we are embedding an agent-based model of mate choice into the framework of a continuous-time microsimulation model. The subsequent description of this agent-based model relies, to large a extent, on the work presented in Zinn (2012).

Concerning the timing of partnership onsets, continuous-time microsimulations pose some problems that discrete-time models can avoid. In discrete-time models, which update information on demographic events at discrete points in time (usually each year or each month), it is convenient to construct mating pools at equidistant time points, for example, for every year. During simulation, individuals enter these mating pools and undergo mate matching.

In continuous-time models, events occur at exact time points and individuals will never experience partnership events at the same time. Therefore, a pool of potential partners cannot as easily be constructed as in discrete-time models. A simple way to avoid this problem would be to use a so-called open model. In this model class, spouses are created as new individuals when needed, rather than being selected from already existing members of the population. Although such “external” partnership formations do happen in real populations, they constitute the minority of cases.

Open mating models therefore would artificially increase the number of individuals, and new individuals would still have to be supplied with individual attributes to allow for a realistic simulation of their remaining life course. Therefore, we focus on so-called closed models, where appropriate spouses must be identified from the current members of a population. In this context, we have to solve three problems: how to construct a feasible partnership market? How to identify a proper spouse? What to do if no proper spouse can be found? The first problem can be tackled by constructing a partnership market that individuals can enter and leave over the complete simulation time range. That is, once a partnership onset event has been simulated for an individual, he/she enters the partnership market. As soon as the individual finds a proper spouse, both spouses leave the market. The remaining two problems (identifying matching spouses and specifying feasible options in case of unsuccessful partner search) will be the topic of the subsequent sections.

We expressly point out that in the field of microsimulation and agent-based modeling already several useful and convincing mate-matching algorithms exist. Comprehensive summaries are given in, for example, Walker (2010, Chapters 2 & 3) and Zinn (2011, Chapters 9 & 10).

5.3.1 Identifying Matching Spouses

5.3.1.1 Overlapping Mating Periods

In a continuous-time microsimulation, it is impossible that two individuals feature identical mating times. However, a partnership has to have a clearly defined unique formation time. Consequently, to determine such a time already computed event times have to be shifted. In doing so, we have to be very careful to keep the distortion of the microsimulation output at a minimum. That is, we have to ensure that the differences between the already scheduled event times and the shifted ones are small. To conform to this requirement, we introduce individual mating periods which are defined to start at maximum one year before the scheduled event time of partnership onset and to stop at maximum one year later. Concretely, let a woman I_1 experience the onset of a partnership at time t_1 and a man I_2 experience the onset of a partnership at time t_2 . Without loss of generality, we assume $t_2 < t_1$. Then the mating period of I_1 and I_2 is

$$\Gamma_i = [\min(t_e, t_i - B), \max(t_i + B, t_E)], \quad i = 1, 2,$$

where t_e being the time of the last event of I_i , $2B$ is the maximum length of the mating period (with B being smaller than one year), and t_E is the simulation stopping time. We determine that the two individuals I_1 and I_2 can only mate if

$$\Gamma_1 \cap \Gamma_2 \neq \emptyset,$$

that is, if they have overlapping mating periods.⁴ Then, the partnership formation time \tilde{t} of I_1 and I_2 is defined to be the mean of the event times t_1 and t_2 :

$$\tilde{t} = t_2 + 0.5(t_1 - t_2).$$

This way of determining partnership formation times is very different to the way of defining them in agent-based models for partner search. There mating times are the outcomes of the matching algorithm itself; see for example Billari et al. (2008), Todd et al. (2005), and Todd and Billari (2003). In our simulation model, mating times are the outcomes of stochastic processes which are parameterized with empirical or hypothesized rates. Thus, the purpose of our mating model is not to reproduce observed or assumed mating times, but to create couples with feasible characteristics while minimizing the distortion of the microsimulation output. This objective is in clear contrast to agent-based models that describe behavior while aiming at resembling certain stylized facts, for example, the age distribution at first marriage.

5.3.1.2 Compatibility of Individual Characteristics

Even if the mating periods of individuals overlap, their characteristics might not match. Therefore, besides event times, also individual characteristics have to be checked for conformance. For this purpose, we use a compatibility measure. The measure transforms female and male attributes into a numeric index that quantifies how well two potential spouses fit together. Values between zero and one are used to express the degree of matching, with a large value indicating high compatibility. Likewise, a small value points to incompatibility. Specifically, we use logit models to predict how well the characteristics of potential spouses agree with one another. In order to account for different types of partnerships (cohabitations and marriages), we apply a separate logit model for each partnership type. The models predict the probability that two individuals, each with given attributes, form a particular kind of partnership. Which covariates will enter the logit models depends on the state space of the actual application. As our microsimulation model is a generic model, the state space is not fixed. However, naturally, individual age and sex should be included. Obviously, we can only include covariates that are mapped by the state space. If, for example, educational attainment, children ever born, or ethnicity are included in the state space, these attributes are natural candidates for covariates in the logit models. Data on observed couples are used to estimate the coefficients of these models. According to the theory of assortative mating, partners tend to have similar attributes such as similar ages and levels of education (Zietsch et al. 2011;

⁴Setting the length of the individual mating periods to two years at maximum is a compromise between minimizing the potential distortion of the microsimulation output and ensuring a feasible number of potential partners available within a certain time period.

Blossfeld 2009; Kalmijn 1998). Therefore, the estimated coefficients are likely to be in accordance with the theory of assortative mating. Nevertheless, depending on the data used the estimated coefficients might also point to deviations from that theory. This makes logit models a very flexible tool to measure compatibility because they are not tailored to one particular theory of mating.

In order to estimate logit models of spousal compatibility, we randomly assign to each male spouse a female who is not his observed partner.⁵ We set the response variable to one if the couple has been observed, and to zero otherwise. This way, we construct a data set with an identical number of observed synthetic couples, resembling the retrospective data design of a case-control study. Unfortunately, when conducting mate matching we are confronted with a prospective problem, that is, we need the likelihood that two individuals with certain attributes are going to mate. Hence, for a prospective problem we have constructed a case-control data set. Prentice and Pyke (1979) show that all nonintercept parameters of a prospective logit regression model are asymptotically correct when using a case-control data set; only the estimator of the intercept is biased. In our mate matching procedure (see subsequent section), we decide on a match between two potential spouses depending on their attributes, and not on the composition of the pool of available candidates. That is, we measure compatibility on a relative scale. Therefore, for our purposes, the estimation of a prospective logit model is suitable.

5.3.1.3 Mate Matching

Once an individual has entered the partnership market, he/she starts to look for an appropriate partner. For this purpose, the seeking individual inspects other individuals in the market. As the human network size is limited to approximately 150 members (Hill and Dunbar 2003), the number of potential spouses is restricted from the outset. We set an upper bound that follows a normal distribution with expectation $\mu = 120$ and standard deviation $\sigma = 30$. Furthermore, we assign to each individual a random value that captures his/her aspiration level regarding a partner. This aspiration level takes values between zero and one. If the compatibility measure between an individual and a potential spouse exceeds this aspiration level, he/she will accept an offer. Every time an individual has been inspected and fails to be chosen, he/she reduces his/her aspiration level by $\delta_A = 0.1$. We use the beta distribution to describe the individual aspiration levels. As is generally known, the degree of choosiness of females and males varies with age (Trivers 1972; Buss 2006). We assume that women's requirements decrease with declining fecundity. This is in accordance with, for example, the findings of Waynforth and Dunbar (1995) and de Sousa Campos et al. (2002) who figure out that with fertility decline women tend to become less demanding. In contrast to men: when they are young,

⁵In order to avoid biased results, we do not assign to a male spouse a female spouse who has attributes identical to those of his observed spouse.

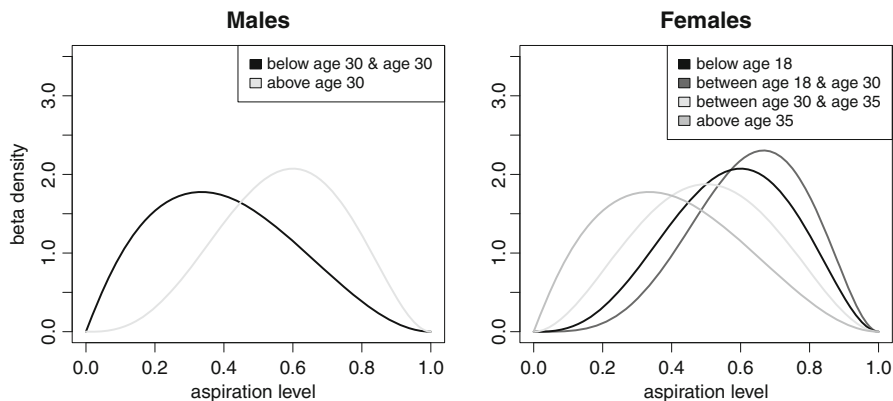


Fig. 5.2 Age and gender-specific densities of the beta distributions that are used to determine aspiration levels regarding partners

they are more involved in short-term relationships and thus less selective concerning their partner's traits. When they are older, men start to look for a long-term relationship and to invest in offspring. Hence, their level of choosiness increases. To account for the variability in the degree of aspiration, we parameterize the beta distribution accordingly, see Fig. 5.2. An important aspect when searching for a mate is the size of the pool of potential spouses. If it is small, it is not reasonable to assume a very selective seeker. Therefore, we assume that, if a seeker faces less than ten potential partners, he/she reduces the aspiration level by $\delta_B = 0.3$.

To implement our agent-based model of mate choice we use a *marriage queue* \mathcal{M} . The marriage queue comprises all unpaired individuals looking for a partner. Moreover, each individual in the queue is equipped with a stamp indicating the time scheduled for the upcoming partnership event. To create couples, we use a sequential mate matching algorithm. Concretely, when an individual I_i enters the partnership market the algorithm performs the following steps.

- The mating period Γ_i of I_i is determined and the aspiration level a_i of I_i is generated.
- If the marriage queue \mathcal{M} is empty (i.e., the partnership market is empty), I_i is inserted into \mathcal{M} .
Otherwise:
 - A. The social network size of I_i is determined by drawing a normally distributed random number N , with mean μ and variance σ^2 . If N is greater than the current number $N_{\mathcal{M}}$ of individuals in the marriage queue, N is set to $N_{\mathcal{M}}$.
 - B. N individuals, whose mating periods overlap with Γ_i , are randomly taken from \mathcal{M} . They are inserted into the working marriage queue \mathcal{W} .
 - C. Individuals of the same sex such as I_i are removed from \mathcal{W} .
 - D. If \mathcal{W} is empty, I_i is directly inserted into \mathcal{M} .

Otherwise:

- (i) If \mathcal{W} contains less than ten individuals, the aspiration level of I_i is reduced to $a_i = \max(0, a_i - 0.3)$.
- (ii) For the potential spouses to inspect, a counter $j = 1$ is initialized.
- (iii) From \mathcal{W} , the j th individual I_j (with aspiration level a_j) is selected and the compatibility measure c_{ij} between I_i and I_j is computed. If $a_i < c_{ij}$ and $a_j < c_{ij}$, the individuals I_i and I_j are paired, and I_j is removed from \mathcal{M} .
- (iv) Otherwise, the aspiration level of I_i is reduced to $a_i = \max(0, a_i - 0.1)$ and the aspiration level of I_j is reduced to $a_j = \max(0, a_j - 0.1)$.
- (v) The value of the counter variable j is incremented by one.

The steps (iii) and (v) are repeated until either I_i is paired or all individuals of \mathcal{W} have been inspected.

If for I_i no appropriate spouse can be found, I_i is enqueued into \mathcal{M} .

5.3.2 Options in Case of an Unsuccessful Search

The presented mate matching algorithm does not guarantee that each searching individual will be paired. Mate matching fails, if an individual is unable to find within his/her mating period a spouse with compatible characteristics. Five options exist to cope with the problem of an unsuccessful search.

1. Form a couple with the most compatible opposite-sex seeker who is searching for the same type of partnership within the same searching period.
2. Extend the mating period.
3. Return the individual to the model population unpaired. That is, the individual is again at risk of experiencing a partnership event (or, alternatively, any other kind of event).
4. Let the individual emigrate.
5. Let a proper spouse immigrate.

The last option is inspired by the processing in open models because an appropriate spouse is taken from “outside”. Each of these options entails a major difficulty. Forming couples between unsuccessful seekers and their most compatible counterparts in the partnership market holds the danger that too many couples with little compatibility are created. Extending the mating period means shifting the time of the scheduled partnership event to a later time point and thus notably distorting the output of the underlying stochastic model of individual behavior. Rejecting a seeker and sending him back unpaired implies that an already scheduled event is completely ignored. Allowing (too many) immigrated spouses will eventually spoil the model population and hence the plausibility of the model outcome. Consequently, searching periods that expire without success should definitely be an

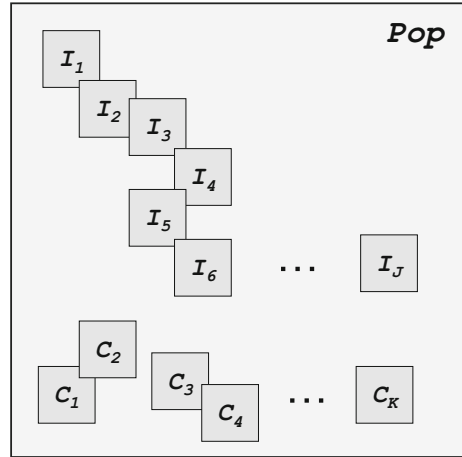
exception. This can only be assured if the model population maps a large proportion of an actual population, so that each seeker can meet at least one potential spouse.⁶

5.4 Implementation and Software

Throughout this paper, we consider a microsimulation model that describes life courses as sequences of discrete events that occur along a continuous time line. That is, our microsimulation model is a discrete event simulation model. For implementation purposes, we intend to deploy a reliable and efficient up-to-date simulation methodology. Moreover, developing new software from scratch should be avoided. Therefore, to ease implementation, we make use of an already established and well-proven discrete event simulation approach, the discrete event specification language DEVS (Zeigler et al. 2000; Uhrmacher et al. 2010; Wainer et al. 2002). It originates in systems theory, defines real systems as composites of submodels (which are either atomic or coupled), and is based on a state-based concept. DEVS does not assume a particular stochastic model to specify system behavior. Thus, it allows us to describe population dynamics by nonhomogeneous continuous-time Markov chains. In other words, DEVS offers all the functionalities required to describe the microsimulation at hand. Several DEVS-based tools exist (e.g., DEVSJava, JDEVS, CD++, DEVS variants in JAMES II) that support implementation. To facilitate the modeling of different kind of problems, various DEVS variants have been developed (Uhrmacher et al. 2010; Zeigler et al. 2000). We use ml-DEVS formalism (Uhrmacher et al. 2007). This formalism is a variant of the classical DEVS model language, extending it by explicitly supporting multilevel modeling. Commonly, a ml-DEVS model consists of micro models embedded in a macro model. The macro level model is described by a coupled DEVS model, equipped with a state and behavioral rules of their own, and the micro models are described by ordinary atomic DEVS models. Communication between micro and macro models is handled by exchanging messages. The propagation of information from the micro level to the macro level is facilitated by equipping micro models with the ability to change their ports. In this way, the macro model can access the information given in the exhibited ports of the micro models, and the micro models can influence macro-level dynamics. The macro model can concurrently activate (several) micro models by signaling messages via value coupling. Value coupling means that, at the macro level, information is mapped to specific port names, and every micro model can access this information by forming input ports with corresponding port names. The ml-DEVS formalism supports variable structure models, that is, models

⁶By nature, the run times of such simulation settings are considerably longer than the run times of simulation settings that only implement sparse marriage markets. This is due to the fact that in sparse marriage markets many seekers might fail to encounter a potential partner. Thus, there is nobody to be screened and nothing to compute.

Fig. 5.3 The ml-DEVS population network **Pop** comprises a finite number J of individual models I_j and a finite number K of couple models C_k



that change their own composition, interaction, and behavior pattern. Structural changes are operated top-down directed by the macro model and executed on the micro level.

We describe our microsimulation model by defining micro models of individuals and couples embedded in a population macro model. Transition rules determine the behavior of each model. The transition rules of the micro models control life course transitions, whereas the transition rules of the population macro model direct mate matching and structural changes of the population (i.e., adding and removing of micro models). To efficiently execute our ml-DEVS model of population dynamics we extend the so-called sequential hierarchical simulator for (parallel) DEVS. The ml-DEVS population model and its simulator are implemented by means of the modeling and simulation framework JAMES II.⁷ The simulation software and illustrative examples are available in the model library of openABM (Zinn 2015) and also from the author upon request.

Our ml-DEVS microsimulation model consists of a macro model **Pop** comprising two types of micro components: individual models **I** and couple models **C**, see Fig. 5.3. These micro components handle the life course dynamics of individuals and the dynamics of married or cohabiting couples. The macro-DEVS model guides the onset of partnerships (marriages or cohabitations), that is, it performs mate matching and instructs the creation of couple models. If a couple model signals an emigration, a dissolution, or a widowhood event, the macro-DEVS model instructs its extinction or dissolution. Likewise, the macro-DEVS model handles death and migration events of individual models as well as the creation of micro models in case of immigration and childbirth events. The full specification of the ml-DEVS population model is given in very detail in Appendix A.

⁷JAMES II can be downloaded at <http://www.jamesii.org> and is distributed under the JAMESLIC which allows free reuse for commercial and noncommercial projects.

Commonly, the execution of a discrete event simulation model demands an event queue, which is a list of events sorted by their scheduled event times. In our setting, the event queue holds the events that are scheduled for the whole population (e.g., successful mating events), as well as for the individuals and the couples of the virtual population. In each simulation step the event with the minimal time stamp is dequeued and executed. Then, for the related model a next event is computed and—if that event is neither death nor emigration—enqueued. In other words, an event queue organizes the scheduling of upcoming events. For the processing of events executable simulation code has to be derived, that is, simulation semantics have to be specified. Generally, the execution semantics of a DEVS model are described by the *abstract simulator*, which comprises *simulators* and *coordinators* (Zeigler et al. 2000). In a ml-DEVS model, simulators correspond to micro models and coordinators to macro models (Uhrmacher et al. 2007). Coordinators are responsible for the execution and the correct synchronization of the simulators of the micro models and for the handling of external events (in our case: the arrival of immigrants). In line with this, a coordinator implements an event queue algorithm managing upcoming events of micro models. Synchronization is guaranteed by communication protocols: if a model consists of only one macro model, as is the case for our population model, the coordinator waits for protocols sent by its subordinate simulators and transmits them to the root-coordinator. The root-coordinator guides the overall simulation processing. It initializes a new simulation and instructs the model execution until some termination criterion is met (e.g., the simulation stop time has been reached). To each ml-DEVS model a corresponding processor tree can be given, which directly maps the hierarchical structure of the model on the architecture of simulators and coordinators. Figure 5.4 displays the processor tree corresponding to the ml-DEVS population model designed. To ensure consistency within each simulation step, messages between the root-coordinator, the coordinator, and the simulators are processed in a well-defined order: if an

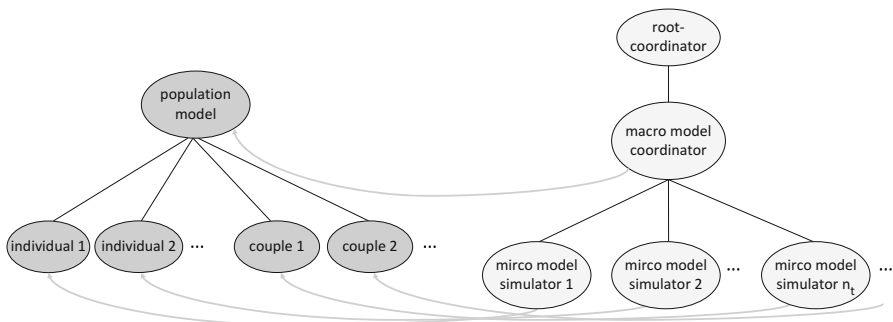


Fig. 5.4 Processor tree corresponding to our ml-DEVS population model

(internal) event is due,⁸ the coordinator of the macro model activates the micro model concerned by sending a \star -message. After having performed the event, the simulator of the micro model forwards to the coordinator a y -message with output data. In our case, such output data might comprise the number of newborns in case of a childbirth event or the states and birth dates of individuals in case of emigration. Afterwards, the simulator of the micro model waits for an x -message comprising input information. For example, an individual micro model is informed about the actual mating time in case of a successful partner search. Finally, a *done*-message signals the completion of a simulation step. The \star -, y -, and x - comprise, besides their regular information, also the actual simulation time t , and the *done* message comprises the time *tonie* of the next event and all approachable ports *ports*.

Uhrmacher et al. (2007), who have developed the ml-DEVS formalism, suggest an abstract simulator, which is essentially a direct implementation of the ml-DEVS processing scheme. That is, it executes a ml-DEVS model by passing messages successively through the model's processor tree. Such processing means that any time a superordinate model sends a message to a subordinate model, it has to wait for the subordinate model to react and to respond. Only then simulation processing can proceed. Such an approach poses problems (Himmelspace and Uhrmacher 2006). On the one hand, depending on the depth of the processor tree and the number of models, a remarkable overload might result. On the other hand, the ml-DEVS simulator of Uhrmacher et al. (ibid.) demands for each micro model one thread.⁹ Regular Java virtual machine implementations, however, do not support more than a few thousand parallel threads. This implies that meaningful population projections would not be possible. To overcome this limitation, we have designed a novel simulator for our ml-DEVS population model. We have extended the *sequential abstract simulator* for parallel DEVS developed by Himmelspace and Uhrmacher (2006). The novel abstract simulator executes parts of the processor tree en bloc which means to exploit computing power in an efficient way. The related communication protocol is displayed in Appendix B. Likewise, this appendix comprises the pseudocode of the respective coordinator and simulator functions as well as a comprehensive description of their functionality.¹⁰

⁸We call the state transition of a model an internal event if it has not been provoked by an input message. Otherwise, the event is denoted as external event.

⁹ Within a program, a thread is a sequentially executed stream of instructions.

¹⁰Note that the specification of the ml-DEVS population model and the newly designed sequential hierarchical simulator for ml-DEVS have already been described in a similar but more extended form in Zinn (2011).

5.5 A Hypothetical Application on Smoking Behavior of Couples

We conduct a small case study to illustrate the capabilities of the novel microsimulation model. For this purpose, we look at a synthetic population that (roughly) resembles the Dutch population. We consider smoking and partnership behavior. In particular, we study how partners influence each other's smoking behavior. Note that the presented application is mainly based on synthetic data, and should not be used to draw conclusions about actual behavior. It only serves as a means to demonstrate the potential of a microsimulation that accounts for interdependencies between life courses of married or cohabiting individuals.

5.5.1 *The Synthetic Population and Parametrization*

Starting on January 1, 2008, we generate life courses of a synthetic population that (roughly) resembles one percent of the Dutch population. The simulation horizon ranges from January 1, 2008 to December 31, 2020. During simulation, we consider individuals aged between zero and 63 years. The state space that we employ for individuals and couples is shown in Table 5.3. If the value 'being single & living alone' is assigned to an individual, he/she lives either alone and never lived in a union before, or he/she lives alone but was cohabiting before, or he/she lives alone and was married before. The value 'dissolved' indicates the separation of a married or cohabiting couple. Spouses that are dissolved or widowed enter the 'being single & living alone' state of the individual model. The events that individuals and couples can experience during simulation are listed in Table 5.4. To describe individual and couple behavior, for simplicity we assume that

Table 5.3 State variables that individuals and couples might occupy, inclusively possible values, separated by commas

State variable	Individual values	Couple values
<i>Sex</i>	Female, male	Opposite-sex couple
<i>Fertility status</i>	Childless, at least one child	Childless, parents
<i>Marital status & living arrangement</i>	Living at parental home & never married/cohabiting before, being single & living alone	Married, cohabiting, dissolved, widowed (she is dead), widowed (he is dead)
<i>Smoking status</i>	Non-smoker, smoker	Non-smoker couple, dual smoker couple, female smoker & male non-smoker, male smoker & female non-smoker
<i>Mortality</i>	Alive, dead	Both dead, she is dead (he is alive), he is dead (she is alive), both alive

Table 5.4 Possible events that individuals and couples can experience

Possible events of individuals	Possible event of couples
Leaving parental home, dying, launching a cohabitation, marrying, quitting to smoke, starting to smoke, giving birth (for females)	Getting divorced or separated, childbirth, female starts smoking, male starts smoking, female quits smoking, male quits smoking, both quit smoking, both start smoking, female dies and male gets widowed, both die, male dies and female gets widowed

- the same fertility rates apply to paired and unpaired women,
- the risk to quit and to start smoking depends on the presence of children and the smoking traits of the partner, but does not depend on the marital status or the living arrangement,
- the extent of the impact that the smoking behavior of the spouse has on the own smoking behavior does not vary with age,
- the divorce risk and the risk to break up depend on the presence of children, but not on the smoking behavior of the spouses,
- to single, married, and cohabiting individuals the same mortality risk applies, and
- the risk that both spouses die at the same time is very low (we set it to 10^{-5}).

The propensity of individuals and couples to experience events is quantified by either empirical or synthetic transition rates. We have estimated (non-parity specific) fertility rates of females and transition rates to change the marital status and/or the living arrangement for single individuals using the Family and Fertility Survey¹¹ for the Netherlands (FFS_NL). For this purpose, we have applied a slightly modified version of the MAPLES estimation procedure¹² (Impicciatore and Billari 2011). The transition rates of couples to change their marital status or the smoking behavior are mainly hypothesized, constructed such that they resemble observed transition patterns. We assume that the mortality rates vary with age, sex and calendar year, and the other rates are age- and sex-specific, but are held constant over calendar time. We use hypothetical death rates and transition rates of changing the smoking behavior. It is well known that smokers have a higher mortality risk than non-smokers (Doll et al. 2004). We account for this fact by adapting the mortality rates of the EuroStat2008 projections for the Netherlands (baseline scenario)¹³ accordingly: mortality rates for smokers are obtained by increasing these rates by 10 % and mortality rates for non-smokers are obtained by reducing them by 10 %. We assume that the mortality rates vary with age, gender and calendar year, and the other rates are age- and sex-specific, but are held constant over calendar time. All transition rates used are given in Appendix C.

¹¹<http://www.unece.org/pau/ffs/ffs.html>

¹²MAPLES estimates age profiles from longitudinal survey data using a generalized additive model and piecewise cubic splines.

¹³Detailed data on EUROPOP 2008 mortality were kindly provided by Eurostat.

To run a microsimulation model we need a base population to start with. For our illustration, we determine the base population to consist of individuals aged between zero and 63, differentiated according to sex, smoking status, fertility status, and marital status/living arrangement. We assume them to resemble one percent of the respective age groups of the Dutch population at January 1, 2008. Hence, the base population comprises 70,295 males and 68,264 females. Marginal distributions concerning the considered age classes and state variables—required to estimate the base population—have been taken from the EuroStat web portal,¹⁴ the Health and Retirement study¹⁵ and the FFS_NL survey. To estimate the base population we use the method of iterative proportional fitting (Deming and Stephan 1940). Figure 5.5 shows the resulting population numbers. Generally, our simulation model allows the consideration of migration; however, for reasons of simplicity we neglect migration for this example.

During simulation, we conduct mate matching as describe in Sect. 5.3. We use two logit models to quantify the compatibility between potential spouses: the first model describes the probability to enter a cohabitation and the second model describes the probability to enter a marriage without cohabiting before. Cohabiting spouses who marry are not considered here because they are already partnered. For estimating the models, we employ the first wave of the Netherlands Kinship Panel Study (NKPS)¹⁶ (conducted in the period from 2002 to 2004). We only consider partnerships that started in the years from 1990 to 2002. Our data set contains a record for each observed couple, which consists of the age of the male spouse, the age difference between the female and the male spouse (in integer years), a variable indicating whether the female or the male spouse were married before, and a variable showing whether the female spouse has children. The NKPS data do not contain any information about smoking behavior. Thus, on the basis of these data we cannot study accordant effects on matching probabilities. We come back to this issue at a later time and suggest a way to account for matching over smoking traits nevertheless. Following the procedure described in Sect. 5.3.1.2, we construct a data set of observed and nonobserved potential couples. By means of these data we estimate the two logit models. The estimated coefficients are given in the Tables 5.5 and 5.6. In both models the direct effect of the age of the male is very small. It is only slightly significant in the case of cohabitation and even insignificant in the case of marriage. For cohabitation and marriage we find—as expected—that individuals with small age differences are more prone to mate. We find no significant effect of the presence of children on a man’s propensity to marry or cohabit a woman. For marriage there is a slightly significant negative effect of whether one or both of the spouses experienced marriage before. The accordant effect is insignificant for cohabitation. We control for possible effects between marriage history (i.e., first and

¹⁴<http://epp.eurostat.ec.europa.eu/portal/page/portal/statistics/themes>

¹⁵<http://hrsonline.isr.umich.edu/>

¹⁶<http://www.nkps.nl/>




Fig. 5.5 The base population of males and females, cross-classified according to age, fertility status, smoking status, and marital status/living arrangement; *PH*: living at parental home, *MA*: married, *CO*: cohabiting, *SI*: being single (never married cohabiting before, dissolved, widowed), *S*: smoker, *nS*: non-smoker, *noChild*: childless, *child*: at least one child

Table 5.5 Results of logit model determining compatibility between potential partners entering cohabitation

Variable	Coefficient	p-value
Age of male	0.025	0.012
Age difference (age of male – age of female)		
<i>Greater than 7</i>	-2.525	< 0.001
<i>From 4 to 7</i>	-0.867	< 0.001
<i>From 2 to 3</i>	0.155	0.505
<i>From -1 to 1</i>	0	
<i>From -3 to -2</i>	-1.138	< 0.001
<i>From -7 to -4</i>	-1.975	< 0.001
<i>Smaller than -7</i>	-3.798	< 0.001
Woman has child(ren)	-1.355	0.179
Woman or man or both were married before	-0.125	0.423
Indirect effect		
<i>Woman has child(ren) & married before</i>	1.905	0.066
Number of pairs: 1472		

Table 5.6 Results of logit model determining compatibility between potential partners entering marriage

Variable	Coefficient	p-value
Age of male	0.013	0.455
Age difference (age of male – age of female)		
<i>Greater than 7</i>	-2.077	< 0.001
<i>From 4 to 7</i>	-0.865	0.066
<i>From 2 to 3</i>	0.194	0.758
<i>From -1 to 1</i>	0	
<i>From -3 to -2</i>	-1.643	0.006
<i>From -7 to -4</i>	-2.113	< 0.001
<i>Smaller than -7</i>	-19.220	0.977
Woman has child(ren)	0.352	0.811
Woman or man or both were married before	-0.892	0.042
Indirect effect		
<i>Woman has child(ren) & married before</i>	0.949	0.537
Number of pairs: 280		

higher order marriages) and the presence of children by introducing an accordant interaction term. The related effect is slightly significant in the case of cohabitation and insignificant in the case of marriage.

The considered setting surely simplifies actual partner choice patterns, in particular as it relates to the effect of smoking behavior on matching probabilities. Several studies exist that analyze and discuss such effects; see for example Clark and Etilé

(2006) and Chiappori et al. (2010). Here, positive assortative matching over smoking because of similar lifestyle preferences seems to be the most common finding. In accordance therewith, we extend the two empirical likelihood equations relating to marriage and to cohabitation by a term supporting matches of individuals with similar smoking traits. Concretely, we add the variable ‘smoking status’ featuring the two categories ‘one partner smokes and the other not’ and ‘non-smoker couple or dual smoker couple’. We assign an effect size of value zero to the first category and an effect size of value one to the second category. The dimension of this ex-post alignment is not based on empirical evidence but chosen to fit to the overall context. There is no doubt that this is only a minimal solution to account for positive assortative matching over smoking traits and the accordant values should be updated if better data are available. Overall, however, we deem our illustration sufficient to show the capability of our population model.

5.5.2 Results

We use a desktop machine equipped with an Intel(R) Core(TM) i7 Duo CPU, 2.80 GHz and 8 GB memory to run our simulation. Due to the large population size considered, the outcome of different simulation runs is very similar. Hence, the standard error due to Monte Carlo variability is negligible, and it is sufficient to concentrate on the results of one single run. During simulation all demographic events (births and deaths, and state transitions of individuals and couples) are tracked. The simulation provides information on the simulated life courses in four files:

1. an ASCII file containing the birth dates of all simulated individuals,
2. an ASCII file containing the dates of transitions and the corresponding destination states for all simulated individuals,
3. an ASCII file containing the dates when seeking individuals enter the partnership market, their current state, their next state, their desired mating time, as well as
4. an ASCII file containing the dates of transitions and the corresponding destination states for all simulated couples.

These files have a well-defined format, which can be accessed and managed further by arbitrary tools. We use R to summarize and visualize the simulation output.

In the run considered, the simulation conducts 294,484 events and creates 67,743 newborns. Furthermore, 38,314 individuals are seeking for a partner. 37,046 of them find a partner, that is, 18,523 couples are built. Approximately three percent of all seekers cannot find a proper spouse along simulation time. This flaw is caused by a surplus of ‘mating-minded’ women in our population. (A way to counteract this unbalance would be to let proper spouses immigrate.) Subsequently, we show a few descriptive statistics for simulated couple data (based on one single simulation run). The results shown here are not comprehensive and only aim at indicating the potential of our microsimulation model. As a first step we convert the simulation

ID	IDFem	IDMale	BirthDateFem	BirthDateMale	TrDate	AgeAtTrFem	AgeAtTrMale	NewState
3296228	1481858	14860	(Aug/29/1986 13:56:19)	(Oct/01/1986 14:56:19)	(Sep/21/2008 02:04:52)	22.06	21.97	(fem. CO. nS. noChild. male. CO. S)
3296228	1481858	14860	(Aug/29/1986 13:56:19)	(Oct/01/1986 14:56:19)	(Nov/07/2009 10:55:34)	23.19	23.10	(fem. SI. nS. noChild. male. SI. S)
3296264	1909901	1923943	(Apr/29/1956 14:56:19)	(May/15/1957 14:56:19)	(Dec/25/2008 07:36:53)	52.66	51.61	(fem. CO. S. child. male. CO. nS)
3297833	944182	918172	(Apr/22/1987 13:56:19)	(Jul/17/1979 14:56:19)	(Aug/28/2008 02:23:23)	21.35	29.11	(fem. CO. S. noChild. male. CO. S)
3297833	944182	918172	(Apr/22/1987 13:56:19)	(Jul/17/1979 14:56:19)	(Mar/20/2019 01:36:36)	31.91	39.67	(fem. MA. S. noChild. male. MA. S)
3297833	944182	918172	(Apr/22/1987 13:56:19)	(Jul/17/1979 14:56:19)	(Jul/28/2019 09:53:58)	32.26	40.03	(fem. MA. S. child. male. MA. S)
3381993	1505012	1503091	(Apr/06/1981 13:56:19)	(Mar/14/1974 14:56:19)	(Sep/04/2009 11:54:08)	28.41	35.47	(fem. CO. nS. child. male. CO. nS)
3381993	1505012	1503091	(Apr/06/1981 13:56:19)	(Mar/14/1974 14:56:19)	(Sep/03/2014 20:55:11)	33.41	40.47	(fem. SI. nS. child. male. SI. nS)
3499462	1129652	1129210	(May/11/1992 13:56:19)	(Nov/09/1987 14:56:19)	(Jul/29/2010 18:05:14)	18.22	22.72	(fem. CO. nS. noChild. male. CO. nS)
3499462	1129652	1129210	(May/11/1992 13:56:19)	(Nov/09/1987 14:56:19)	(Apr/27/2012 17:36:34)	19.96	24.47	(fem. CO. S. noChild. male. CO. nS)
3499462	1129652	1129210	(May/11/1992 13:56:19)	(Nov/09/1987 14:56:19)	(Oct/04/2015 18:09:34)	23.40	27.90	(fem. CO. nS. noChild. male. CO. nS)
3499462	1129652	1129210	(May/11/1992 13:56:19)	(Nov/09/1987 14:56:19)	(Apr/13/2017 03:03:41)	24.92	29.42	(fem. CO. nS. child. male. CO. nS)
3499462	1129652	1129210	(May/11/1992 13:56:19)	(Nov/09/1987 14:56:19)	(May/04/2018 04:18:51)	25.98	30.48	(fem. MA. nS. child. male. MA. nS)
3499515	1196394	1174260	(Jan/26/1987 14:56:19)	(Mar/12/1986 14:56:19)	(Aug/28/2010 19:27:48)	23.59	24.46	(fem. CO. nS. noChild. male. CO. S)
3499515	1196394	1174260	(Jan/26/1987 14:56:19)	(Mar/12/1986 14:56:19)	(May/19/2012 20:27:29)	25.31	26.19	(fem. SI. nS. noChild. male. SI. S)
3425577	1152670	1146210	(Mar/06/1990 14:56:19)	(Feb/25/1978 14:56:19)	(Oct/18/2009 08:49:48)	19.62	31.64	(fem. CO. nS. noChild. male. CO. nS)
3425577	1152670	1146210	(Mar/06/1990 14:56:19)	(Feb/25/1978 14:56:19)	(Jun/27/2012 00:51:59)	22.31	34.33	(fem. CO. nS. child. male. CO. nS)
3425577	1152670	1146210	(Mar/06/1990 14:56:19)	(Feb/25/1978 14:56:19)	(Aug/25/2012 13:57:23)	22.47	34.50	(fem. SI. nS. child. male. SI. nS)
3425666	27661	258963	(Feb/19/1980 14:56:19)	(Feb/10/1975 14:56:19)	(Nov/22/2009 03:40:35)	29.76	34.78	(fem. CO. nS. child. male. CO. nS)

Fig. 5.6 The simulated life courses of eight couples

Table 5.7 Frequency distributions of age differences between partnered men and women (age of man minus age of woman), given according to age intervals used in modeling spousal compatibility

		[-1, 1]	(1, 3]	(3, 7]	[-3, -1]	[-7, -3)
Cohabitation	Simulated	0.29	0.21	0.29	0.12	0.10
	Observed	0.32	0.25	0.29	0.08	0.07
Marriage	Simulated	0.70	0.08	0.11	0.05	0.05
	Observed	0.73	0.08	0.13	0.03	0.03

output into a format resembling event history data. This format eases further computation. In Fig. 5.6 shows typical life courses of eight simulated couples. Each record shows an event that a couple has experienced during simulation. It gives the ID of the couple and the birth times of both spouses ('BirthDateFem' and 'BirthDateMale'). Furthermore, it contains the transition date ('TrDate'), the transition age of the female spouse ('AgeAtTrFem'), the transition age of the male spouse ('AgeAtTrMale'), and the state that the couple has entered ('NewState'). The first transition of a couple corresponds to the onset of the marriage or cohabitation. If a couple experiences a dissolution event, the spouses return to the population of single individuals. Likewise, in case of a widowhood event, the surviving spouse is handled as a single individual.

It is essential for the usefulness the proposed mate matching strategy that it resembles actual characteristics of partners in couples. Therefore, in order to validate our mate matching strategy, we analyze the distribution of age differences of cohabiting and married spouses (age of male minus age of female). We find that the simulated and observed frequency distributions are very similar. Consequently, we deem the proposed mate matching algorithm suitable to produce reasonable results.

Having a partner who smokes can influence the spouse's initiation of smoking. That is, a smoking spouse might incite his/her non-smoking partner to start smoking, or prevent his/her smoker partner from quitting smoking. It is also possible that a nonsmoking partner urges his/her spouse to stop smoking. Likewise, the presence of children has very likely a strong impact on a person's smoking behavior. We study




Fig. 5.7 Histograms of age differences between partnered men and women (age man minus age woman in years). *Black bars* mark observed data and *grey bars* depict simulated results

Table 5.8 Percentage of male and female spouses who quit smoking, conditioned on the partner’s smoking status and the presence of children

	Childless couple		Couple with children	
	Nonsmoking partner (%)	Smoking partner (%)	Nonsmoking partner (%)	Smoking partner (%)
Males	6.42	4.72	5.95	5.30
Females	8.04	5.64	11.71	6.45

whether in our simulation output we can find accordant behavior. For this purpose, we determine how many partnered smokers quit smoking, conditional on the smoking status of the partner and the presence of children. Table 5.8 shows the results.

That is, during simulation 6.42 % of the smoking males who are part of a childless couple and who are partnered with a nonsmoking woman quit smoking. This is contrasted by 4.72 % of the smoking men without children being partnered with a smoking woman. A similar pattern is evident for smoking women without children. In couples with children, the effect of the smoking trait of the partner is more pronounced for females than for males. Generally, women seem to have a stronger propensity to stop smoking. Only few spouses start smoking during simulation, in total 0.88 % of the male and 0.86 % of the female spouses. Table 5.9 gives the percentages of female and male spouses who start smoking, conditioned on the partner’s smoking behavior and the presence of children. We find that 1.18 % of the nonsmoking male spouses who are part of a childless couple and who are partnered with a smoking woman start smoking. In contrast, only 0.57 % of the nonsmoking males who have children and who are partnered with a nonsmoking woman start smoking. Almost no mothers start smoking. This is opposed by 3.11 % of the female childless spouses who are partnered with a smoking man and who start smoking. Overall, both the presence of children and the smoking behavior of the partner have a significant effect on an individual’s propensity to quit or to start smoking. All these results are in accordance with the input transition rates for smoking behavior.

Table 5.9 Percentage of male and female spouses who start smoking, conditioned on the partner’s smoking status

	Childless couple		Couple with children	
	Nonsmoking partner (%)	Smoking partner (%)	Nonsmoking partner (%)	Smoking partner (%)
Male	1.01	1.18	0.57	1.06
Female	1.14	3.11	0.01	0.03

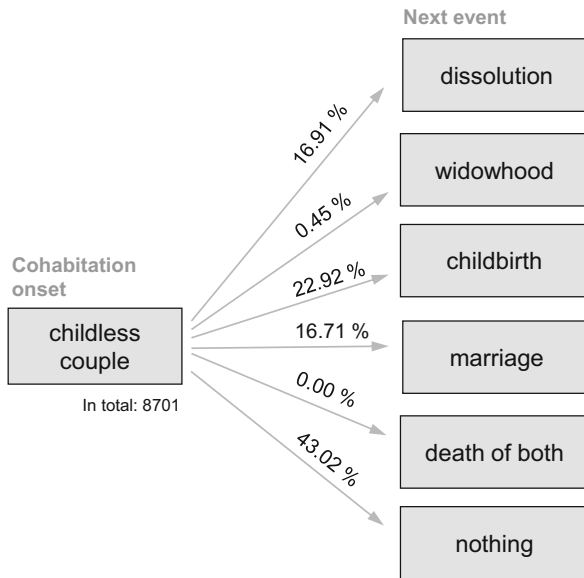


Fig. 5.8 First event after entering cohabitation

As a further aspect, we study the behavior of childless couples once they have launched a cohabitation. In sum, during simulation 17,402 individuals undergo a cohabitation event. We study dissolution and childbirth events, and restrict our consideration to the first event that happens after partnership onset. Figure 5.8 shows the respective results. After having entered cohabitation, 22.92 % of the childless couples experience a childbirth event, 16.71 % marry, 16.91 % undergo a dissolution event, and 0.45 % survive their partner.

As mentioned before, numerous summaries and insights can be derived from the simulation output of a population model as the one described, but here we restrict ourselves to the few examples given above. Again, please note that the results shown are only illustrative and do not reflect observed behavior, because they are based on input rates and model assumptions that are not solely derived from empirical data and established social theories.

5.6 Summary and Conclusion

We describe a microsimulation model that accounts for individual life courses, as well as for couple life courses, and for mate matching. To define individual and couple behavior we use a continuous-time multi-state model, that is, we use a continuous-time microsimulation model. This kind of model specifies life courses as sequences of discrete events; hence it is a discrete event model. For sake of simplicity, we assume that the propensity of individuals and couples to experience certain demographic events depends only on their current state, on possible next states, on age and calendar time. Our microsimulation model performs mate matching. For this purpose, we apply a two-sided stochastic mate matching procedure. In our microsimulation, individual life courses are specified as sequences of waiting times to a next event. That is, we can determine in advance when individuals will experience the onset of a marriage or a cohabitation. The waiting time until such an event is used for scanning potential partners. Concretely, once an individual is scheduled to experience a partnership onset he or she enters the marriage market. Each individual remains in the market for a specific period of mate searching and matching. In order to build up synthetic couples in the market a two-fold stochastic approach is used. First, we assign to each individual a random value that captures the aspiration level regarding a partner. An empirical likelihood equation reveals the probability that a given woman and a given man would mate. Subsequently, we simulate a decision making process whether two individuals date each other applying individuals' aspiration levels and their mating probability. A couple is formed if a positive decision has been made and the timing of the couple's partnership event is consistent regarding their individual mating periods. Individuals that are inspected, but rejected, lower their level of aspiration.

For implementation purposes, we opt for a model formalism that supports discrete event simulation. That way we can rely on existing modeling and simulation tools as well as on up-to-date simulation methodology. Concretely, we use the ml-DEVS formalism to specify our microsimulation model and the modeling and simulation framework JAMES II to implement it. The corresponding ml-DEVS population model is designed so that a macro-DEVS model guides the onset and the dissolution of partnerships and deals with structural population changes such as migration and childbirth events. In accordance therewith, ml-DEVS micro components are used to map life course dynamics of individuals and couples. Furthermore, we design simulation semantics that facilitates an efficient execution of the population ml-DEVS model.

In order to illustrate the capability of the presented microsimulation, we run a simulation projecting a synthetic population based on the population of the Netherlands. We parameterize the model using micro data from different data sources, among others, data from Statistics Netherlands and EuroStat. We study partnership and smoking behavior; particularly, we analyze how partners influence each other's smoking behavior. The application indicates that the simulation model produces feasible results. However, it should be noted that our illustration does not reflect observed behavior. It only serves to highlight the potential of the developed approach.

Although our demographic microsimulation is capable to conduct realistic population projections at a very detailed level, it shows limitations. First, a general impediment for the usage of microsimulations is their demand for data. To run meaningful microsimulation applications a lot of micro data is needed—which might be hard to access. Second, using the DEVS metaphor to specify the microsimulation bears some problems. Although the ml-DEVS formalism allows specifying population dynamics in the intended way, it introduces some modeling overhead. The formalism is very rigid which results in a bulky and longish definition of the ml-DEVS population model. Unfortunately, currently no modeling formalism exists that entirely copes with the requirements of demographic simulation models. Nonetheless, this is work in progress, see Steiniger et al. (2014). Third, currently individuals enter the partnership market based on empirical marriage rates or based on rates indicating cohabitation propensities. To correctly determine partnership events, however, instead of marriage or cohabitation rates, rates indicating the willingness to mate would have to be used. We know of no data source that allows estimating such rates. A way to anyhow obtain those rates would be to hypothesize them based on external knowledge of the phenomenon. Fourth, so far we restrict our consideration to binary linkages, that is, our model contains only individuals and couples. From a pure technical point of view, the construction of interaction networks comprising more than two individuals such as families and households is straightforward: we model individuals who are linked as being part of a “larger simulation unit that contains all individual units that are mutually dependent” (Galler 1997, p.14). However, parameterizing such models can become very difficult. This is mainly due to the fact that factors driving interactions and causal relationships are widely unobservable.

Appendices

A. The ml-DEVS Population Model

The ml-DEVS microsimulation model consists of a macro model **Pop** comprising two types of micro components: individual models **I** and couple models **C**, see Fig. 5.3. The special structure of the considered microsimulation model implies that not the entire functionality of the original ml-DEVS formalism is covered. For example, we do not employ any couplings between micro models. Therefore, for reasons of clarity, we adapt the original ml-DEVS approach such that in the subsequent description we leave out any functionalities that we do not demand.

A.1 The Population Macro Model

We formulate the population macro model **Pop**^{*} as structure

$$\mathbf{Pop} = \langle X, Y, S, s_{\text{init}}, \mathbf{I}, \mathbf{P}, \delta, \lambda_{\text{down}}, \lambda, sc \rangle$$

where the input port comprises the following information

$$X = \begin{cases} [\psi_1^I, \dots, \psi_{n_1}^I, \psi_1^C, \dots, \psi_{n_2}^C] & \text{if } n_1 \text{ single immigrants and/or } n_2 \text{ couples enter,} \\ \emptyset & \text{otherwise,} \end{cases}$$

ψ_i^I describes the state of an immigrant ($i = 1, \dots, n_1$), and

ψ_j^C describes the state of an immigrating couple ($j = 1, \dots, n_2$),

Y is an output port for emigrants leaving the population,

S is the set of possible states of **Pop**, a state $s = [s_1, \dots, s_5]$, $s \in S$, indicates

- s_1 : whether the transition of an individual model leads to a structural model change; indicators of structural model changes invoked by individual models are $actionOfInd = \{immigrating, emigrating, dying, childbirth\}$,
- s_2 : whether the transition of a couple model leads to a structural model change; indicators of structural model changes invoked by couple models are $actionOfCouple = \{immigrating, emigrating, dying, childbirth, dissolution, widowhood\}$,
- s_3 : the two individuals last found to form a proper match, and their mating time, otherwise $s_3 = \emptyset$,
- s_4 : all searching individuals that were included in the latest mate matching round and could not properly be matched, otherwise $s_4 = \emptyset$,
- s_5 : two individual models that are due to form a couple, otherwise $s_5 = \emptyset$.

$s_{init} = [\emptyset, \dots, \emptyset]$ is the initial state of **Pop**,

$\delta : X \times S \times \mathcal{J} \times \mathcal{C} \rightarrow S$ is the state transition function of **Pop**, where \mathcal{J} is the index set of all individual models **I** and \mathcal{C} is the index set of all couple models **C**; δ is composed of four component functions:

- δ_1 : if individuals are immigrating, or an individual signalizes emigration, death, or a childbirth event, δ_1 updates the first component s_1 of the state of **Pop** accordingly, otherwise $s_1 = \emptyset$,
- δ_2 : if couples are immigrating, or a couple signalizes emigration, death, childbirth, dissolution, or a widowhood event, δ_2 updates the first component s_2 of the state of **Pop** accordingly, otherwise $s_2 = \emptyset$,
- δ_3 : if an individual model signalizes the onset of a partner search, then δ_3 executes a mate matching algorithm
 - if two individuals can be identified as forming a proper match, then δ_3 reports these two individuals and their corresponding mating time in s_3 ,
 - if no individuals can be identified as forming a proper match, δ_3 sets s_3 to \emptyset ;
 - besides this, δ_3 reports in s_4 all searching individuals who were inspected during the mate matching process and could not properly be matched;
 otherwise $s_3 = s_4 = \emptyset$,
- δ_4 : if two individuals signalize that their mating time is due, δ_4 reports this in s_5 , otherwise $s_5 = \emptyset$.

$\lambda_{\text{down}} : S \rightarrow 2^{\bigcup_{i \in \mathcal{I}} X_i}$ is the downward output function to inform individual models (via their input ports X_i)

- about upcoming mating times (reported by s_3), and
- searching individuals, who were inspected during the latest mate matching round and for whom no proper match could be found (reported by s_4), about lowering their aspiration level;

$\lambda : S \rightarrow Y$ is the output function; it forwards the states and the birth dates of emigrating individuals and couples to Y .

$sc : S \rightarrow \mathbf{I} \times \mathbf{C}$ is the structural change function working on the set \mathbf{I} of individual models and the set \mathbf{C} of couple models; sc is composed of seven component functions:

sc_1 : *creates individual models:*

if s_1 indicates immigrating individuals, sc_1 creates n_1 new individual models for immigrants,

if s_1 or s_2 indicate childbirth, sc_1 creates for each newborn an individual model,

sc_2 : *deletes an individual model:*

if s_1 indicates that an individual is dying or emigrating, sc_2 deletes the corresponding individual model,

sc_3 : *creates couple models:*

if s_2 indicates immigrating couples, sc_3 creates n_2 new couple models for immigrants,

sc_4 : *deletes a couple model:*

if s_2 indicates that a couple is emigrating or both partners of a couple die, sc_4 deletes the corresponding couple model,

sc_5 : *creates a couple model and deletes two individual models:*

if s_5 indicates a mating event, sc_5 creates a new couple model for the mating individuals and deletes the corresponding individual models,

sc_6 : *deletes a couple model:*

if s_2 indicates the dissolution of a couple: sc_6 deletes the concerned couple model and creates for the separating partners two individual models,

sc_7 : *creates an individual model and deletes a couple model:*

if s_2 indicates a widowhood event: sc_7 deletes the corresponding couple model and creates for the surviving partner an individual model; otherwise sc stays idle.

$ta : S \rightarrow \mathbb{R}_0^+ \cup \{\infty\}$ is the time-advanced function: $ta(s) = 0$ if at least one $s_i \neq \emptyset$ ($i = 1, \dots, 5$), and $ta(s) = \infty$ otherwise.

The macro model **Pop** handles structural changes such as adding and removing individual and couple models. Via its input X the population model receives information about incoming immigrants.¹⁷ Several immigrants may enter the population simultaneously, for example, family members or couples. Once **Pop** receives the information about immigration events, it creates as many individual and couple models as immigrate. In an analogous manner, **Pop** creates new individual models if individual or couple models report birth events. If individual or couple models indicate the occurrence of emigration events, **Pop** forwards their states and birth dates via its output port Y and removes the related models.¹⁸ Besides migration, the macro model **Pop** guides the onset of partnerships, that is, the creation of couple models. To this end, **Pop** performs the mate matching algorithm described in Sect. 5.3. The mate matching procedure involves all individuals who signalize their disposition to mate, that is, exhibit a related output port. If two individuals are found to form a proper couple, **Pop** records this in its state and informs the individuals concerned immediately ($ta(s) = 0$) about the upcoming mating time (by carrying out λ_{down}). These matched individuals receive the accordant information on their input port $X_i = \{foundMate\}$. Similarly, **Pop** instructs individuals who were unsuccessfully inspected during a mate matching round to lower their level of aspiration (via their input port $X_i = \{redAspLevel\}$). When two individual models signalize the due date of their partnership onset, **Pop** replaces them by a couple model. Accordingly, if a couple model informs **Pop** about a dissolution event, **Pop** replaces this couple model by two individual models describing the separated partners. Likewise, if a couple model signalizes a widowhood event, **Pop** replaces this couple model by one individual model that describes the surviving partner. Once **Pop** receives a message about a structural population change, it processes that information immediately (i.e., $ta(s) = 0$) and empties the output ports of the related micro models. Apart from immigration, Figs. 5.9 and 5.10 illustrate the different types of structural model changes that **Pop** carries out.

A.2 The Individual Micro Model

I is the set of all individual models I . We formulate I as structure

$$\langle X, Y, {}^I\Psi, \psi_0, p, \delta, \lambda, ta \rangle$$

where

X is the input port of I ; it is $X = \{foundMate, redAspLevel\}$;

¹⁷The generation of immigration events might rely on empirical data about immigration dates and number of immigrants or on hypothetical data.

¹⁸Modeling migration in this way allows us to extend the population model to study the migration behavior between different populations. For this purpose different population models would have to be coupled.

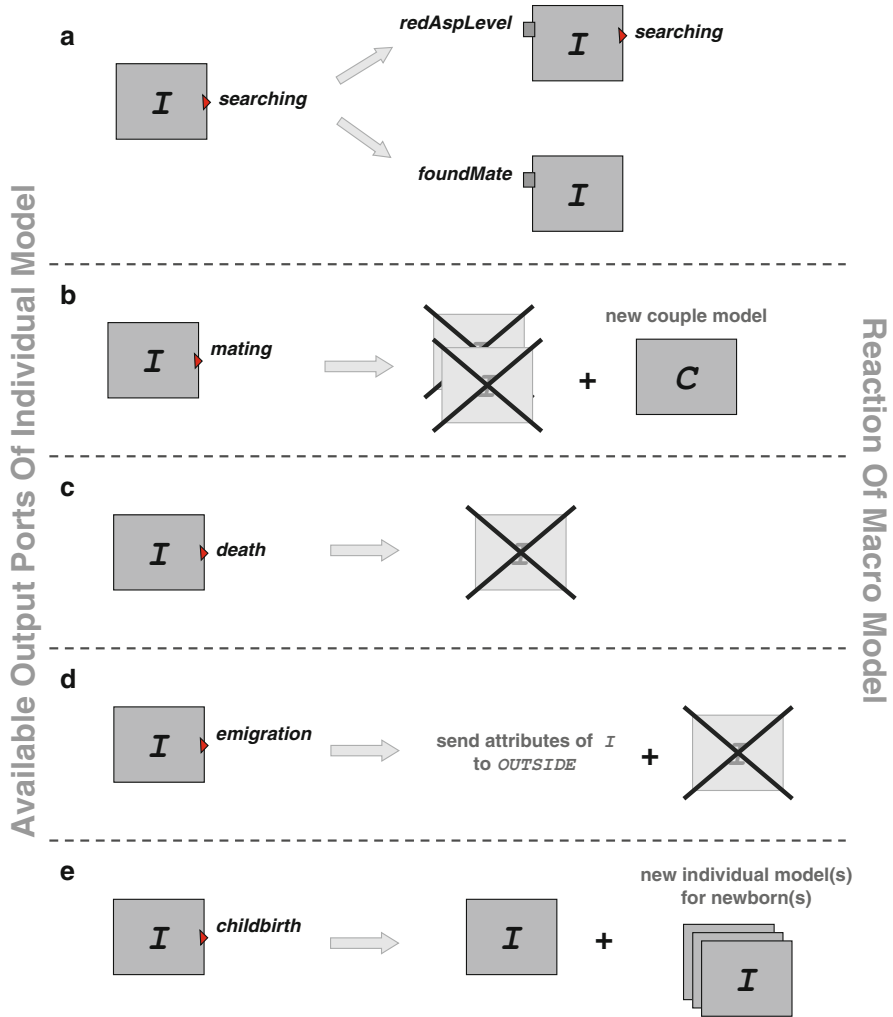


Fig. 5.9 Possible output messages of the individual models and reactions of the population macro model

1. **Pop** informs *I* via *foundMate* about an upcoming partnership onset, and
2. **Pop** instructs *I* via *redAspLevel* to lower the aspiration level.

Y is the set of output ports of *I*; $Y = \{searching, childbirth, emigration, death, mating\}$; we differ between two types of output ports:

1. the port *searching* that is permanently exhibited when *I* is searching for a mate, and
2. the output ports that indicate structural model changes:




Fig. 5.10 Possible output messages of the couple models and reactions of the population macro model

- (i) if a childbirth event is due, I signals this on the port *childbirth*,
- (ii) if a death event is due, I signals this event on the port *death*,
- (iii) if an emigration event is due, I signals this on the port *emigration*, and
- (iv) if a mating event is due, I signals this on the port *mating*.

${}^I\Psi$ is the set of possible states that I can occupy; ${}^I\Psi$ comprises

- (i) I 's current state s_c and I 's upcoming state s_u ($s_c, s_u \in \Psi$ where Ψ is the state space of the stochastic process that maps the individual life course),
- (ii) I 's birth date b and I 's age a ,
- (iii) the complete waiting time τ that I has to spend in s_c as well as
- (iv) I 's aspiration level la concerning a potential spouse ($la = \emptyset$ indicates that I is not searching for a mate).

$\psi_0 = [s_0, s_0, b, a_0, \infty, \emptyset]$, comprises I 's initial state s_0 , I 's birth date b , I 's age a_0 at model initialization, the entire waiting time that I has to spent in his/her current state (in the initialization phase we set this time to ∞), and the initial level of aspiration concerning a potential spouse is \emptyset .

To facilitate the subsequent description of the ml-DEVS functions δ , p , λ and ta , we define the following auxiliary functions:

$$\begin{aligned} search &: \Psi \times \Psi \rightarrow \{true, false\}, \\ death &: \Psi \rightarrow \{true, false\}, \\ emigration &: \Psi \rightarrow \{true, false\}, \\ childbirth &: \Psi \times \Psi \rightarrow \mathbb{N}_0. \end{aligned}$$

The function *search* shows whether a transition (from state s_c to state s_u) implies the onset of a partnership, the functions *death* and *emigration* indicate a death and an emigration event (identified by inspecting s_u), and the function *childbirth* gives the number of newborns that an event implicates (identified by inspecting the transition from s_c to s_u).

$\delta : X \times {}^I\Psi \rightarrow {}^I\Psi$ is the state transition function of I (x indicates input messages and $\psi \in {}^I\Psi$):

$$\delta(x, \psi) = \begin{cases} [s_u, s_n, b, a + \tau, \tau, \emptyset] & \text{if } search(s_u, s_n) = false \\ [s_u, s_n, b, a, \infty, la_0] & \text{if } search(s_u, s_n) = true \\ [s_c, s_u, b, a, \infty, \min(0, la - \delta_A)] & \text{if an input message has been received} \\ & \text{via the port } redAspLevel \\ [s_c, s_u, b, t^{on} - b, t^{on} - b - a, \emptyset] & \text{if the actual mating time } t^{on} \text{ has been} \\ & \text{received via the port } foundMate \\ \text{not defined} & \text{otherwise,} \end{cases}$$

where s_n is the next state of I and τ is the (random) waiting time between s_c and s_n , la_0 is the aspiration level at the moment when an individual enters the mate-searching phase, and δ_A is the decrement to lower the aspiration level in case of an unsuccessful search round.

$p : {}^I\Psi \rightarrow \mathcal{P}_I$ selects the output port available in a given state (with $\mathcal{P}_I = \{searching, mating, death, emigration, childbirth\}$ is the set of all possible output ports of I):

$$p(\psi) = \begin{cases} \textit{searching} & \text{if } \textit{search}(s_c, s_u) = \textit{true} \text{ and } la \neq \emptyset, \\ \textit{mating} & \text{if } \textit{search}(s_c, s_u) = \textit{true} \text{ and } la = \emptyset, \\ \textit{death} & \text{if } \textit{death}(s_u) = \textit{true}, \\ \textit{emigration} & \text{if } \textit{emigration}(s_u) = \textit{true}, \\ \textit{childbirth} & \text{if } \textit{childbirth}(s_c, s_u) > 0, \\ \emptyset & \text{otherwise,} \end{cases}$$

$\lambda : {}^I\Psi \rightarrow Y$ is the output function; it is composed by five component functions that fill the output ports *searching*, *mating*, *childbirth*, *death*, and *emigration*:

- λ_1 : If $\textit{search}(s_c, s_u) = \textit{true}$ and $la \neq \emptyset$, λ_1 forwards the information $(s_u^m, b, t^{\text{des}} = b + a + \tau^*, la)$ necessary for mate searching, where $s_u^m \subseteq s_u$ comprises the individual attributes that are relevant for finding a proper spouse (like partnership type and age), and t^{des} is the desired mating time; otherwise $\lambda_1 = \emptyset$.
- λ_2 : If I is due to mate (i.e., $\textit{search}(s_c, s_u) = \textit{true}$ and $la = \emptyset$), λ_2 forwards via the output port *mating* the upcoming state and the birth date of I , otherwise $\lambda_2 = \emptyset$.
- λ_3 : If I is due to give birth (i.e., $\textit{childbirth}(s_c, s_u) > 0$), λ_3 forwards via the output port *childbirth* the number of newborns, otherwise $\lambda_3 = \emptyset$.
- λ_4 : If I is due to die (i.e., $\textit{death}(s_u) = \textit{true}$), λ_4 signalizes this to the output port *death*, otherwise $\lambda_4 = \emptyset$.
- λ_5 : If I is due to emigrate (i.e., $\textit{emigration}(s_u) = \textit{true}$), λ_5 forwards via the output port *emigration* the current state and the birth date of I , otherwise $\lambda_5 = \emptyset$.

$ta : {}^I\Psi \rightarrow \mathbb{R}_0^+ \cup \{\infty\}$ is the time-advanced function: $ta(\psi) = \tau$.

The state of I captures all the information necessary to describe I 's attributes and behavior. It comprises I 's current state s_c and I 's upcoming state S_u as well as I 's birth date and age when entering state s_c . Additionally, it contains the information about I 's aspiration la concerning a potential spouse if I is in a mate-searching phase. The occurrence of an event (provoked by a state transition of the corresponding stochastic process) that does not imply the onset of a partnership results in an (ordinary) state transition of I . Therefore, I 's current state, upcoming state, and age at last transition is redefined accordingly. If the event implies the onset of a partnership, I immediately enters a phase of mate searching. In this phase, I features a positive aspiration level la concerning the traits of a potential spouse and permanently exhibits his/her *searching* output port. This port holds I 's birth date, the desired mating time, and the attributes of I that are relevant for the mate matching procedure, for example, whether I wants to marry or enter a cohabitation. The permanent exhibition of the *searching* port ensures that every time **Pop** is conducting a mate matching round, it can retrieve those data that are relevant for

mating. Meanwhile, every time **Pop** considers (the searching) I in a mate matching round and no proper spouse could be detected for I , **Pop** instructs I (via its input port *redAspLevel*) to lower his/her aspiration level. During each mate matching round **Pop** checks whether I 's desired mating time t^{des} is expired. If this is the case, the options listed in Sect. 5.3.2 might be applied. For reasons of convenience, we opt for matching I with the most compatible opposite-sex candidate who is seeking for the same partnership type. If no such candidate is available, we shift the preferred mating time by half a year, and pair off I as soon as possible. If a proper spouse for I is then found, **Pop** informs I and the new spouse about the upcoming mating event. In line with this, I receives an accordant message on its input port *foundMate*. Besides the *searching* port, I possesses the four other output ports: *death*, *emigration*, *childbirth*, and *mating*. The port *death* signals that I makes a transition to 'death'. Likewise, *emigration* shows an emigration event, and *childbirth* indicates the number of children that a transition implicates. On the *mating* port I signals that he/she enters a partnership. Figure 5.9 displays the output ports of I . It further illustrates the operations that **Pop** conducts in response.

A.3 The Individual Couple Model

C is the set of all couple models C . We formulate C as structure

$$\langle Y, {}^C\Psi, \psi_0, \delta, \lambda, ta \rangle$$

where

Y is the output port of C ; $Y = \{\textit{dissolution}, \textit{widowhood}, \textit{death}, \textit{emigration}, \textit{childbirth}\}$; it forwards structural changes such as dissolution (port: *dissolution*), widowhood (port: *widowhood*), death of both partners (port: *death*), emigration (port: *emigration*), or childbirth (port: *childbirth*),

${}^C\Psi$ is the set of possible states that C can occupy; ${}^C\Psi$ comprises

- (i) C 's current state s_c and C 's upcoming state s_u ($s_c, s_u \in \Psi$ with Ψ being the state space of the stochastic process that maps the dynamics of the couple),
- (ii) the birth date b_m of the male spouse and the birth date b_f of the female spouse
- (iii) the age a_m of the male spouse and the age a_f of the female spouse, as well as
- (iv) the complete waiting time τ that C has to spent in s_c .

$\psi_0 = [s_0, s_0, b_m, b_f, a_m, a_f, \infty]$, comprises C 's initial state s_0 , the birth date b_m of the male spouse, the birth date b_f of the female spouse, the age a_m of the male spouse at model initialization, and the age a_f of the female spouse at model initialization, as well as C 's waiting time in s_0 to which we assign ∞ in the initialization phase.

To ease the subsequent description of the ml-DEVS functions δ, p, λ and ta , we define the following auxiliary functions:

$$\begin{aligned}
dissolution &: \Psi \times \Psi \rightarrow \{true, false\}, \\
widowhood &: \Psi \times \Psi \rightarrow \{true, false\}, \\
deathOfBoth &: \Psi \rightarrow \{true, false\}, \\
emigration &: \Psi \rightarrow \{true, false\}, \\
childbirth &: \Psi \times \Psi \rightarrow \mathbb{N}_0.
\end{aligned}$$

The function *dissolution* indicates whether a state transition implies the dissolution of the partnership, the function *widowhood* indicates the death either of the male or the female spouse, and *deathOfBoth* shows the death of both partners; *emigration* indicates an emigration event, and the function *childbirth* gives the number of newborns that a state transition implicates.

$\delta : {}^C\Psi \rightarrow {}^C\Psi$ is the state transition function of C (with $\psi \in {}^C\Psi$):

$$\delta(\psi) = (s_u, s_n, b_m, b_f, a_m + \tau, a_f + \tau, \tau),$$

where s_n is the next event of C and τ is the (random) waiting time between the states s_u and s_n .

$p : {}^C\Psi \rightarrow \mathcal{P}_C$ selects the port available in a given state of the couple model (with $\mathcal{P}_C = \{dissolution, widowhood, death, emigration, childbirth\}$ being the set of all possible output ports of C):

$$p(\psi) = \begin{cases}
dissolution & \text{if } dissolution(s_c, s_u) = true, \\
widowhood & \text{if } widowhood(s_c, s_u) = true, \\
childbirth & \text{if } childbirth(s_c, s_u) > 0, \\
death & \text{if } deathOfBoth(s_u) = true, \\
emigration & \text{if } emigration(s_u) = true, \\
\emptyset & \text{otherwise,}
\end{cases}$$

$\lambda : {}^C\Psi \rightarrow Y$ is the output function; it is composed of five component functions that fill the output ports *dissolution*, *widowhood*, *childbirth*, *death* and *emigration*:

- λ_1 : If C is due to dissolve (i.e., $dissolution(s_c, s_u) = true$), λ_1 forwards via the output port *dissolution* the upcoming state of C and the birth dates of both spouses, otherwise $\lambda_1 = \emptyset$.
- λ_2 : If C is due to experience a widowhood event (i.e., $widowhood(s_c, s_u) = true$), λ_2 forwards via the output port *widowhood* the state and the birth date of the surviving partner, otherwise $\lambda_2 = \emptyset$.
- λ_3 : If C is due to experience a childbirth event (i.e., $childbirth(s_c, s_u) > 0$), λ_3 forwards via the output port *childbirth* the number of newborns, otherwise $\lambda_3 = \emptyset$.

- λ_4 : If C is due to experience the death of both partners (i.e., $deathOfBoth(s_u) = true$), λ_4 signals this to the output port $death$, otherwise $\lambda_4 = \emptyset$.
- λ_5 : If C is due to emigrate (i.e., $emigration(s_u) = true$), λ_5 forwards via the output port $emigration$ the current state of C and the birth dates of both spouses, otherwise $\lambda_5 = \emptyset$.

$ta : {}^C\Psi \rightarrow \mathbb{R}_0^+ \cup \{\infty\}$ is the time-advanced function: $ta(\psi) = \tau$.

The specification of the couple model C is very similar to the specification of the individual model I . It records in its state the attributes, the ages, and the birth dates of the female and the male spouse. Output ports are used to inform the macro model **Pop** about structural changes. State transitions are specified in the same way as in the case of I . To inform **Pop** about structural changes, the couple model C features five output ports (*dissolution*, *widowhood*, *death*, *emigration*, and *childbirth*). If C experiences a dissolution event, it forwards via the port *dissolution* to **Pop** the birth dates and the upcoming attributes of both spouses. In response, **Pop** creates for the separated partners two individual models, and deletes C . The processing in case of a widowhood event is similar. The only difference is that the individual model is created for the surviving partner only. If C experiences an emigration event, it forwards to **Pop** the current attributes and the birth dates of the female and the male spouse. In response **Pop** forwards C 's data via its output port and deletes C . Equally, the death of both spouses causes the deletion of C . In case of a childbirth event, C forwards the number of newborns to **Pop**. **Pop** reacts by creating as many individual models as newborns have been reported. Figure 5.10 shows the output ports of C . The figure further illustrates the structural model changes that **Pop** conducts in response to activated ports.

B. Sequential Abstract Simulator for ml-DEVS

The sequential abstract simulator for ml-DEVS executes parts of the DEVS processor tree en bloc. Therefore, it implements the two methods *getOutputs* and *doRemainder*. During simulation processing these methods are successively called. In doing so, it still complies with the original ml-DEVS communication protocol, compare Figs. 5.11 and 5.12.

The simulator and coordinator of a ml-DEVS model realize the methods *getOutputs* and *doRemainder* differently, see Algorithms 1 and 2. If its *getOutputs* method is called, the macro model coordinator activates the simulator of all imminent micro




Fig. 5.11 The communication protocol between micro model simulators and the macro model coordinator of the original abstract simulator of ml-DEVS




Fig. 5.12 The communication protocol between micro model simulators and the macro model coordinator of the sequential abstract simulator of ml-DEVS

Algorithm 1 Pseudocode of the coordinator of the sequential abstract simulator for the ml-DEVS population model

```

function Messages getOutputs()
  for each microModel in imminents do
    micro.msgs = union (micro.msgs, microModel.getOutputs())
  if  $t$  is getTonie(macroModel)
    macro.msgs = macroModel.lambda
    macro.downMsgs = macroModel.lambdaDown
    msgs = union (micro.msgs,macro.msgs,macro.downMsgs)
  end if
  return msgs
end function

function (double, ports) doRemainder(msgs)
  if  $t$  is getTonie(macroModel)
    macro.downMsg = getMsgFromMacroLambda(msgs)
    influencees = union (influencees, sendMessagesToMicroModels(macro.downMsgs))
  else influencees is empty
  end if
  for each microModel in union (influencees, imminents) do
    (microModel.Tonies, ports)=microModel.doRemainder()
  macro.msgs = getMsgFromMacro(msgs)
  micro.msgs = getMsgFromMicro(msgs)
  msgs = union (micro.msgs,macro.msgs)
  if (msgs is not empty) or  $t$  is getTonie(macroModel)
    execute macroModel.stateTransition(msgs,ports)
    execute macroModel.timeAdvance
    macroModel.processStructuralChanges
    ports = macroModel.availablePorts
  end if
  tonie = min(microModel.Tonies, getTonie(macroModel))
  return union (tonie, ports)
end function
  
```

models (i.e., all micro models for whom the next event is due). Subsequently, it reads their output messages. If the current model time t corresponds to the next event time of the macro model (i.e., to the coordinator's actual *tonie*), the coordinator executes the output function of the macro model and its downward output function. Then, *getOutputs* forwards the output message of the macro model

Algorithm 2 Pseudocode of the simulator of the sequential abstract simulator for the ml-DEVS population model

```

function Messages getOutputs()
  execute model.lambda
  return getMsgs(model)
end function

function (double, ports) doRemainder(msgs)
  execute model.stateTransition(msgs)
  execute model.timeAdvance
  ports = model.availablePorts
  return union (tonie, ports)
end function

```

to the root-coordinator. The root-coordinator maintains that information for further processing.¹⁹ Afterwards, the *doRemainder* method of the macro model coordinator is called. As a first step, it identifies all micro models who are influenced by an external event (i.e., by an upcoming partnership event or the instruction to lower the aspiration level) and feeds that information into their input ports. Then, if the current model time t equals the coordinator's *tonie*, it executes all influenced micro models (via calling their *doRemainder* method). Thereafter, the coordinator executes all imminent models. For this purpose, it calls their *doRemainder* method. This method computes for all imminent models new next events—however, only if the models are not exposed to structural changes implying their extinction such as death or dissolution events. Then, the *doRemainder* method of the coordinator requests all (newly computed) next event times (*tonies*) as well as the currently exhibited ports of the micro models, and stores them. Subsequently, it checks whether it has received a message about an external event (i.e., about immigrants) or a messages about structural changes from its subordinate micro models. If so, or if its *tonie* is due, it executes the transition function and the time-advanced function of the macro model. Eventual reported structural model changes are subsequently executed. The *doRemainder* method exits by updating the available ports *ports* of all subordinate micro models, and by determining and forwarding the system's next event time. Note that the *searching* port of all individuals willing to mate is permanently exhibited by the related micro models. Only if an individual is paired off is it retracted.

The simulator of a micro model employs the method *getOutputs* to call the model's output function and to forward the respective output information, for example, death or mating events. By *doRemainder* the simulator performs the state transition of the model and computes its next internal event determined by its own state transition function. Finally, it forwards to the coordinator of the macro model the ports exhibited in the current state of the model and its next event time. Opposed

¹⁹For example, the root-coordinator might send information on emigrants to another coupled population model; see footnote 7.

to a direct implementation of the ml-DEVS processing scheme, the sequential abstract simulator for ml-DEVS requires only two threads to execute a ml-DEVS model: one for the macro model coordinator and one for the root-coordinator. Hence, we are not in danger of facing any limitations concerning population size due to the restricted numbers of parallel threads supported by the Java virtual machine being used.

C. Transition Rates Used in the Application

In this appendix we present the age-profiles of the transition rates used in the application shown in Sect. 5.5. The transition rates that describe the propensity of unlinked females and males to change the marital status are depicted in Figs. 5.13 and 5.14. Transition rates of unlinked individuals to change their smoking behavior and (non-parity specific) fertility rates of females are also depicted in Fig. 5.14. Figure 5.15 shows the transition rates of spouses to change their smoking behavior. Figure 5.16 depicts the log-mortality rates of female and male non-smokers and smokers. The transition rates of couples to change marital status are given in Fig. 5.17.

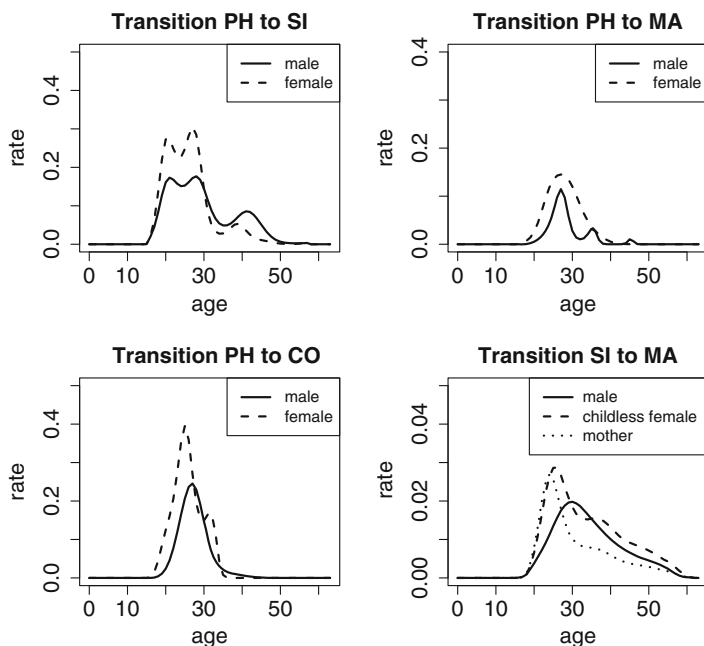


Fig. 5.13 Transition rates of unlinked females and males to change the marital status; *PH*: living at parental home & never married/cohabitating before, *SI*: being single & living alone, *MA*: married, *CO*: cohabiting

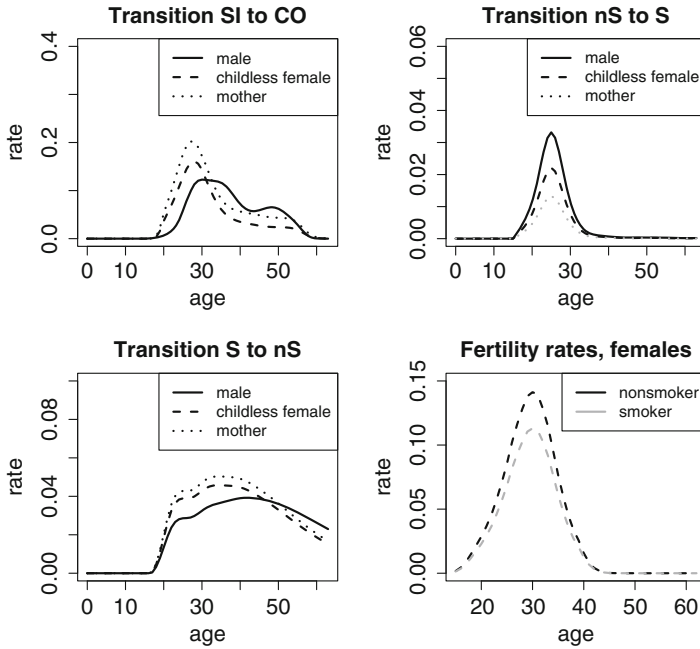


Fig. 5.14 Transition rates of unlinked females and males to change the marital status and smoking behavior, and fertility rates of females; *SI*: being single & living alone, *MA*: married, *CO*: cohabiting, *S*: smoker, *nS*: non-smoker

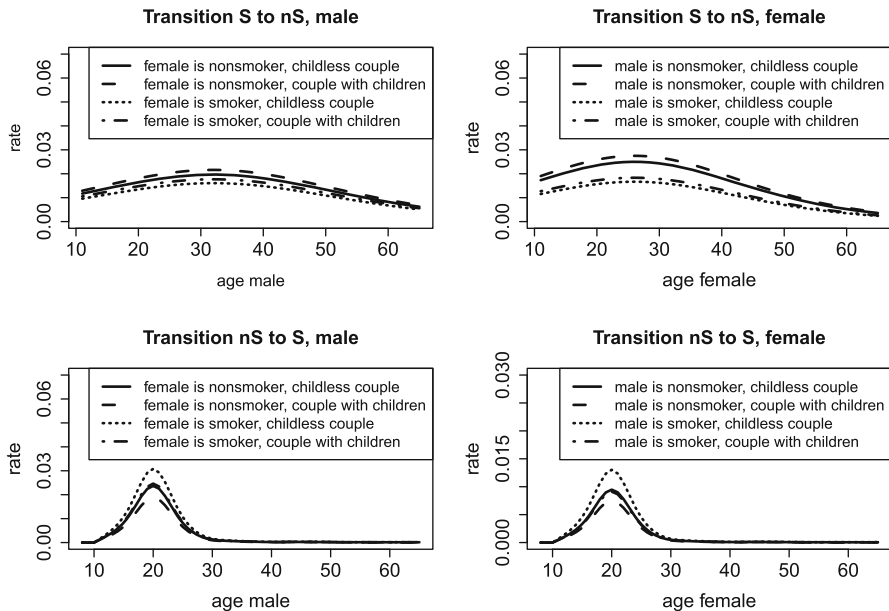


Fig. 5.15 Transition rates of spouses to change the smoking status; *S*: smoker, *nS*: non-smoker

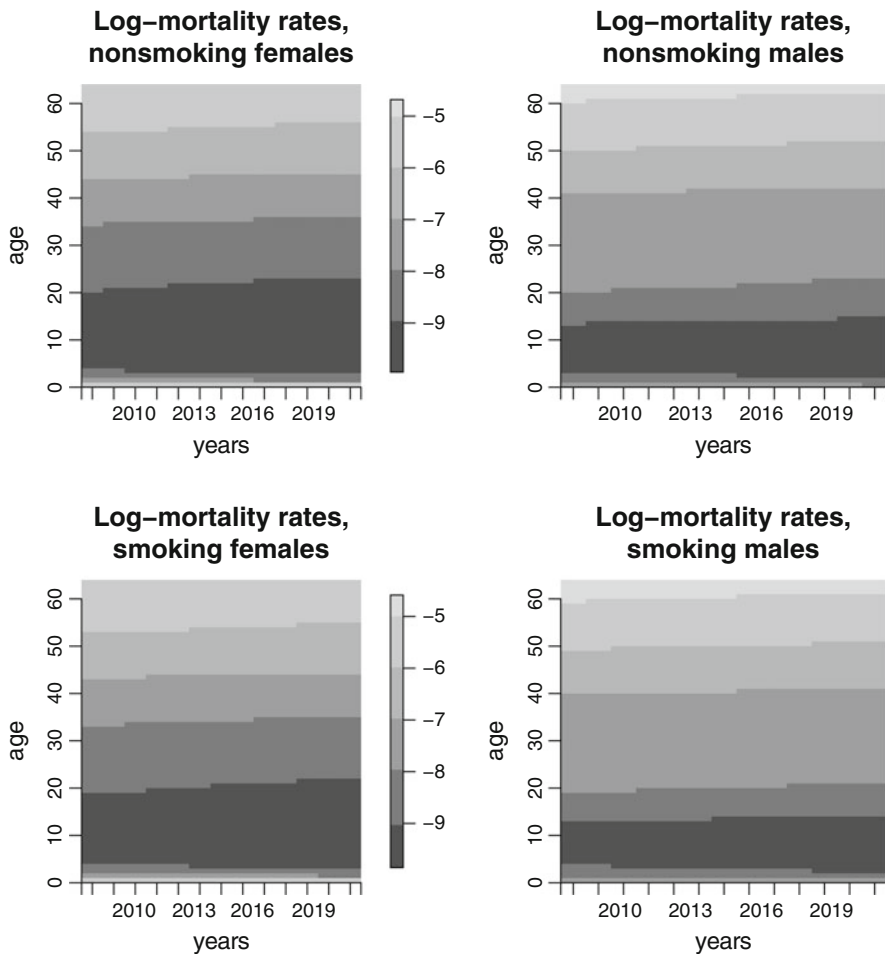


Fig. 5.16 Log-mortality rates by age, sex, and smoking status for the period 2008–2020 and ages from 0 to 63

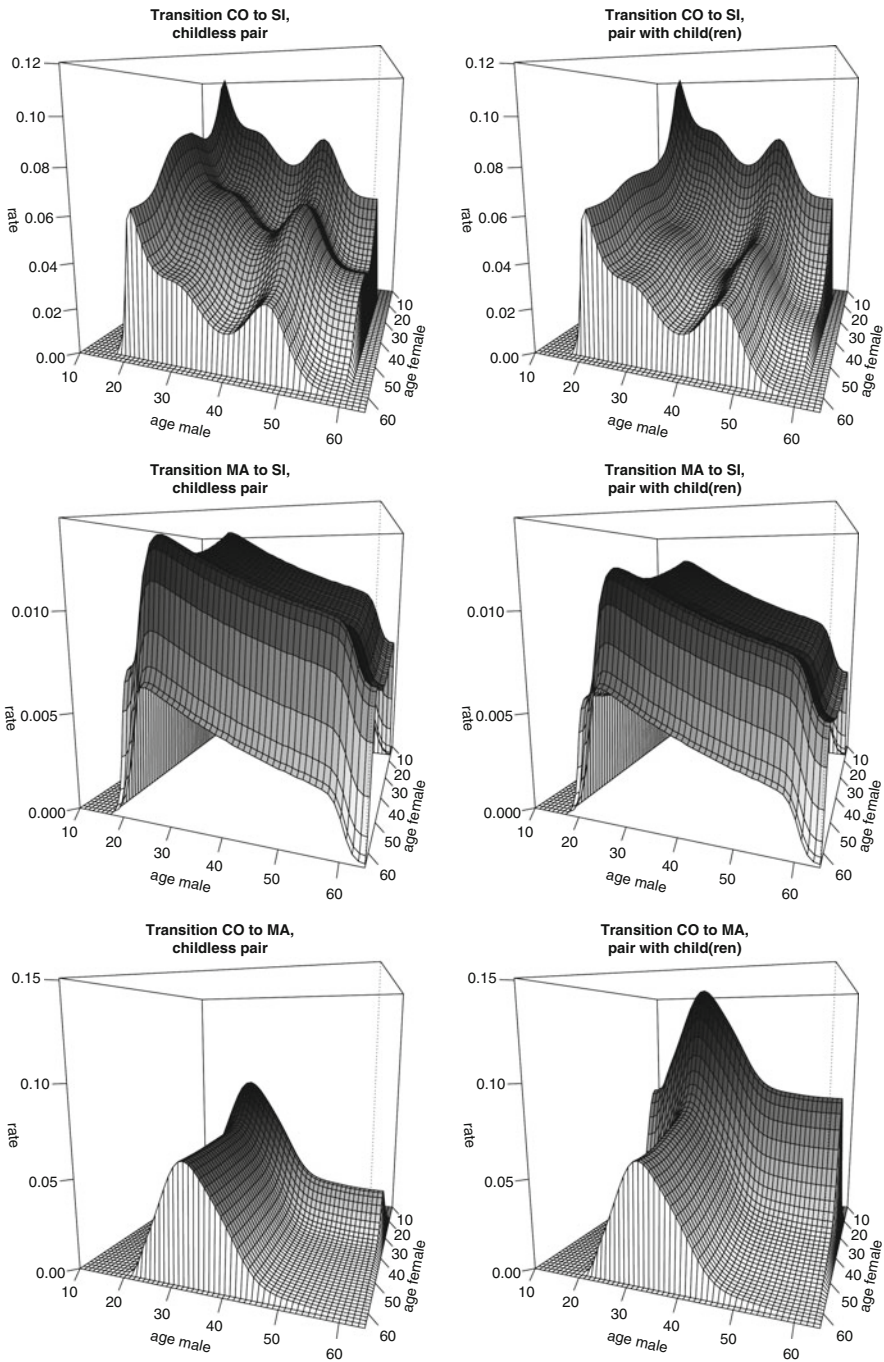


Fig. 5.17 Transition rates of couples to change the marital status; *MA*: married, *CO*: cohabiting, *SI*: dissolved

References

- Billari, F. C., Prskawetz, A., Aparicio Diaz, B., & Fent, T. (2008). The “wedding-ring”: An agent-based marriage model based on social interaction. *Demographic Research*, *17*, 59–82.
- Blossfeld, H.-P. (2009). Educational assortative marriage in comparative perspective. *Annual Review of Sociology*, *35*, 513–530.
- Buss, D. M. (2006). Strategies of human mating. *Psychological Topics*, *15*(2), 239–260.
- Cellier, F. E. (1991). *Continuous system modeling*. New York: Springer.
- Chiappori, P. A., Oreffice, S., & Quintana-Domeque, C. (2010). *Matching with a handicap: The case of smoking in the marriage market* (IZA Discussion Paper No. 5392). Retrieved from <http://ssrn.com/abstract=1731196>
- Clark, A. E., & Etilé, F. (2006). Don't give up on me baby: Spousal correlation in smoking behaviour. *Journal of Health Economics*, *25*(5), 958–978.
- de Sousa Campos, L., Otta, E., & de Oliveira Siqueira, J. (2002). Sex differences in mate selection strategies: Content analyses and responses to personal advertisements in Brazil. *Evolution and Human Behavior*, *23*(5), 395–406.
- de Vos, S., & Palloni, A. (1989). Formal models and methods for the analysis of kinship and household organisation. *Population Index*, *55*(2), 174–198.
- Deming, W. E., & Stephan, F. F. (1940). On a least square adjustment of a sampled frequency table when the expected marginal totals are known. *Annals of Mathematical Statistics*, *11*, 427–444.
- Doll, R., Peto, R., Boreham, J., & Sutherland, I. (2004). Mortality in relation to smoking: 50 years' observations on male British doctors. *BMJ*, *328*(7455), 1519.
- Galler, H. P. (1997). *Discrete-time and continuous-time approaches to dynamic microsimulation (reconsidered)* (Tech. Rep.). University of Canberra: NATSEM, National Center for Social and Economic Modelling, University of Canberra.
- Gill, R., & Keilman, N. (1990). On the estimation of multidimensional demographic models with population registration data. *Mathematical Population Studies*, *2*(2), 119–143.
- Hill, R. A., & Dunbar, R. I. M. (2003). Social network size in humans. *Human Nature*, *14*(1), 53–72.
- Himmelspach, J., & Uhrmacher, A. M. (2006). Sequential processing of PDEVS models. In A. Bruzzone, A. Guasch, M. Piera, & J. Rozenblit (Eds.), *Proceedings of the 3rd EMSS* (pp. 239–244), Barcelona.
- Hougaard, P. (1999). Multi-state models: A review. *Lifetime Data Analysis*, *5*(3), 239–264.
- Huinink, J., & Feldhaus, M. (2009). Family research from the life course perspective. *International Sociology*, *24*(299), 299–324.
- Impicciatore, R., & Billari, F. C. (2011). Maples: A general method for the estimation of age profiles from standard demographic surveys (with an application to fertility). *Demographic Research*, *24*, 719–748.
- Kalmijn, M. (1998). Intermarriage and homogamy: causes, patterns, trends. *Annual Review of Sociology*, *24*, 395–421.
- Mayer, K. U. (2009). New directions in life course research. *Annual Review of Sociology*, *35*, 413–433.
- Murphy, M. (2006). The role of assortative mating on population growth in contemporary developed societies. In F. C. Billari, T. Fent, A. Prskawetz, & J. Scheffran (Eds.), *Agent-based computational modelling: Applications in demography, social, economic and environmental sciences* (pp. 61–84). Heidelberg: Physica-Verlag.
- Prentice, R., & Pyke, R. (1979). Logistic disease incidence models and case-control studies. *Biometrika*, *66*(3), 403–411.
- Rubinstein, R. Y., & Kroese, D. P. (2008). *Simulation and the Monte Carlo method* (2nd ed.). New York: Wiley.
- Ruggles, S. (1993). Confession of a microsimulator: Problems in modeling the demographic kinship. *Historical Methods*, *26*(4), 161–169.

- Steiniger, A., Uhrmacher, A. M., Zinn, S., Gampe, J., & Willekens, F. (2014). The role of languages for modeling and simulating continuous-time multi-level models in demography. In A. Tolk, S. Diallo, I. Ryzhov, L. Yilmaz, S. Buckley, & J. Miller (Eds.), *Proceedings of the 2014 Winter Simulation Conference*, Savannah.
- Todd, P., & Billari, F. C. (2003). Population-wide marriage patterns produced by individual mate-search heuristics. In F. C. Billari & A. Prskawetz (Eds.), *Agent-based computational demography: Using simulation to improve our understanding of demographic behaviour* (pp. 118–137). Heidelberg/New York: Physica-Verlag.
- Todd, P., Billari, F. C., & Simão, J. (2005). Aggregate age-at-marriage patterns from individual mate-search heuristics. *Demography*, 42(3), 559–574.
- Trivers, R. (1972). Sexual selection and the descent of man, 1871–1971. In B. Campbell (Ed.), *Parental investment and sexual selection* (pp. 136–179). Chicago: Aldine.
- Uhrmacher, A. M., Ewald, R., John, M., Maus, C., Jeschke, M., & Biermann, S. (2007). Combining micro and macro-modeling in DEVS for computational biology. In *Proceedings of the 2007 Winter Simulation Conference*, Washington DC.
- Uhrmacher, A. M., Himmelspach, J., & Ewald, R. (2010). Effective and efficient modeling and simulation with DEVS variants. In G. Wainer & P. Mosterman (Eds.), *Discrete-event modeling and simulation: Theory and applications*. Boca Raton: CRC/Taylor and Francis.
- van Imhoff, E., & Post, W. (1998). Microsimulation methods for population projection. *Population: An English Selection, New Methodological Approaches in Social Sciences*, 10(1), 97–138.
- Wainer, G., Morihama, L., & Passuello, V. (2002). Automatic verification of DEVS models. In *Proceedings of the 2002 Spring Simulation Interoperability Workshop*, San Diego.
- Walker, L. (2010). *Modelling inter-ethnic partnerships in New Zealand 1981–2006: A census-based approach*. Ph.D. thesis, University of Auckland. Retrieved from <https://researchspace.auckland.ac.nz/handle/2292/5823>
- Waynforth, D., & Dunbar, R. (1995). Conditional mate choice strategies in humans: Evidence from 'lonely hearts' advertisements. *Behaviour*, 132(9), 755–779.
- Zeigler, B. P., Praehofer, H., & Kim, T. G. (2000). *Theory of modelling and simulation* (2nd ed.). New York: Academic Press. (Reissued by Krieger Pub. Co., Malabar, Fla. 1985.)
- Zietsch, B., Verweij, K., Heath, A., & Martin, N. (2011). Variation in human mate choice: Simultaneously investigating heritability, parental influence, sexual imprinting, and assortative mating. *The American Naturalist*, 177(5), 605.
- Zinn, S. (2011). *A continuous-time microsimulation and first steps towards a multi-level approach in demography*. Ph.D. thesis, University Rostock. Retrieved from http://rosdok.uni-rostock.de/file/rosdok_derivate_0000004766/Dissertation_Zinn_2011.pdf
- Zinn, S. (2012). A mate-matching algorithm for continuous-time microsimulation models. *International Journal of Microsimulation*, 5(1), 31–51.
- Zinn, S. (2015, January). *Demographic microsimulation for individuals and couples (version 1)*. CoMSES Computational Model Library. Retrieved from <https://www.openabm.org/model/4490/version/1>

Chapter 6

Using Survey Data for Agent-Based Modeling: Design and Challenges in a Model of Armed Conflict and Population Change

Nathalie E. Williams, Michelle L. O'Brien, and Xiaozheng Yao

6.1 Introduction

Agent-based modeling is a relatively new methodology that holds immense promise for demographic and social-behavioral research. Many of the methodological tools available to quantitative demographers and social scientists consist of statistical approaches that allow for precise modeling of micro or macro phenomena and can investigate important but essentially static relationships. In contrast, agent-based models (ABMs), using a complex systems approach, provide a method for examining dynamic interactions of social and demographic actors at both micro and macro levels. As such, ABMs provide a new perspective towards understanding an immense variety of outstanding questions in demography, such as how macro-level shocks like armed conflict, natural disaster, climate change, economic crises, and policy changes affect population growth and change. In other words, they can help us understand relationships between macro- and micro-level processes that are intimately and interactively linked. Given the scientific advances that are possible

N.E. Williams (✉)

Department of Sociology, University of Washington, Seattle, WA, USA

Jackson School of International Studies, University of Washington, Seattle, WA, USA

Center for Studies in Demography and Ecology, University of Washington, Seattle, WA, USA

e-mail: natw@uw.edu

M.L. O'Brien

Department of Sociology, University of Washington, Seattle, WA, USA

Center for Studies in Demography and Ecology, University of Washington, Seattle, WA, USA

e-mail: shannml@uw.edu

X. Yao

Virginia Tech University, Blacksburg, VA, USA

© Springer International Publishing Switzerland 2017

A. Grow, J. Van Bavel (eds.), *Agent-Based Modelling in Population Studies*,

The Springer Series on Demographic Methods and Population Analysis 41,

DOI 10.1007/978-3-319-32283-4_6

with this modeling strategy, it is encouraging to find an increasing number of studies, including those in this volume, that use ABMs to generate key demographic insights (Aparacio Diaz and Fent 2006; Aparacio Diaz et al. 2011; Billari et al. 2007; Bruch and Mare 2006; Ewert et al. 2003; Heiland 2003; Žamac et al. 2009).

A quick review of existing demographic studies that use ABMs reveals that the vast majority model one or possibly two demographic processes (such as marriage or fertility) in a largely hypothetical population or scenario. Such models are useful for providing broad insights into the relationship between particular demographic processes that might generally relate to any population around the world. However, they are less useful for understanding population change in any particular population or region and how multiple demographic, social, economic, and political processes interact.

Another type of ABM attempts to model specific populations, includes multiple interactive social and demographic processes, and uses empirical survey data to populate and operationalize model procedures. This strategy, commonly found in the geographic and land use sciences (An 2012; An et al. 2001, 2014; Manson 2005; Parker et al. 2008; Zvoleff and An 2014), is almost absent in demography and sociology. This modeling strategy, which we call survey data based ABM, can be characterized as a case study approach, providing more accurate insights into a single population, but possibly less relevance for other areas outside the study setting. A classic example of this type of model is one by An and Liu (2010) that analyzed how changes in fertility policies influenced destruction of panda habitat in China. The authors find that an increase in the number of households was more destructive to panda habitat than an increase in the population size overall.

Survey data based ABMs are also useful for testing the mechanisms that influence a particular outcome and the potential effects of policy changes, outstanding issues in many areas of demography. For example, beyond just looking at households and population size, the An and Liu (2010) study also models different hypothetical scenarios restricting the fertility rate, age at marriage, and the upper childbearing age. Using the simulated model, the authors were able to determine which changes to fertility would have the largest impact and which would have the quickest impact in reducing population growth.

Despite the many possible benefits, because these models utilize empirical data and are increasingly complex, they are immensely time consuming and cumbersome to design, test, and analyze. This might be one reason that they are seldom used in the demographic sciences to date. In this context, our purpose in this chapter is to develop methods and encourage the use of survey data based ABMs in demography. We present an ABM designed to analyze the impact of armed conflict on population change in rural Nepal. This model uses empirical survey data from the Chitwan Valley Family Study throughout the modeling process, for initializing a population and parameterizing behavioral rules. We discuss design challenges and suggest methods for addressing each challenge. It is our hope that this methodological presentation will help to streamline the long process of developing survey data based ABMs for future studies.

We write this chapter with the assumption that the reader has a basic understanding of ABMs. Given that ABMs are still poorly understood in the demographic and social behavioral literature, this assumption is clearly not reasonable. Regardless, instead of providing a lengthy description of ABMs, we guide the reader to reference existing literature on ABMs (An et al. 2005; Billari et al. 2007; Mena et al. 2011; Kniveton et al. 2011; Aparacio Diaz et al. 2011; Jackson et al. 2008; Heiland 2003; Ewert et al. 2003; Bruch and Mare 2006; Walsh et al. 2013).

6.2 Survey Data Based ABMs as Experimental Models

Before presenting our ABM of armed conflict and in order to understand the design that we do present, it is important to discuss the foundation of survey data based ABMs, what researchers should aim for, and what they can and cannot expect such an ABM to do.

ABMs that are informed by survey data can simulate real populations with real individual, household, and community characteristics and create behavioral rules that are modeled from real behaviors. However, they cannot, and should not be expected to model reality in all aspects. Social reality is much too complex to be modeled. Furthermore, social reality changes constantly; people's behaviors, beliefs, and relationships change; households change; communities change; weather, politics, and economies change. The primary problem with this almost constant, multidimensional, and non-linear change in social reality is that it makes it difficult to attribute differences over time in any one process to another particular process. For example, in Nepal, the context upon which our ABM is based, the conflict changed over time, with increasing and decreasing numbers of gun battles, states of emergency and ceasefires. At the same time, fertility rates decreased and household wealth generally increased. Thus, it is difficult to attribute a change in migration during the time period to the conflict, when fertility, wealth, and other processes varied significantly at the same time.

Instead of attempting to re-create reality, we argue that survey data based ABMs should be designed as experimental laboratories. From an analytical point of view, the ideal experiment is a situation with two or more groups, where everything is the same for the groups except one experimental factor.¹ Translating this concept to the ABM situation where the impact of a temporally changing factor is being tested, an ideal ABM would have relatively stable demographic and social processes, except for the one process of interest. Then, the researchers can simulate the model population with this process and without this process. By isolating the process, researchers can rule out spurious effects. Any difference in population growth between the two simulations would arguably be caused by the process being tested.

¹For a useful guide to experimental research design for the social sciences, see Adler and Clark (2008).

For example, with our ABM of armed conflict, the most desirable situation is one where most processes, such as marriage, birth, and migration are relatively stable over time. In this case, a simulation without the conflict acts as a control, while a simulation with the armed conflict acts as the experimental treatment. Differences between the two simulations can reveal the isolated influence of armed conflict on population change. Thus we designed our model with this ideal in mind, and sought to create a generally stable population using empirical survey data from a population that is not inherently stable.

6.3 Agent-Based Model for Armed Conflict and Population Change

The ABM we describe in this chapter is designed to investigate the effects of macro-level crises, such as armed conflicts, natural disasters, and economic crises, on population dynamics. We focus on armed conflict here, but with a variety of behaviors in the model, predictors of each behavior, and modules that include livelihoods, other crises can easily be simulated. Previous research has shown large impacts of armed conflict on individual demographic behaviors, such as marriage, childbearing, migration, and mortality (Agadjanian and Prata 2002; Apodaca 1998; Czaika and Kis-Katos 2009; Davenport et al. 2003; Eloundou-Enyegue et al. 2000; Gibney et al. 1996; Heuveline and Poch 2007; Jayaraman et al. 2009; Lindstrom and Berhanu 1999; Melander and Oberg 2006; Moore and Shellman 2004; Schmeidl 1997; Shemyakina 2009; Stanley 1987; Weiner 1996; Williams et al. 2012; Williams 2013, 2015; Winter 1992; Zolberg et al. 1989). Although these behaviors together comprise population change, a straightforward projection would arguably be inappropriate for estimating overall change in the population, because each of these behaviors affects other behaviors. For example, if someone migrates then they are less likely to get married or have children. People also interact. For example, if one person migrates, then they likely influence the probability of other household and community members migrating as well. Thus, if armed conflict affects most, if not all, people in a community and it influences all of these demographic behaviors, a complex interactive model is necessary to thoroughly examine the impact of armed conflict on population change. We discuss this further, and demonstrate the differences between single-behavior regression output and ABM results towards the end of this chapter.

6.3.1 Setting

Our model of armed conflict and population change is based on survey data from the western Chitwan Valley of south-central Nepal during the armed conflict of 1996–2006. The administrative district of Chitwan borders India and is about 100 miles from Kathmandu. As shown in Fig. 6.1, there is one large city, Narayanghat, and the

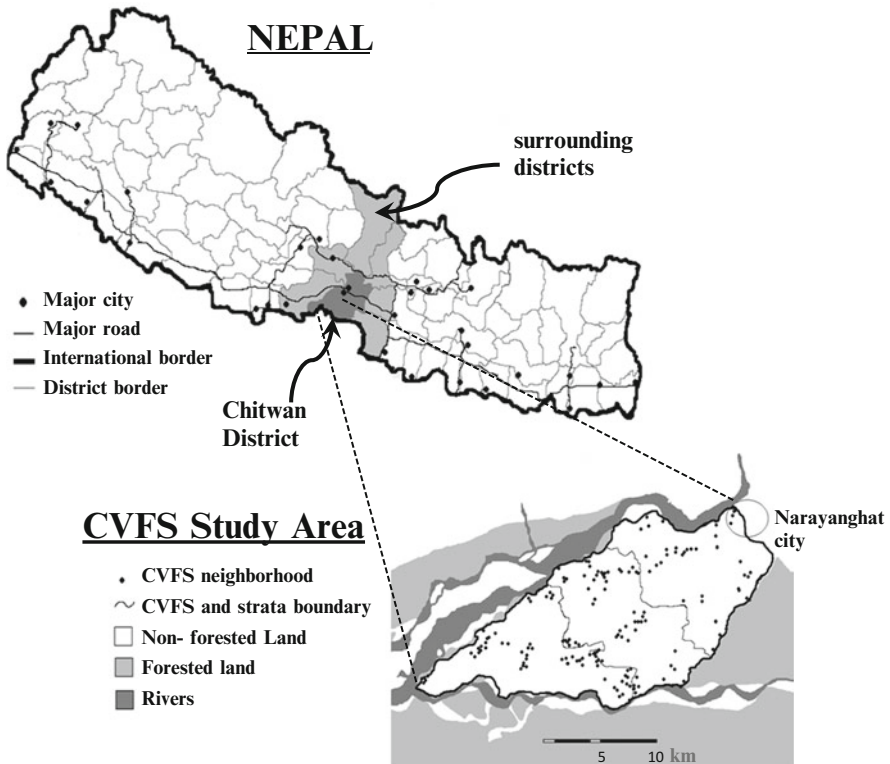


Fig. 6.1 Map of Nepal and Chitwan Valley Family Study area

rest of Chitwan's population, like much of Nepal, lives in small, rural villages. The valley is dominated by agriculture with about 80% of households using farming as their primary livelihood in 1996. Since then, market production of agricultural goods has increased, with continued production of agricultural goods for household consumption. From Narayanghat moving south-west, the study area is progressively more rural, poorer, and less involved in market agriculture.

Prior to the time period of this study, there was a massive proliferation of public and private services in Chitwan, including paved roads, off-farm employers, markets, schools, health services, banks, and many other services. This provision of services has led to increasing rates of education and non-family employment. Evidence also connects these community changes to drastically decreasing fertility rates, and changes in marriage and household residence patterns (Axinn and Yabiku 2001). Notably, childbearing still takes place almost exclusively within marriage (Axinn and Yabiku 2001; Ghimire and Axinn 2010).

Migration has also increased. The selectivity in who migrates has changed, and the destinations to which people migrate have expanded dramatically during this time period (Williams 2009; Williams et al. 2012; Massey et al. 2010). Migration

has long been a common livelihood strategy in Chitwan and continues to be mostly short-term and used to supplement regular farm incomes (Kollmair et al. 2006; Thieme and Wyss 2005). Nepal and India share an open border, so international migration to India, in addition to domestic migration, is common (United Nations 2013). Evidence shows that men have historically been more likely to migrate than women and that migration rates are lower among the married and people with children (Massey et al. 2010; Williams 2009). However, with political and economic changes on local and global scales in recent decades, migration to a variety of international destinations, most notably the Persian Gulf, has increased immensely and evidence suggests that gender and ethnic gaps in migration are decreasing at the same time (Williams 2009, 2015).

The armed conflict began in 1996 when the Communist Party of Nepal (Maoist) made a declaration of war with the intention to unseat the monarchy and install a people's republic. The early stages of the conflict were contained primarily in several midwestern districts and involved damage to government installations. From mid-2000, however, the Maoists progressively expanded their campaign across the country, including to Chitwan, and the Nepalese government responded by creating a special armed force to fight the Maoists. In 2006, the government and Maoists signed a comprehensive peace agreement declaring an end to the conflict.

The conflict was staged mainly as a guerrilla war. With no true "frontline," it was largely unknown where fighting would break out, and civilians were often caught up in violence. Reported violent acts by the Maoists and government forces against civilians include torture, assassinations, bombings, gun fights, abductions, forced conscription, billeting, taxing, and general strikes (Hutt 2004; Pettigrew 2004; South Asia Terrorism Portal 2006). A variety of political events also characterized this conflict, including states of emergency, ceasefires, depositions of the prime minister, and multiple nationwide strikes and protests that severely affected the day-to-day life of the general population and spread considerable unrest and fear nationwide. Evidence suggests that both violent and political events had significant influences on residents' marriage, contraception, and migration behaviors (Williams et al. 2012).

Figure 6.2 shows a timeline of the conflict and the violent and political events in the Chitwan Valley and surrounding districts. As you can see, there were relatively low levels of violence and political upheaval until 2002. This increased in 2003 through mid-2005, which was the height of the conflict. During this time, there were some gun battles, with up to 4 in 1 month, and even more bomb blasts reaching a high of 12 in 1 month.

6.3.2 Data

Data that were used to inform the ABM come from several sources. Survey data, which were used to create the initial population of the model and to operationalize the behavioral equations that define the probabilities of marriage, childbearing, death, migration, and other behaviors, come from the Chitwan Valley Family Study

consisted of multiple focus groups, informal and formal in-depth interviews, and observation in the study area. We collected this data from a variety of people that generally represented the gender, caste, and wealth distribution of the study area. The information we collected with this fieldwork includes household behaviors, such as splitting, and inheritance. Norms for and patterns of household splitting are notoriously difficult to measure and examine with existing survey data due to difficulties in defining and tracking households through time. Inheritance norms (such as when inheritances are given, to whom, and how much) are rarely collected by surveys. As such, our ABM rules for household inheritance and splitting are based on the ethnographic information we collected, instead of survey data. We also determined purchase prices of durable goods, livestock, and land, and selling prices of agricultural and livestock goods (such as meat and eggs) by creating price lists for multiple vendors of these items in the study site, as well as from government and NGO reports available on the internet.

6.3.3 Overview of the Agent-Based Model

In this section we provide a basic description of the ABM. Because the model is detailed and employs many different behaviors, full equations, decision-making processes and a full list of variables are provided online at www.bitly.com/NepalABM. Also online is a description following the ODD (Overview, Design concepts, Details) protocol that was designed explicitly to guide description of ABMs so that clear and comparable information could be provided for different models (Grimm et al. 2006, 2010). Additionally, we have included online a link to information about the software used to develop this model, *Repast Simphony*.

The model is organized through modules at the individual, household, and neighborhood levels. Each module contains a series of probabilistic decisions and deterministic processes that each individual, household, and neighborhood go through. Some modules take place on a monthly basis, and some on an annual basis. Figure 6.3 shows the overall structure of the ABM and the timing of each module.

We initialize the model based on data from the CVFS described in the previous section above. The agents in year 0 of the model have the characteristics of respondents in the data. Some new agents are created at the onset, because the CVFS did not interview the entire population, leaving out children and migrants who were away during the baseline survey. We discuss this detail later. The primary agents are individuals, who live in households, situated in 151 villages in the Chitwan Valley. Each agent undertakes behaviors based on probabilistic or deterministic equations that are described in more detail below.

Decision-making occurs on the individual level. Household and village characteristics affect individual decision-making, but are by and large aggregated from the individual characteristics, with few exceptions discussed below. Decisions occur at monthly and annual time points. At the beginning of each simulated time period (month or year), an individual or household module begins, moving agents

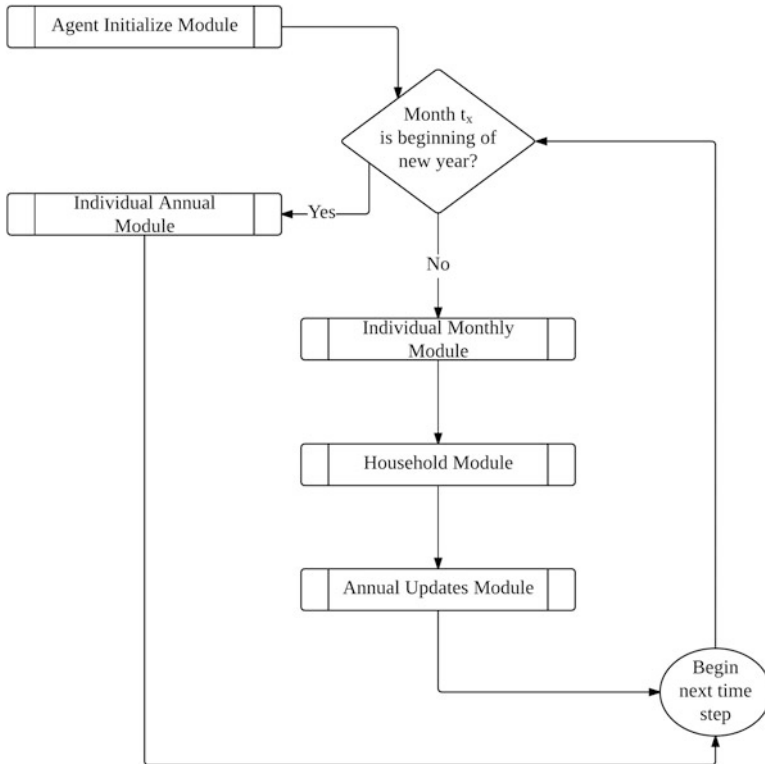


Fig. 6.3 Overall model flow with monthly and annual modules for individual, household, and neighborhood behaviors and processes

along a decision pathway. Agents and their collectives can undertake any of several demographic, social, and economic behaviors, including marriage, childbirth, out-migration, return-migration, death, and the accumulation of resources. Some modules simply add time onto a process, such as receiving one more month of pregnancy, or receiving one more year of education. Other modules are more complex and involve probabilistic decision-making, such as migration and marriage.

6.3.4 Decision-Making and Deterministic Processes

Processes in the model can be categorized as either probabilistic decisions or deterministic processes. Probabilistic decisions are based on regression equations estimated on the CVFS survey data. Each equation consists of multiple weighted factors, such as age, sex, and household assets, where the weights are determined by the coefficients from the regression equation. Figure 6.4 shows an example of such a probabilistic equation, which calculates the probability of out-migration for eligible men in the model. This equation results in a calculated probability of undertaking

Eligibility: At least 16 years old.

Probability of Out-Migration Men

$$\begin{aligned} \text{Log (P/(1-P))} &= (-7.9827) \\ &+ (\text{Number of gun battles this month} * (0.1225)) \\ &+ (\text{Number of bomb blasts this month} * (-0.0101)) \\ &+ (\text{State of Emergency this month} * (-0.1349)) \\ &+ (\text{Major government instability this month} * (0.668)) \\ &+ (\text{Major strikes and protests this month} * (-0.2558)) \\ &+ (\text{Number of gun battles this month} * \text{Had a salaried job} * (-0.1904)) \\ &+ (\text{Number of gun battles this month} * \text{Natural log of amount of land owned by household} * (0.0664)) \\ &+ (\text{Number of bomb blasts this month} * \text{Married and living with spouse} * (0.0386)) \\ &+ (\text{Number of bomb blasts this month} * \text{Married and living apart} * (0.00749)) \\ &+ (\text{State of Emergency this month} * \text{High caste} * (-0.109)) \\ &+ (\text{State of Emergency this month} * \text{Married and living with spouse} * (0.0526)) \\ &+ (\text{State of Emergency this month} * \text{Number of children} * (0.0294)) \\ &+ (\text{State of Emergency this month} * \text{Had a salaried job} * (-0.0169)) \\ &+ (\text{Major government instability this month} * \text{High caste} * (-0.175)) \\ &+ (\text{Major government instability this month} * \text{Married and living with spouse} * (-0.4733)) \\ &+ (\text{Major government instability this month} * \text{Number of children} * (0.0762)) \\ &+ (\text{Major government instability this month} * \text{Had a salaried job} * (-0.0425)) \\ &+ (\text{Major government instability this month} * \text{Distance to nearest urban area} * (-0.0386)) \\ &+ (\text{Major strikes and protests this month} * \text{Number of organization in neighborhood} * (0.1088)) \\ &+ (\text{Age} * (-0.0359)) \\ &+ (\text{High caste} * (-0.2088)) \\ &+ (\text{Education} * (0.0253)) \\ &+ (\text{Married and living with spouse} * (0.0655)) \\ &+ (\text{Married and living apart} * (-1.3785)) \\ &+ (\text{Widowed} * (0.257)) \\ &+ (\text{Number of children} * (-0.0125)) \\ &+ (\text{Had a salaried job} * (0.3053)) \\ &+ (\text{Natural log of amount of land owned by household} * (0.0252)) \\ &+ (\text{If household in bottom third of assets distribution} * (-0.1133)) \\ &+ (\text{If household in bottom third of income distribution} * (0.2791)) \\ &+ (\text{If household in middle third of income distribution} * (0.2006)) \\ &+ (\text{Distance to nearest urban area} * (0.0257)) \\ &+ (\text{Number of organizations in neighborhood} * (-0.0566)) \\ &+ (\text{Natural log of percent migrants in neighborhood} * (1.1578)) \\ &+ (\text{Ever migrated before start of model} * (0.2093)) \\ &+ (\text{Number of migrations since start of model} * (0.3514)) \\ &+ (\text{Number of months back from most recent migration trip} * (0.00206)) \\ &+ (\text{Number of months away on most recent migration trip} * (0.0299)) \\ &+ (\text{July} * (-0.028)) \\ &+ (\text{August} * (0.2451)) \\ &+ (\text{September} * (0.0615)) \\ &+ (\text{October} * (-0.4897)) \end{aligned}$$

Fig. 6.4 Migration equation in the ABM

migration during that month. The probability is compared to a random number between 0 and 1. If the probability is greater than the random number, the agent takes the action.

Deterministic processes do not involve a probability schema, but are calculated monthly or annually in the model. For example, at birth individuals are assigned an educational attainment level of 2 years greater than their parents. When the individual reaches the attainment level, he or she cannot accrue any more education. Income is also deterministic, where a household accrues a specified amount of income, depending on household members' migration, salaried employment, and land and livestock holdings.

6.3.5 Individual Module

At the beginning of each simulated month, the individual module, shown in Fig. 6.5, begins. Each individual experiences the possibility of death, based on his or her age. If the individual dies, they are removed from population and model. If the agent lives, then the model checks if he or she is married. If the agent is not married, then he or she experiences the possibility of marriage. All marriage is exogamous to the

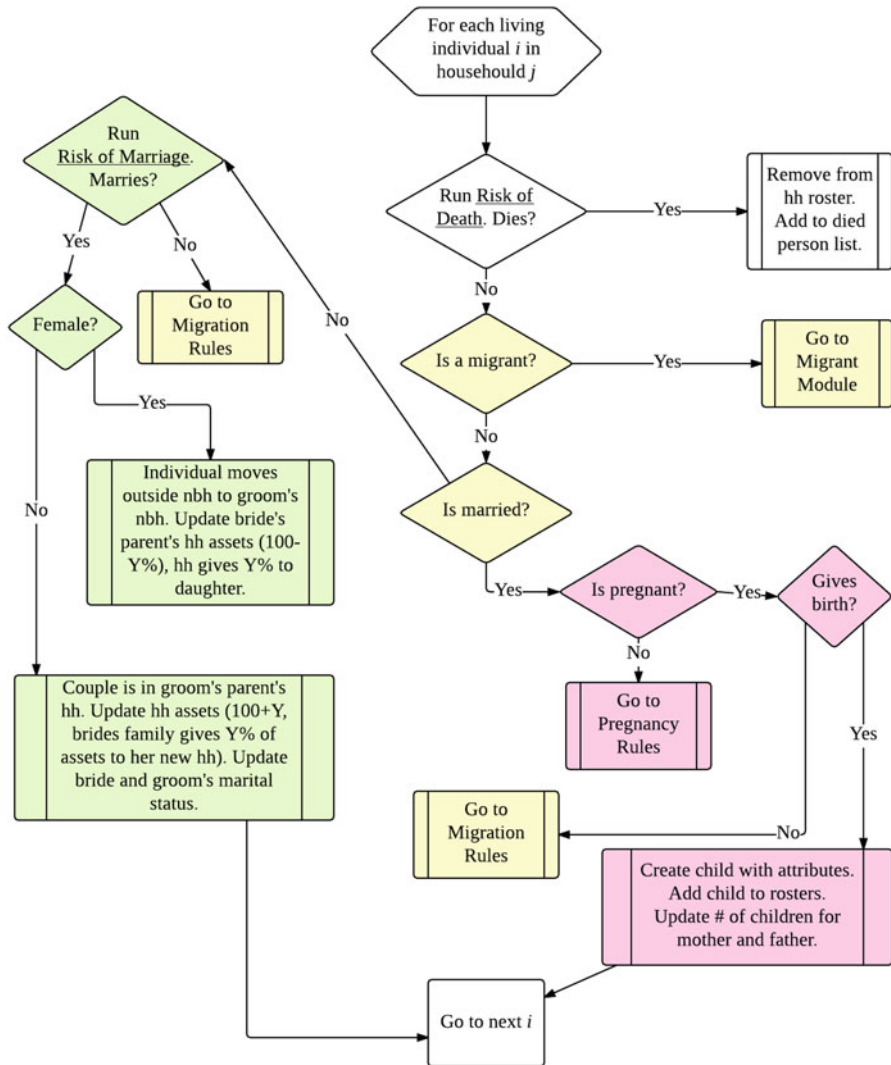


Fig. 6.5 Individual monthly module (Part I) in the ABM, for death, marriage, and birth behaviors

model.² If a woman marries, she “moves out” of the model, to a hypothetical spouse who is not amongst the agents in the model. In being removed from the model, she no longer contributes to the population. If a man marries, a spouse is created and enters the model, from a hypothetical natal family that is not amongst the agents in the model. This new spouse is assigned characteristics based on her new husband: she is 2 years younger, has the same migration experience and caste, and comes from the same economic background.

If an agent is married and female, she may get pregnant, and if pregnant, she will give birth in the ninth month of the pregnancy. Children born this month are added to the household roster and will run through the individual module during the next month.

If an individual does not get married or give birth, he or she can migrate. If the agent migrates, the model chooses a destination region from Nepal/India, other Asian countries, the Middle East, or Europe/North America, Australia, or New Zealand.

If an agent is a migrant this month, they cannot marry, and female agents cannot get pregnant. Migrants can experience the possibility of remitting money to their origin households and return migrating. A migrant who returns moves back into the origin household.

Other behaviors occur on an annual basis. At the beginning of each year, agents’ ages are updated. Each individual in the household who was enrolled in school last year either reaches their educational attainment (which is assigned at birth) or receives an additional year of education. If the agent reaches attainment, he or she experiences the possibility of working a salaried job for the year.

6.3.6 Household Module

At the end of each year, each household goes through the household module. Households can split into separate households, die (when all household members move out or die), and accrue assets. Sub-families, which are comprised of married sons and their spouses and children, can split from the primary household and create their own separate household within the same neighborhood. When they initially

²Exogamous marriage allows for a less computationally burdensome model and in this case, reflects the common marital patterns of the context. Here we consider marriage patterns for Nepal, wherein women most often leave their childhood homes to reside with their new husband’s natal family. In this context where women leave their villages at marriage, it is reasonable to program our model for exogamous marriage, where our female agents leave their model villages (and thus entirely leave the model). At the same time, new wives (female agents) enter the model for every male agent who gets married. It would be possible to program the model to allow endogamous marriage (where a female agent from one model village marries a male agent from another model village), but this would create an unwieldy model and increase computational time immensely. Further, endogamous and exogamous marriage patterns create almost exactly the same gender ratios in the model. When utilizing the simpler process of exogamous marriage, modelers should populate the life history of the new agent with characteristics as appropriate to the social context of the study area and requirements of the particular model.

split from a parent household, new households can purchase land and livestock. The amount of these items that they purchase is randomly selected from a distribution of the ownership of these items, with separate distributions for high- and low-caste groups. The cost of these items is subtracted from the new household's assets.

If the last remaining parent dies, the household is considered dead and household inheritance rules are applied. Land and assets are split among the sons. If there are no sons, the assets are split among the daughters.

Finally, household income from all sources is calculated for the year and added to the accrued assets. Sources of income include salaried work, remittances from migrants, and agricultural and livestock products. Costs of household sustenance, migration, and purchasing land and livestock are subtracted from income. We use a basic (additive) accounting system that is based on actual prices of goods in Chitwan, priced on site in 2013. This is more manageable than a regression based income system that allows assets to grow geometrically and in turn inordinately influence many other behavioral processes in the model. We then calculate the distribution of assets each year, and use a household's position in one of the three terciles of the distribution as the key household economic indicator in probabilistic decision-making equations. We do this to adjust for the fact that tracking assets over a long period of time is implausible, and to stress the importance of relative income to other households, instead of real asset accumulation, which may be inflated over long periods of time. We discuss the use of assets in more detail in the Challenges and Solutions section below.

6.3.7 *Neighborhood Module*

Neighborhoods are not static, although their data are aggregated from households and individuals. At the end of each year, neighborhoods are updated with the number and percent of migrants. The log of the number of migrants in a neighborhood is one factor in the outmigration decision.

6.3.8 *Interactions and Interdependence*

Interactions between agents are a key defining feature of ABMs. The ability for agents to interact differentiates ABMs from other types of micro-simulation models and is responsible for many of the unique results that ABMs are capable of finding. In fact, intra- and inter-agent interactions are likely the reason that we find very different results from our ABM compared to the regression-based predictions that we present in Figs. 6.7 and 6.8 and discuss further below.

There are multiple *inter-agent* interactions in our model, between individuals within the same households, as well as individuals within the same neighborhoods. Many of these interactions can be identified through the characteristics (or variables) in the equations that govern each behavior. For example, in the equation that

calculates the probability of men's migration, shown in Fig. 6.4, the variables for *married and living with spouse* and *married and living apart from spouse* create an interaction between the agent and his spouse: the residential location of a man's spouse affects his likelihood of migration. In addition, the number of children ever born to the man (which is entirely dependent on his wife's characteristics) affects his likelihood of migration. Another example of an influential characteristic that generates interactions between individuals is the indicator for household assets. The level of household assets affects a number of processes, including the probability of men's migration, as well as the destination of migrants, and the likelihood of giving birth and return migration. Because household assets are influenced by the work history (of both salaried jobs and migration) of every household member, this is one major way in which individuals influence the behaviors of others in the same household. In a similar way, the variable for the *natural log of the percentage of migrants in the neighborhood* tracks all people who are or were living in a particular neighborhood and records the percentage of them who are currently migrants living elsewhere. This creates interdependence between individuals from the same neighborhood. An individual who migrates out of a neighborhood contributes to the percent of migrants in that neighborhood, which in turn affects the probability of migration for other individuals in that neighborhood. The eligibility rules for some behaviors also create inter-agent interactions. For example, in order to be eligible for a salaried job, an individual's spouse cannot already be working a salaried job.

Another form of interaction in our ABM is *intra-individual*. In this case – what an individual does in one time period influences their likelihood of undertaking many other behaviors in subsequent time periods. As an example, look again at the men's migration equation, shown in Fig. 6.4. A man's marital status, number of children, work status, migration history, and household assets together have large influences on his likelihood of migration. In other words, his previous marriage, childbearing, work, and migration behaviors influence his future migration behavior. Intra-agent interactions are also present in the eligibility rules that govern many behaviors. For example, in order to get pregnant, a woman must be married, currently living with her spouse, not already be pregnant, and not have given birth in the last 4 months. The interactions and interdependencies in our ABM between individuals and their past behavioral history, their households, and their neighborhoods are numerous. We can mention only some of them here, but also refer the reader to the detailed model diagrams and behavioral rules online at www.bitly.com/NepalABM.

6.3.9 Armed Conflict in the ABM

One of the benefits of an ABM based on survey data is our ability to include exogenous shocks in the model, while continuing to use regression-based methods for decision making. Because the CVFS data were collected throughout the conflict between the Government of Nepal and the Maoists, we can include conflict events in regression analysis of demographic processes. Likewise, we simulate armed

conflict through event-based effects on the probability of marriage, pregnancy, out-migration, and return-migration. We employ an event-centered approach (Williams et al. 2012), using specific events—gun battles, bomb blasts, states of emergency, government instability, ceasefires, and strikes and protests, which influence the likelihood of the demographic behaviors each month. In this way, different conflict scenarios can be simulated by changing the number of each of these events each month.

6.4 Challenges and Solutions to ABM Design with Survey Data

As described above, we use empirical survey data to inform all aspects of our ABM, including initialization of the population of agents, parameterization of behavioral rules, and verification and validation. At each of these stages, there are multiple challenges in using survey data. In the remainder of this chapter, we describe some of these challenges and how we addressed them, with the aim to streamline the design and modeling processes for other demographic researchers.

6.4.1 *Initializing the Model*

The primary challenge of initializing an ABM with survey data is that there are almost no surveys that include every person in a population. In the CVFS case, the survey included only people between 15 and 59 years old who were resident in sampled neighborhoods. It thus excluded younger people, older people, and migrants who were temporarily away when the baseline survey was taken. These groups are commonly excluded from most surveys. A full population for ABM simulation requires that all these types of people be included in the initialization.

We used the CVFS survey data to create these missing agents. The CVFS, as with many other surveys, included a complete household roster, providing information on all people living in the household, regardless of age or current residence. In addition, surveyed adults were asked how many children they had and their sex and age. Using the parent report, we created agents for all children ever born up to the age of 15 and placed them in the parents' household and assumed that older children would have moved away by that age.³ The CVFS household roster information also allowed

³Household rosters also allow for the creation of older agents. However, for some models of demographic processes it will not be necessary to create older people if they cannot undertake birth, marriage, or migration. The necessity of creating older people entirely depends on the behaviors they can undertake and to what extent their presence influences the behaviors of other agents.

us to create migrants who were living away at the beginning of the model. They were accorded age and sex based on the household roster and other characteristics to match those of household members.

6.4.2 *Parameterizing the Model*

Probabilistic decisions in the model, such as the outmigration example we provided, can be based on analyses of survey data. The process is fairly simple, where regression equations on survey data are used to predict a certain behavior in the ABM. However, regression techniques are designed to examine the influence of factors on a behavior and are not meant to predict that behavior. As such, there are certain limitations and challenges to using them in generating ABM rules.

One challenge is in determining which variables to use in a regression equation. First, we identified parameters that would allow an individual to interact with her changing status (such as being married or having ever migrated) as well as her household and neighborhood environment. Second, we used variables that were statistically significant to at least the $p < 0.10$ level. Third, we used variables that produced theoretically sound results. For example, we would expect increased assets to decrease the likelihood of migration.

A second challenge is to determine when regression equations are appropriate. Ordinary least squares equations can result in geometric growth or decrease in an outcome overtime. When simulated over long periods, this type of growth can create unrealistic and drastic outcomes on an entire model and simulated results. For example, after many trials with a regression based equation to determine income in our model, we found that this resulted in almost exponentially increasing assets, which unduly influenced almost all other parts of the model. An additive model, where income was added and a series of expenses were subtracted annually, created a much more realistic, and stable, change in household assets over time. However, even the additive model of assets created dramatic increases in wealth, beyond what we believe is reasonable. Because we are not concerned with the amount of assets, but rather the effect of assets on decision-making, such as migration, we changed our regression equations for behaviors to rely on the relative distribution of assets per year to inform decisions. This is accomplished by using terms for the top, middle, and lowest third of the asset distribution, instead of absolute assets, to predict migration, return migration, and pregnancy. With this focus on the relative household income versus other households in Chitwan, we were able to model the relationship between assets and demographic processes consistently over time.

6.4.3 *Adjusting the Model to Create an Experimental Situation*

As discussed above, the ideal ABM for studying the influence of macro-level change over time would create an experiment-like situation where key processes are generally stable. Because demographic processes, such as population growth, number of households, and age and sex structure, likely affect almost any behavior under study, it is necessary to create generally stable population processes. Unstable population processes (such as extremely rapid population growth or decline in household size) could create large changes in any behavior under study. If this is not recognized, then the change in one behavior could be attributed to the macro process of interest, instead of the unstable process that actually created it. In short, the necessity of creating generally stable population processes cannot be understated.

Population stabilization can be easily assessed with population pyramids, household size, number of households, and birth, death, and migration rates. The processes we examined, and found to be useful, are listed in Fig. 6.6. Based on Stable Population Theory (Preston et al. 2001), we can expect population pyramids to stabilize in shape after about 80–100 years if a population is generally stable. If population pyramids do not stabilize after this time, then the population processes are likely not stable and must be adjusted to create an analytically useful ABM. For other processes, such as household size and birth, death, and migration rates, we also recommend examining at least 100 years of simulated data. Although analysis of the final model might not extend past one generation (about 30 years), irregularities in population processes can geometrically increase or decrease starting small but becoming much bigger after several decades. Thus much longer simulations are required to make some problems visible.

In the case of our model, preliminary simulation results without conflict or other disturbances showed that the population grew rapidly and household sizes grew from an average of five to an average of 11 individuals per household. This could have been overlooked as an interesting (if possibly theoretically significant) finding. However, our verification process allowed us to pinpoint irregularities such as: significantly decreasing marriage and fertility rates and increasing death rates, many fewer boys were born than girls, households were not splitting properly when a married couple moved out to form their own household, and assets were increasing exponentially with a consequent exponential decrease in migration.

Some of these problems were mistakes in the model programming and were thus easily fixed. Other problems, such as decreasing marriage rates, were based in the regression equations derived from empirical survey data. The root of the problem appears not to be biased survey data, but the fact that the data were collected during a period when marriage and fertility rates were changing dramatically. Such changes in demographic behaviors are common, but are most often period effects and are seldom sustained for long periods of time. To address these problems, we adjusted the constants in the marriage and fertility equations and the age coefficient in the death equation. All adjustments were within the 95 % confidence intervals for the constants and coefficients and are thus statistically appropriate.

Cumulative composition and rates

Population composition

- Population size
- Population age-sex structure (population pyramids)
- Number of households
- Mean household size
- Median household size
- Deaths per capita

Fertility and marriage

- Proportion of women/men ages 25-40 who are married
- Proportion of married women over 40 who have ever had a birth
- Mean number of children ever born per married woman over 40
- Mean age at first birth for women over 40

Migration

- Proportion of women/men over age 16 who have ever migrated
- Mean number of times migrated for female/male migrants over age 16
- Number of months away during last migration spell for female/male migrants

Education and salaried work

- Mean years of education

Income and assets

- Mean and median assets per household
- Median assets per household in the bottom, middle, and top terciles of asset distribution
- Mean poultry, livestock, and land owned per household

Annual Rates

Fertility and marriage

- Proportion of unmarried men over age 14 who got married this year
- Proportion of unmarried women over age 14 who got married this year
- Proportion of married women ages 14-45 who had a birth this year
- Proportion of women over age 14 with first child born this year
- Proportion of women over age 14 with second child born this year

Migration

- Proportion of households with at least one migrant out this year
- Mean number of migrants out this year per household with 1+ migrants out
- Proportion of eligible men who migrated this year
- Proportion of eligible women who migrated this year

Education and salaried work

- Proportion of population over age 16 with salaried job this year

Income and assets

- Mean and median income per household

Fig. 6.6 Demographic processes that can be used for assessment and analysis of ABMs

6.4.4 Burning-in the Model to Allow for New Selection Processes

In seeking to create stable population processes, burning-in the model for an appropriate amount of time is also necessary. This is particularly the case when survey data are used to initially populate the model, but simulated agents replace them over time. The issue here is that no matter how well the data based behavioral rules (such as our regression equations for marriage, pregnancy, and migration) are specified, they are each models. And models are always simplified versions of reality and subject to unobserved heterogeneity in numerous ways. For example, it is likely that beauty is a key factor in the likelihood of marriage in Nepal, but our marriage regression equations do not and cannot take this into account. Consequently, the selection processes programmed in the model are different from the selection processes that affected the behaviors of the surveyed population. The result of these different selection processes is initial instability in model results. This happens when the model selection processes replace the real population selection processes and simulated agents (those born during the model) replace agents who are based on real survey data.

Fortunately, there is a simple process to address this concern: model burn-in. Burning-in is defined as allowing the ABM to run for several time steps, to allow the simulated processes and simulated agents to populate the model as the survey-based agents who were subject to real behavioral processes and selection age out of the model population or age out of a behavioral process. A model should be entirely burned-in before experimental scenarios are enacted. In some cases, the length of burn-in is evident, when initially wildly unstable results stabilize within a few years. In other cases, as with the marriage rates calculated by our ABM (as shown in Fig. 6.7 and discussed more below), results are not wildly unstable in the first years of the simulation. Instead, rates for men and women increase steadily for about 14 years, after which they level off, until our conflict scenarios begin in year 17. Notably, 14 years is the age at which simulated agents in our model are eligible for marriage. In other words, marriage rates stabilize after 14 years because this is when the entirely simulated population (those “born” during the model) and their simulated marriage processes overwhelm the survey-based population and their differently selective marriage processes from before the model simulation began. Thus, we defined our burn-in period as the length of time after which entirely simulated agents are eligible to undertake all key behaviors. This happens at year 16 of the simulation, as it is age 16 at which all agents are eligible to marry, migrate, and give birth. We begin the simulation of conflict scenarios just after, at year 17 of the simulation.

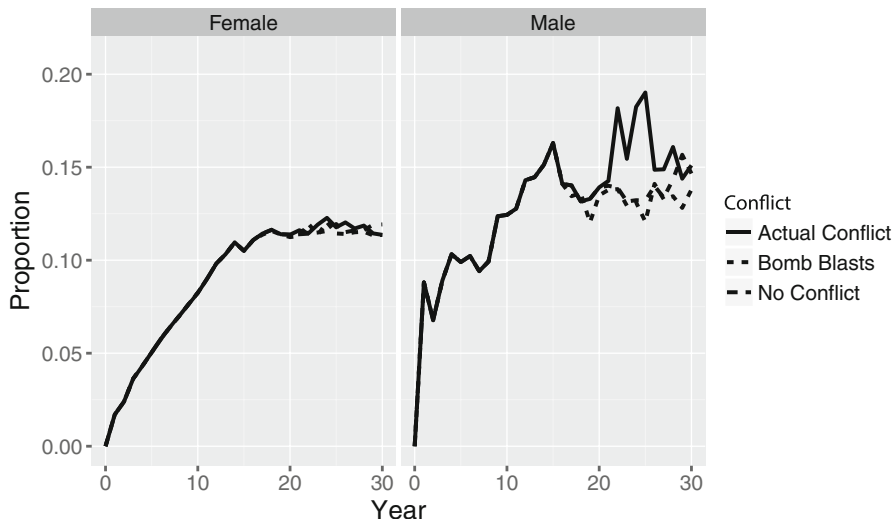


Fig. 6.7 ABM-based results showing male and female marriage rates (proportion of unmarried who married this year) in three different conflict scenarios for 30 years of simulation (Note: The conflict simulation is years 17–20 of the simulation)

6.5 Implications of ABM Compared to Regression Based Analysis

As discussed earlier in this chapter, ABMs are relatively new to the social sciences, particularly when compared to the more common and widely understood tool of regression analysis. Although we have identified several applications for which ABMs are more useful or appropriate than regression based analysis, it is often difficult to fully understand these differences without an explicit example. This is the purpose of this section. We show results from our ABM, of men’s and women’s marriage rates, in the scenarios of no conflict, the conflict as it actually happened in Nepal, and a hypothetical conflict scenario. These scenarios are the virtual experiments which we mention above. In the scenario with no conflict, we did not simulate any conflict events. In the scenario with the conflict as it actually happened, we simulate the actual number of each type of conflict related event that happened each month. The events include gun battles, bomb blasts, strikes and protests, government instability and states of emergency. In the scenario with a hypothetical conflict, we simulate a conflict in which there was one bomb blast per month for the duration of the conflict period. For each of these scenarios we allow a 17-year period before the conflict (in order for the model to stabilize, or “burn-in” as described above in 6.4.4), followed by a 48-month conflict. After the conflict ends, we continue to run the model simulation for a total of 30 years.

We also show results from a calculation of predicted men’s and women’s probability of marriage for each of these scenarios, based only on the regression equations for men’s and women’s marriage. These calculations were undertaken

by using the characteristics (values for each variable in the regression equation) of an “average” 25 year old male and female in the population—in other words, the mean population value for each variable. These values were input into the regression equation, along with the monthly number of conflict-related events in each scenario, to calculate a monthly predicted probability of marriage. The monthly probabilities were then converted into annual probabilities, which we present here. The comparison of these results provides a more clear understanding of the differences between these two analytical tools (ABM and regression) for prediction of population processes.

Figure 6.7 shows the results from our ABM simulated annual marriage rates for men and women. Specifically, the graphs show the proportion of unmarried men and women over the age of 14 who got married during each year of the simulations. As you can see, marriage rates start very low, then progressively increase throughout all scenarios, for both men and women. Further, marriage rates change each year throughout the 30 years. Amongst the most notable pattern in these graphs is that marriage rates for men in the actual conflict scenario are much higher than in the no conflict and bomb blast scenarios. Alternately, bomb blasts produce a similar men’s marriage rates to the no conflict scenario. A second particularly notable result is that men’s marriage rates after the conflict period remain quite different in the actual conflict scenario until about 10 years after the conflict. In other words, we find long-term effects of the conflict on men’s marriage rates, well after the conflict ends. We do not find long-term effects for women’s marriage rates.

Figure 6.8 shows the results from our regression based annual predicted probabilities of marriage. Note that these are probabilities instead of rates. This is of course




Fig. 6.8 Annual predicted probabilities of marriage for men and women, based on regression equation alone

the first major difference between regression and ABMs. ABMs are designed to simulate population processes and results can easily be calculated as rates (such as the number of people who got married in a year divided by the number of people who were eligible to get married at the beginning of that year). Regression on the other hand is designed to calculate the effect of any particular characteristic (or variable) on an outcome and is not suited to calculating rates within a population. Indeed, predicted probabilities, which are shown here, are the closest possible calculation to a rate that can be achieved with a regression equation. The only way that regression equations can reasonably yield a rate calculation at population level is through including the equations in an ABM.

A second notable outcome of this comparison is the stability of the regression-based probabilities. As Fig. 6.8 shows, the probability of marriage is exactly the same in every year of the no conflict scenario, and is stable throughout the conflict (as well as before and after the conflict) in the bomb blast scenario. This is because the “average” hypothetical person who is used to calculate these probabilities is the same each year of the calculation. The actual conflict scenario yields varying probabilities during the conflict, entirely due to the changing number of conflict-related events each month of the actual conflict. Note however that the probability of marriage in the actual conflict scenario returns to the exact same level as that of the no conflict scenario after the conflict ceases. In contrast, remember that the marriage rates in the ABM-based results (in Fig. 6.7) constantly changed, before, during, and after the conflict simulations. This is because the ABM simulates a population that constantly changes and interacts. In fact, the ABM platform allowed us to find long-term effects of the conflict on men’s marriage rates, while such a result is clearly not possible with the regression-based calculations.

A third notable result of our comparison here is that the regression based probability of marriage for women is lower in the bomb blast scenario and higher in the actual conflict scenario, compared to the no conflict scenario. Alternately, with the ABM we find generally similar rates of women’s marriage in all three scenarios, with slightly higher rates in the bomb blast scenario compared to the no conflict scenario. In the case of men’s marriage, the ABM and regression predict similar comparative differences between the actual and no conflict and bomb blasts scenarios, with actual conflict producing much higher rates of marriage in the ABM, just as it does in the regression.

6.6 Conclusion

This chapter focuses on a particular type of agent-based model, one that leverages survey data to initialize a population and operationalize behavioral rules. Although this type of model is reasonably common in the geographic and land use sciences

(An et al. 2001, 2014; Manson 2005; Parker et al. 2008; An 2012; Zvoleff and An 2014), it is rarely used in demography. This situation is unfortunate, given the many possibilities for this methodology to contribute to demography and particularly to our understanding of how macro-level events, such as armed conflict, natural disaster, climate change, economic crises, and policy changes, influence population growth and change.

While the ability to address outstanding questions in these areas is an immense benefit to survey data based ABMs, there is one key limitation. As they use a case study approach, this type of ABM is comparatively weaker in providing broad insights that are relevant regardless of geographic setting. Broader conclusions can only be developed when multiple studies in different areas reach similar conclusions. For example, if a case study of armed conflict and population change in Nepal finds similar outcomes to studies of the same subject in Colombia, Sudan, and Afghanistan, then we can begin to develop broader conclusions about the nature of armed conflict and population change. As demography and other social science disciplines are moving more towards the case study approach, as compared to cross-country models, regardless of the methodological tools used, this situation with ABMs is not unusual.

In this context, our broad aim is to encourage the use of survey data based ABMs in the demographic sciences. In addition to not being widely known in demography, this type of model is extremely time consuming to design, test, and analyze. In this chapter, we presented the design of our ABM, which uses detailed survey data to simulate population dynamics during armed conflict in the Chitwan Valley of Nepal. The combination of this chapter and more detailed description of our model online at www.bitly.com/NepalABM should provide key guidance for the development of future survey data based ABMs. To further streamline the long and difficult design and testing process, we also discuss several challenges we faced and how they can be addressed. The primary method we used to find and address problems in our model is careful examination of simulated outcomes of several demographic processes, as listed in Fig. 6.6. In this way, not only can ABMs contribute to demography, but demography can also contribute to ABM methodology.

One of the key points we hope to instill here is that survey data based ABMs should not be thought of as attempts to model reality. They are simply models, just as other ABMs and statistical procedures are models that are not replications of reality but can nonetheless be useful. We argue that the most useful way to design survey data based ABMs is as experimental laboratories and we describe methods for doing so. Because we can almost never ethically experiment on real human populations (and certainly not with armed conflicts!) experimental designs using survey data based ABMs have immense promise for contributing to the demographic sciences. We hope to find more of this type of model in the literature in coming years.

Acknowledgements This research was supported by generous grants from the National Institute of Child Health and Human Development including a research grant (HD067587) and a research infrastructure grant (R24HD042828) to the Center for Studies in Demography and Ecology. We thank Barbara Entwisle, William Axinn, and Dirgha Ghimire for their assistance, the Institute for

Social and Environmental Research – Nepal for collecting the survey and violent event record data, Meeta Sainju-Pradhan for collecting the political event data, the South Asia Terrorism Portal and Informal Service Sector for collecting records of violent events, and the CVFS respondents for sharing their lives and experiences. All errors or omissions remain the responsibility of the authors.

References

- Adler, E. S., & Clark, R. (2008). Experimental research. In E. Adler & R. Clark (Eds.), *How it's done: An invitation to social research* (pp. 185–214). Belmont: Thomson Wadsworth.
- Agadjanian, V., & Prata, N. (2002). War, peace, and fertility in Angola. *Demography*, 39(2), 215–231.
- An, L. (2012). Modeling human decisions in coupled human and natural systems: Review of agent-based models. *Ecological Modelling*, 229, 25–36.
- An, L., & Liu, J. (2010). Long-term effects of family planning and other determinants of fertility on population and environment: Agent-based modeling evidence from Wolong Nature Reserve, China. *Population and Environment*, 31(6), 427–459.
- An, L., Liu, J., Ouyang, Z., Linderman, M., Zhou, S., & Zhang, H. (2001). Simulating demographic and socioeconomic processes on household level and implications for giant panda habitats. *Ecological Modelling*, 140(1–2), 31–49.
- An, L., Linderman, M., Qi, J., Shortridge, A., & Liu, J. (2005). Exploring complexity in a human-environment system: An agent-based spatial model for multidisciplinary and multiscale integration. *Annals of the Association of American Geographers*, 95(1), 54–79.
- An, L., Zvoleff, A., Liu, J., & Axinn, W. G. (2014). Agent-based modeling in coupled human and natural systems (CHANS): Lessons from a comparative analysis. *Annals of Association of American Geographers*, 104(4), 723–745.
- Aparacio Diaz, B., & Fent, T. (2006). An agent-based simulation model of age-at-marriage norms. In F. C. Billari, T. Fent, A. Prskawetz, & J. Scheffran (Eds.), *Agent-based computational modelling: Applications in demography, social, economic and environmental sciences* (pp. 85–116). Heidelberg: Physica-Verlag.
- Aparacio Diaz, B., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography*, 48(2), 559–579.
- Apodaca, C. (1998). Human rights abuses: Precursor to refugee flight? *Journal of Refugee Studies*, 11(1), 80–93.
- Axinn, W. G., & Yabiku, S. T. (2001). Social change, the social organization of families, and fertility limitation. *American Journal of Sociology*, 106(5), 1219–1261.
- Axinn, W. G., Barber, J. S., & Ghimire, D. J. (1997). The neighborhood history calendar: A data collection method designed for dynamic multilevel modeling. *Sociological Methodology*, 27(1), 355–392.
- Axinn, W. G., Pearce, L. D., & Ghimire, D. (1999). Innovations in life history calendar applications. *Social Science Research*, 28(3), 243–264.
- Barber, J. S., Shivakoti, G. P., Axinn, W. G., & Gajurel, K. (1997). Sampling strategies for rural settings: A detailed example from Chitwan Valley Family Study, Nepal. *Nepal Population Journal*, 6(5), 193–203.
- Billari, F. C., Prskawetz, A., Aparacio Diaz, B., & Fent, T. (2007). The “Wedding-Ring”: An agent-based marriage model based on social interaction. *Demographic Research*, 17(3), 59–82.
- Bruch, E. E., & Mare, R. D. (2006). Neighborhood choice and neighborhood change. *American Journal of Sociology*, 112(3), 667–709.
- Czaika, M., & Kis-Katos, K. (2009). Civil conflict and displacement: Village-level determinants of forced migration in Aceh. *Journal of Peace Research*, 46(3), 399–418.
- Davenport, C., Moore, W., & Poe, S. (2003). Sometimes you just have to leave: Domestic threats and forced migration, 1964–1989. *International Interactions*, 29(1), 27–55.

- Eloundou-Enyegue, P. M., Stokes, C. S., & Cornwell, G. T. (2000). Are there crisis-led fertility declines? Evidence from central Cameroon. *Population Research and Policy Review*, 19(1), 47–72.
- Ewert, U. C., Roehl, M., & Uhrmacher, A. M. (2003). Consequences of mortality crises in pre-modern European towns: A multiagent-based simulation approach. In F. Billari & A. Prskawetz (Eds.), *Agent-based computational demography. Using simulation to improve our understanding of demographic behavior* (pp. 175–196). Heidelberg: Springer Physica-Verlag.
- Ghimire, D. J., & Axinn, W. G. (2010). Social organization, land use and the hazard of first birth. *Rural Sociology*, 75, 478–513.
- Gibney, M., Apodaca, C., & McCann, J. (1996). Refugee flows, the internally displaced and political violence (1908–1993): An exploratory analysis. In A. Schmid (Ed.), *Whither refugee? The refugee crisis: Problems and solutions*. Leiden: Ploom.
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S. K., Huse, G., Huth, A., Jepsen, J. U., Jørgensen, C., Mooij, W. M., Müller, B., Pe'er, G., Piou, C., Railsback, S. F., Robbins, A. M., Robbins, M. M., Rossmannith, E., Rüger, N., Strand, E., Souissi, S., Stillman, R. A., Vabø, R., Visser, U., & DeAngelis, D. L. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2), 115–126.
- Grimm, V., Berger, U., DeAngelis, D. L., Gary Polhill, J., Giske, J., & Railsback, S. F. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, 221(23), 2760–2768.
- Heiland, F. (2003). The collapse of the Berlin Wall: Simulating state-level East to West German migration patterns. In C. F. Billari & A. Prskawetz (Eds.), *Agent-based computational demography. Using simulation to improve our understanding of demographic behavior* (pp. 73–96). Heidelberg: Springer Physica-Verlag.
- Heuveline, P., & Poch, B. (2007). The phoenix population: Demographic crisis and rebound in Cambodia. *Demography*, 44(2), 405–426.
- Hutt, M. (2004). Monarchy, democracy and Maoism in Nepal. In M. Hutt (Ed.), *Himalayan people's war: Nepal's Maoist revolution* (pp. 1–20). Bloomington: Indiana University Press.
- Jackson, J., Forest, B., & Sengupta, R. (2008). Agent-based simulation of urban residential dynamics and land rent change in a gentrifying area of Boston. *Transactions in GIS*, 12(4), 475–491.
- Jayaraman, A., Gebreselassie, T., & Chandrasekhar, S. (2009). Effect of conflict on age at marriage and age at first birth in Rwanda. *Population Research and Policy Review*, 28, 551–567.
- Kniveton, D., Smith, C., & Wood, S. (2011). Agent-based model simulations of future changes in migration flows for Burkina Faso. *Global Environmental Change*, 21(1), S34–S40.
- Kollmair, M., Manandhar, S., Subedi, B., & Thieme, S. (2006). New figures for old stories: Migration and remittances in Nepal. *Migration Letters*, 3(2), 151–160.
- Lindstrom, D. P., & Berhanu, B. (1999). The impact of war, famine, and economic decline on marital fertility in Ethiopia. *Demography*, 36(2), 247–261.
- Manson, S. M. (2005). Agent-based modeling and genetic programming for modeling land change in the Southern Yucatán Peninsular Region of Mexico. *Agriculture Ecosystems and Environment*, 111(1–4), 47–62.
- Massey, D. S., Axinn, W. G., & Ghimire, D. J. (2010). Environmental change and out-migration: Evidence from Nepal. *Population & Environment*, 32, 109–136.
- Melander, E., & Oberg, M. (2006). Time to go? Duration dependence in forced migration. *International Interactions*, 32, 129–152.
- Mena, C. F., Walsh, S. J., Frizzelle, B. G., Xiaozheng, Y., & Malanson, G. P. (2011). Land use change on household farms in the Ecuadorian Amazon: Design and implementation of an agent-based model. *Applied Geography*, 31(1), 210–222.
- Moore, W., & Shellman, S. (2004). Fear of persecution: Forced migration, 1952–1995. *Journal of Conflict Resolution*, 40(5), 723–745.

- Parker, D. C., Entwisle, B., Rindfuss, R. R., VanWey, L. K., Manson, S. M., Moran, E., An, L., Deadman, P., Evans, T. P., Marc, L., Mohammad Mussavi Rizvi, S., & Malanson, G. (2008). Case studies, cross-site comparisons, and the challenge of generalization: Comparing agent-based models of land-use change in frontier regions. *Journal of Land Use Science*, 3(1), 41–72.
- Pettigrew, J. (2004). Living between the Maoists and the army in rural Nepal. In M. Hutt (Ed.), *Himalayan people's war: Nepal's Maoist revolution* (pp. 261–284). Bloomington: Indiana University Press.
- Preston, S. H., Heuveline, P., & Guillot, M. (2001). *Demography: Measuring and modeling population processes*. Oxford: Blackwell Publishers.
- Schmeidl, S. (1997). Exploring the causes of forced migration: A pooled time-series analysis, 1971–1990. *Social Science Quarterly*, 78(2), 284–308.
- Shemyakina, O. (2009). *The marriage market and Tajik armed conflict*. HiCN Working Paper 66. Accessible at <http://www.hicn.org/wordpress/wp-content/uploads/2012/06/wp66.pdf>.
- South Asia Terrorism Portal. (2006). *Major incidents of terrorist violence in Nepal, 1999–2006*. New Delhi: South Asia Terrorism Portal.
- Stanley, W. D. (1987). Economic migrants or refugees? A time-series analysis of Salvadoran migration to the United States. *Latin American Research Review*, 22(1), 132–154.
- Thieme, S., & Wyss, S. (2005). Migration patterns and remittance transfer in Nepal: A case study of Sainik Basti in Western Nepal. *International Migration*, 43(5), 59–98.
- United Nations, Department of Economic and Social Affairs. (2013). *Trends in international migrant stock: Migrants by destination and origin*. United Nations database, POP/DB/MIG/Stock/Rev.2013.
- Walsh, S. J., Malanson, G. P., Entwisle, B., Rindfuss, R. R., Mucha, P. J., Heumann, B. W., McDaniel, P. M., Frizzelle, B. G., Verdery, A. M., Williams, N. E., Yao, X., & Deng, D. (2013). Design of an agent-based model to examine population-environment interactions in Nang Rong District, Thailand. *Applied Geography*, 39, 183–198.
- Weiner, M. (1996). Bad neighbors, bad neighborhoods: An inquiry into the causes of refugee flows. *International Security*, 21(1), 5–42.
- Williams, N. E. (2009). Education, gender, and migration in the context of social change. *Social Science Research*, 38(4), 883–896.
- Williams, N. E. (2013). How community organizations moderate the effect of armed conflict on migration in Nepal. *Population Studies*, 67(3), 353–369.
- Williams, N. E. (2015). Mixed and complex mixed migration during armed conflict: Multi-dimensional empirical evidence from Nepal. *International Journal of Sociology*, 45(1), 44–63.
- Williams, N. E., Ghimire, D. J., Axinn, W. G., Jennings, E. A., & Pradhan, M. S. (2012). A micro-level event-centered approach to investigating armed conflict and population responses. *Demography*, 49(4), 1521–1546.
- Winter, J. M. (1992). War, family, and fertility in twentieth century Europe. In J. R. Gillis, L. A. Tilly, & D. Levine (Eds.), *The European experience of declining fertility, 1850–1970: The quiet revolution* (pp. 291–309). Cambridge, MA: Blackwell.
- Žamac, J., Hallberg, D., & Lindh, T. (2009). Low fertility and long-run growth in an economy with a large public sector. *European Journal of Population*, 26(2), 183–205.
- Zolberg, A. R., Suhrke, A., & Aguayo, S. (1989). *Escape from violence: Conflict and the refugee crisis in the developing world*. New York: Oxford University Press.
- Zvoleff, A., & An, L. (2014). Analyzing human-landscape interactions: Tools that integrate. *Environmental Management*, 53(1), 94–111.

Chapter 7

Regression Metamodels for Sensitivity Analysis in Agent-Based Computational Demography

André Grow

7.1 Introduction

Over the last years, an increasing number of scholars advocate the use of agent-based computational (ABC) modelling for the study of demographic phenomena (e.g. Billari et al. 2003, 2007; Silverman et al. 2013; Todd et al. 2013). One reason for this is the recognition that population-level outcomes can often not be reduced to a simple aggregate of individual decisions. Instead, human populations are complex systems in which individuals' demographic choices are constrained by the social environment and feed back into this environment (Smaldino and Schank 2011). ABC modelling makes it possible to explicate such feedback mechanisms and enables us to study their implications by means of computational simulation (Bonabeau 2002; Epstein 1999; Macy and Flache 2009; Macy and Willer 2002).

While it is true that agent-based models can greatly facilitate the study of social complexity, it is also true that simulation models themselves can be complex and this can make it “[difficult] to know which relationships and processes are driving model behavior” (Coutts and Yokomizo 2014, p. 7). *Sensitivity analysis* is an important tool for dealing with this problem. In sensitivity analysis, we seek to understand how one or more model parameters affect model outputs through simulation experiments

Electronic supplementary material The online version of this chapter (doi: [10.1007/978-3-319-32283-4_7](https://doi.org/10.1007/978-3-319-32283-4_7)) contains supplementary material, which is available to authorized users.

A. Grow (✉)

Centre for Sociological Research, University of Leuven (KU Leuven), Parkstraat 45, 3000 Leuven, Belgium

e-mail: andre.grow@soc.kuleuven.be

© Springer International Publishing Switzerland 2017

A. Grow, J. Van Bavel (eds.), *Agent-Based Modelling in Population Studies*, The Springer Series on Demographic Methods and Population Analysis 41, DOI 10.1007/978-3-319-32283-4_7

185

(Kleijnen 2008; Law and Kelton 2000). A simple approach to sensitivity analysis is to study model outputs over an exhaustive set of parameter combinations. This is feasible for small simulation models with few parameters, but the larger the number of parameters, the less feasible it becomes. To illustrate this, consider a simulation model with three parameters that each has five different levels. In this case, there are $5^3 = 125$ possible parameter combinations. However, if there are ten parameters, the number of possible combinations increases to $5^{10} = 9,765,625$. Thus, depending on the computational expensiveness of the model (i.e. the time it takes to conduct one simulation run), an increase in the number of parameters can quickly make this approach infeasible. Furthermore, the large amount of data it tends to create can be very difficult to interpret (Coumts and Yokomizo 2014; McCarthy et al. 1995).

A more efficient approach to sensitivity analysis of complex simulation models is the use of *metamodels*. Metamodels, also called emulators or surrogate models (Kleijnen 2005, 2008), are simplified statistical representations of simulation models that aim at quantifying the relation between model parameters and model outputs. One major advantage of metamodels is that they can drastically reduce the computational effort that is needed to gain comprehensive insights into the behaviour of a simulation model. Despite this advantage, metamodels have so far largely been neglected in agent-based computational demography. An exception from this is a recent series of papers by Bijak and colleagues (e.g. Bijak et al. 2013; Silverman et al. 2013; see also the chapter by Hilton and Bijak in this volume), who illustrated the use of metamodels by applying Gaussian process emulators to their re-implementation of Billari et al.'s (2007) model of entry into first marriage. Similarly, De Mulder et al. (2015) recently illustrated how Gaussian process emulators can be used for calibrating demographic agent-based models.

Gaussian process emulators are just one of a number of statistical approaches to metamodeling that all have their specific advantages and disadvantages (for an overview of different approaches see Simpson et al. 2001). The purpose of this chapter is to introduce the reader to *regression metamodels* and to illustrate how they can be used for sensitivity analysis of complex simulation models. Regression metamodels have a long and successful track record in simulation research (Kleijnen 2005) and I argue that this type of metamodel is particularly attractive for sensitivity analysis of agent-based models in demographic research. The reason is that most demographers have at least a basic understanding of regression analysis and this makes regression metamodels highly accessible and easy to communicate.

In what follows, I first present regression metamodels and discuss experimental designs that can be used to collect the data that is necessary for estimating such models. Subsequently, I illustrate the use of regression metamodels by applying this method to Grow and Van Bavel's (2015) model of educational assortative mating in the context of Belgian marriage markets. I close the chapter with a discussion of the benefits and limitations of regression metamodels and point the reader to additional topics in the literature on metamodeling. Throughout the chapter, I assume that the reader has a basic understanding of agent-based computational modelling and ordinary least squares regression analysis. I therefore restrict my mathematical expositions to those aspects of regression metamodels that deviate from standard multiple regression.

7.2 Regression Metamodels

A metamodel treats a simulation model as a *black box* and only describes the observed relations between simulation parameters and outputs, without any reference to the inner workings of the simulation model (Kleijnen and Sargent 2000; Kleijnen et al. 2005). A black box representation of a simulation model can be given by

$$z = f(x_1, \dots, x_k, r), \quad (7.1)$$

where z are the observed simulation outputs, x_i with $i = 1, \dots, k$ refers to a set of k parameters of the simulation model, r are the pseudo random number seeds in the different simulation runs, and $f(\cdot)$ represents the mathematical function that is implicitly defined in the simulation model and connects model parameters to model outputs (Kleijnen 2008). The goal is to find a statistical function that approximates $f(\cdot)$ well and therefore can be used as a surrogate of $f(\cdot)$. Once we have found such a function, we can use it to answer questions such as: which parameters affect model outcome z ? Does a change in x_1 lead to an increase or a decrease in z ? Does the effect of x_1 depend on the value of x_2 ? We can also use it to predict z for hitherto unobserved parameter combinations without needing to actually run the simulation model for these combinations.

Finding a statistical function that approximates $f(\cdot)$ well always requires experimentation with the simulation model (Kleijnen 2005, 2008). That is, we always need to run the simulation model several times, while systematically varying the values of its parameters between the different runs. Yet, different types of metamodels have different data requirements and the selection of the type of metamodel therefore guides the data collection effort (Kleijnen 2005). Hence, in this section, I first discuss the statistical approach that is used to estimate regression metamodels and discuss how we can assess whether a given regression model approximates $f(\cdot)$ sufficiently well. Subsequently, I discuss experimental designs that are suitable for estimating regression metamodels.

7.2.1 Statistical Approach

In regression metamodels, the function that is used to approximate $f(\cdot)$ is a polynomial, typically of the first or second order (Kleijnen 2008). A standard first-order polynomial (i.e. a simple additive model with linear effects) is given by

$$z = \beta_0 + \sum_{i=1}^k \beta_i x_i + \varepsilon, \quad (7.2)$$

where β_0 is the intercept, β_i is the effect of model parameter x_i , and ε is the approximation error (i.e. residual). Equation 7.2 focuses only on the main effect

of each parameter and assumes that the k parameters do not interact with each other in affecting z . Furthermore, it assumes that the effects of all parameters are linear. If we expect that some of the k model parameters interact with each other, we can augment Eq. 7.2 with multiplicative terms, so that the polynomial takes the form

$$z = \beta_0 + \sum_{i=1}^k \beta_i x_i + \sum_{i=1}^{k-1} \sum_{j=i+1}^k \beta_{ij} x_i x_j + \varepsilon, \quad (7.3)$$

where β_{ij} is the estimated effect of the interaction between x_i and x_j . If we additionally expect that the relation between some parameters and z is subject to curvature, we can estimate a full second-order model of the form

$$z = \beta_0 + \sum_{i=1}^k \beta_i x_i + \sum_{i=1}^{k-1} \sum_{j=i+1}^k \beta_{ij} x_i x_j + \sum_{i=1}^k \beta_{ii} x_i^2 + \varepsilon, \quad (7.4)$$

where β_{ii} is the estimated effect of the quadratic term of parameter x_i .


The functions defined in Eqs. 7.2, 7.3, and 7.4 are identical to multiple regression models and can therefore be estimated with the ordinary least squares method (Kleijnen 2008). Polynomials of an even higher order have been used in some applications, but interpreting the estimated regression coefficients can be difficult and the estimated effects of higher-order terms are often small (Kleijnen 2005). I therefore limit my discussion to first- and second-order polynomials as defined in Eqs. 7.2, 7.3, and 7.4.

Once the regression metamodel has been estimated, its coefficients can be used to assess the significance and relative importance of the different parameters in affecting the output. To assess the significance of the effect of a given parameter x_i , we can use a standard t -test for the magnitude of β_i (Myers and Montgomery 1995). To assess the relative importance of the different parameters, we can compare the magnitudes of their respective coefficients. However, such a comparison is complicated by the fact that different parameters might have different scales. It is therefore common practice to standardize each model parameter relatively to its minimal and maximal value in the experiment for estimating the metamodel. In this case, the values of the different β_i indicate the relative importance of the different x_i , given their ranges in the experiment (Kleijnen 1992). The minimal and maximal values of a parameter are typically represented by -1 and $+1$, respectively. Given these boundaries, the standardized value (also called *coded unit*) of x_i can be calculated by

$$x_i = \frac{\rho_i - \bar{\rho}_i}{(\rho_{i,\max} - \rho_{i,\min}) / 2}, \quad (7.5)$$

where ρ_i refers to the value of the parameter x_i on its original scale, $\rho_{i,\max}$ and $\rho_{i,\min}$ refer to the maximal and minimal values of x_i used in the experiment, and $\bar{\rho}_i$ is defined as $(\rho_{i,\max} + \rho_{i,\min}) / 2$ (Kleijnen 2005, p. 290).

Fig. 7.1 Example of predicted output of simulation model with two parameters in coded units. The *surface* shows predicted values based on the metamodel (\hat{z}), and the *points* show the outcomes of the simulation model (z); *grey/black* colouring of the points indicates that the observed outcome is lower/higher than the predicted outcome



We can also use the coefficients to predict the output of the simulation model. Figure 7.1 illustrates this for the following second-order polynomial, which was estimated from the output of a fictive simulation model with two parameters:

$$\hat{z} = 1.09 - .18x_1 + .10x_2 + .05x_1x_2 + .16x_1^2 - .01x_2^2. \tag{7.6}$$

The figure shows as points the observed output (z) that was used to estimate the metamodel and as a surface the predicted output (\hat{z}) over the ranges of x_1 and x_2 used in the experiment (in coded units). Figure 7.1 suggests that the metamodel given in Eq. 7.6 is a valid approximation of the behaviour of the underlying simulation model, given that the predictions are close to the observed values. In the next section, I describe how we can assess the validity of metamodels more formally.

7.2.2 Validation


The discussion up to this point suggests that there is some degree of freedom in selecting the precise form of the polynomial and different forms will perform differently in approximating $f(\cdot)$. Finding the appropriate form is crucial for drawing valid conclusions about the simulation model. For example, if we estimate $f(\cdot)$ with a first-order polynomial under the assumption that there are no interactions between model parameters, we might draw invalid conclusions from the resulting function if the parameters actually interact with each other in affecting the output. Given this freedom, Kleijnen and Sargent (2000) and Kleijnen (1992) highlighted that specifying a regression metamodel is often an iterative process. Before we run experiments with the simulation model, we typically have acquired some prior knowledge (e.g. based on knowledge about the real social system, experiences

during debugging the model, etc.) that enables us to specify a “*tentative regression model*” (Kleijnen 1992, p. 299). If it turns out that this model describes the observed relations poorly, we can augment the metamodel (e.g. by adding or removing quadratic effects) and assess the fit of the new model. This procedure stops when we have found a function that accurately describes the observed relations.

There are three criteria that are commonly used for assessing the validity of a regression metamodel. First, the estimated model should not violate the assumptions of ordinary least squares regression (Myers and Montgomery 1995, p. 41ff). That is, (1) the distribution of the approximation errors should be normal, (2) the variance of the approximation errors should be constant for all parameter combinations, (3) the average of the approximation errors should be zero for all parameter combinations, and (4) the expected correlation between the approximation errors of any two observations should be zero. The first three assumptions can be assessed with standard methods for assessing the distribution of residuals described in text books on multiple regression; the third assumption can additionally be assessed with a formal *lack of fit test* (Rao 1959, see details below). The fourth assumption is typically satisfied when each simulation run is initialized with a different random number seed (Kleijnen 2008). A metamodel that violates one or more of these assumptions might lead to incorrect conclusions and therefore needs to be modified. For example, if there is evidence that the assumption of constant error variance is violated, this might be due to an interaction between some model parameters that has not been included in the metamodel yet. Alternatively, the simulation model might generate outcomes that are inherently heteroscedastic (i.e. the variance in the outcomes systematically in-/decreases as certain parameters increase). In this case, we might consider transforming the output (e.g. by a logarithmic transformation) or using the method of weighted least squares instead of ordinary least squares (Kleijnen 2008; Tunali and Batzman 2000).

Second, the metamodel should fit the observed data and should explain a large and significant part of the variance in the observed output. That is, the value of the coefficient of determination (R^2 and $R^2_{adjusted}$) should be high and the associated F -test should be significant (Kleijnen 2008). Furthermore, the fit of a regression metamodel can be assessed with a formal lack of fit test. Such tests assess whether predictions of the metamodel systematically deviate from the observed simulation outcomes. To illustrate this, consider Fig. 7.2. The points in the figure show the outcome of a fictive simulation model that was observed at five levels of the parameter x_1 . The black line shows the predictions of a first-order polynomial, which was estimated from this data. Intuitively speaking, the metamodel lacks fit given that it neglects the nonlinearity in the relation between x_1 and z and therefore systematically under-estimates z at high and low levels of x_1 , but over-estimates it at intermediate levels of x_1 . More formally (cf. Myers and Montgomery 1995), we can assess whether there is evidence for such lack of fit by partitioning the total approximation error (SS_E) in the regression model into pure error (SS_{PE}) that can be attributed to variation in the simulation outcomes (e.g. due to stochastic processes in the simulation model) and into error that can be attributed to a lack of fit of the regression model (SS_{LOF}), so that

Fig. 7.2 Example of lack of fit of regression metamodel. The line shows predicted values based on a first-order polynomial (\hat{z}), and the black points show the outcomes of the simulation model (z) that underlie this metamodel



$$SS_E = SS_{PE} + SS_{LOF}. \tag{7.7}$$

For calculating SS_{PE} , the outcome of the simulation model needs to be observed several times for at least some parameter values (see Sect. 7.2.3 for details on how/where to collect multiple observations). This is illustrated in Fig. 7.2, given that the model outcome was observed three times for each of the five parameter levels. Intuitively speaking, the variation in the outcome at each parameter level provides information for calculating the approximation error that can be attributed to random variation in model outcomes. More formally, SS_{PE} and SS_{LOF} are calculated as follows. Assume that there is a simulation model with k parameters. Let $l = 1, \dots, v$ be the unique parameter combinations that have been observed. Furthermore, let n be the total number of simulation runs that we have conducted and let $q = 1, \dots, n_l$ be the number of simulation runs that have been conducted at each of the different parameter combinations, so that $n = \sum_{l=1}^v n_l$. Based on this, the value of SS_{PE} is calculated by

$$SS_{PE} = \sum_{l=1}^v \sum_{q=1}^{n_l} (z_{lq} - \bar{z}_l)^2, \tag{7.8}$$

where \bar{z}_l is the average of the observed output over all observations for the l^{th} parameter combination. Thus, Eq. 7.8 holds that SS_{PE} is based on the squared deviations of the individual outputs for a given parameter combination from the average output for this combination, summed over all unique combinations. The value of SS_{LOF} is calculated by

$$SS_{LOF} = \sum_{l=1}^v n_l (\bar{z}_l - \hat{z}_l)^2. \tag{7.9}$$

This means that SS_{LOF} is based on the squared deviations of the predicted output from the observed average output for a given parameter combination, weighted by the number of observations for the combination, summed over all unique combinations. Based on this, the test statistic for the formal lack of fit test is calculated by

$$F_{LOF} = \frac{SS_{LOF}/(v-p)}{SS_{PE}/(n-v)}, \quad (7.10)$$

in which p is the number of regression coefficients in the metamodel (including the intercept) and which follows an F -distribution with $v-p$ degrees of freedom for SS_{LOF} and $n-v$ degrees of freedom for SS_{PE} . When F_{LOF} is significant, we cannot reject the null hypothesis that there is lack of fit.

Third, and finally, the regression metamodel should be able to accurately predict the outcome of the simulation model for parameter combinations that fall within the boundaries of the parameter space that was used in the experiment for estimating the metamodel. That is, the model should have a high level of predictive adequacy (Kleijnen 2008). This adequacy can be assessed by collecting data for additional parameter combinations that were not included in the original experiment and comparing the observed outputs with the predictions from the metamodel that was estimated from the original data. If the observed outcomes are close to the predictions, the metamodel has high predictive adequacy.

7.2.3 Experimental Designs

The choice of the type and form of the metamodel determines the design of experiments (DOE) for collecting output data. DOE is the process of planning experiments so that the metamodel can be estimated effectively and efficiently (Antony 2003, p. 7). Effectively means that we collect the data necessary to draw valid conclusions about the behaviour of the simulation model from the selected metamodel; efficiently means that we collect this data with as little computational effort as possible (Kleijnen et al. 1992; Lorscheid et al. 2012). In the literature on DOE, the output of the simulation model is commonly referred to as *response* and the parameters of the simulation model are referred to as *factors*; the levels of a given parameter are referred to as *factor levels*. In the remainder of this chapter, I use these terms (i.e. output/response, parameter/factor, and parameter level/factor level) interchangeably. Furthermore, the schedule of the combinations of different factor levels that are included in a simulation experiment is called *experimental design*; the different factor combinations at which the simulation model needs to be run are called *design points*. Finally, the highest and lowest values chosen for each factor determine the *experimental region* that the design covers.

Any regression metamodel can only be valid for the experimental region for which it has been estimated, and smaller the experimental region, the more accurate

the model is likely to be (Myers and Montgomery 1995). The goals of the analysis should therefore guide the selection of the experimental region. To illustrate this, consider a simulation model with the two parameters x_1 and x_2 that both can vary from 0 to 100. If we choose 0 and 100 as the lower and upper boundaries for both x_1 and x_2 in our simulation experiment, the experimental region coincides with the operational boundaries of the simulation model. The resulting least squares estimators will therefore provide insights into the general behaviour of the simulation output over the entire parameter space, while smoothing out deviations from this general behaviour that might exist in some parts of the parameter space. If we want to learn in more detail about the behaviour of the simulation model in a smaller portion of the parameter space (say between $x_1 \in [20; 40]$ and $x_2 \in [70; 90]$), we might benefit from reducing the experimental region to this area. The resulting metamodel is potentially more precise, at the cost of being valid only for a smaller region of the parameter space.

Once the experimental region has been defined, it needs to be determined what data need to be collected within this region. In general, the more complex the polynomial, the more points we need to include in the experimental design (Kleijnen 2005; Kleijnen et al. 2005; Simpson et al. 2001). One of the most commonly used designs for estimating first-order polynomials is the full two-level factorial design (Kleijnen et al. 2005). In this design, each factor has two levels and there is one design point for each possible combination of these levels across the factors so that there are 2^k design points. This design makes it possible to estimate the main effect of each parameter and of all possible two-way interactions. To be able to conduct a formal lack of fit test, it is common practice to augment this design with a number of n_c centre runs ($n_c > 1$), which are located at the 0-coordinates of each factor in terms of coded units. For illustration, Table 7.1 shows the design points of a full two-level factorial design based on two factors, which has been augmented with two centre runs; panel (a) of Fig. 7.3 shows the experimental region that this design covers.

One of the most commonly used designs for estimating second-order polynomials is the central composite design. A standard central composite design consists of a full two-level factorial design which is augmented with $2k$ axial points and n_c centre runs. The axial points are located at distance α from the centre of the design, which is typically determined by $\sqrt[4]{2^k}$, locating the axial points outside the -1 and 1 borders of the original experimental region (see Table 7.1 and panel (b) of Fig. 7.3 for an illustration). This makes the estimation of quadratic effects within the experimental region maximally efficient, but can be problematic when the boundaries of the experimental region correspond with the operational boundaries of the simulation parameters. We can solve this problem by using a standard central composite design and scaling it down, so that the original boundaries of the design move closer to the centre and the axial points are located at the original $-1/1$ boundaries. Alternatively, we can set $\alpha = 1$, so that the original factorial design remains unchanged and the axial points are located at the $-1/1$ boundaries of the original experimental region. In the first case, the design becomes an *inscribed* central composite; in the second case it becomes a *face centred* central composite (see Table 7.1 and panels (c) and (d) of Fig. 7.3 for illustrations). In direct


Table 7.1 Examples of different types of experimental design in coded units

Design point	Full two-level factorial		Standard central composite ($\alpha = 1.414$)		Inscribed central composite ($\alpha = 1.414$)		Face centred central composite ($\alpha = 1$)	
	x_1	x_2	x_1	x_2	x_1	x_2	x_1	x_2
1	-1	-1	-1	-1	-0.707	-0.707	-1	-1
2	-1	1	-1	1	-0.707	0.707	-1	1
3	1	-1	1	-1	0.707	-0.707	1	-1
4	1	1	1	1	0.707	0.707	1	1
5 ^a	0	0	0	0	0	0	0	0
6 ^a	0	0	0	0	0	0	0	0
7 ^b	-	-	-1.414	0	-1	0	-1	0
8 ^b	-	-	1.414	0	1	0	1	0
9 ^b	-	-	0	-1.414	0	-1	0	-1
10 ^b	-	-	0	1.414	0	1	0	1

^aCentre run

^bAxial point

Fig. 7.3 Examples of (a) full two-level factorial design, (b) standard central composite design, (c) inscribed central composite design, and (d) face centred central composite design in coded units. The grey area illustrates the experimental region that the original full two-level factorial design covers, which underlies the designs shown in panels (a), (b), and (d).
 ● factorial point, □ centre run, ○ axial point



comparison, an inscribed central composite design is more effective for estimating quadratic effects, but a face centred design is more accurate in describing the observed relations at the corners of the experimental region (for details see Myers and Montgomery 1995).¹

¹For a discussion of additional popular designs see the chapter by Hilton and Bijak in this volume.

Finally, agent-based computational models are typically stochastic, so that the observed response for a given parameter combination tends to vary over repeated simulation runs. To deal with this variation, it is common practice to repeat the simulation multiple times at a given design point and to use the average of the simulation outputs for estimating the metamodel (Lorscheid et al. 2012; Thiele et al. 2014), so that the metamodel is fitted to the mean simulation outcome. In this case each average represents one observation. To illustrate this, consider the full two-level factorial design shown in Table 7.1. If we conduct ten simulation runs for each of the six design points and use the average of these runs for estimating the metamodel, we still have only six observations of the output.

7.3 Applied Example

In this section, I illustrate how metamodels can be used to learn about the behaviour of a simulation model and what decisions need to be taken in this process. I first outline the simulation model and subsequently present the details of the analysis. I have conducted all analyses in the statistical programming environment R (R Core Team 2014) and have estimated the regression metamodels with the package `rsm` (Lenth 2009). In the supplementary material that accompanies this chapter, I provide all files necessary to reproduce the results that I present below, including the code of the simulation model.

7.3.1 Simulation Model

Grow and Van Bavel (2015) present an agent-based computational model that explicates some of the social mechanisms that might have linked the recent reversal of gender inequality in education with changes in educational assortative mating (EAM) across Europe. The notion of assortative mating refers to any pattern of relationship formation based on personal attributes that deviates from the pattern that we might expect if relations were formed purely at random (Schwartz 2013). Across twentieth century Europe, EAM has been mostly homogamic (i.e. partners were similarly educated) but if there was a difference in educational attainment within couples, the man tended to be more educated than the woman. Until the 1970s, this pattern was compatible with the surplus of highly educated men on the marriage market. However, since then the relative educational attainment of men and women has changed drastically: over the years, the number of highly educated women has increased relatively to that of men and this has made the traditional pattern of EAM infeasible (Van Bavel 2012). As a consequence, the number of couples in which the man is more educated than the woman (hypergamic couples) has decreased and the number of couples in which the woman is more educated

than the man (hypogamic couples) has increased, whereas the share of homogamic couples remained largely unchanged (De Hauw et al. 2015; Esteve et al. 2012).

To better understand the mechanisms that have linked changes in educational attainment among men and women with changes in patterns of EAM across Europe, Grow and Van Bavel (2015) developed a model that builds on earlier research on the interplay between the composition of marriage markets and patterns of assortative mating (e.g. Oppenheimer 1988) and on earlier simulation work that has centred on human mate search (e.g. Simão and Todd 2002). At the core of the model are the preferences that underlie individual mating decisions. More specifically, the model focuses on individuals' preferences for the educational attainment, earnings prospects, and age of prospective partners. It assumes that both men and women look for spouses who are similar to them in educational attainment and who have high earnings prospects (e.g. Kalmijn 1994). Additionally, it assumes that women tend to look for spouses who are somewhat older than themselves, whereas men tend to look for spouses who are in their mid-twenties (e.g. England and McClintock 2009).

A detailed description of the model is provided in Grow and Van Bavel (2015). Here, I provide a brief outline of the model and highlight the elements that are relevant for illustrating the use of regression metamodels. The model consists of male (m) and female (f) agents who try to find a heterosexual partner for a long-term relation in the form of marriage. The search for a spouse takes place on a marriage market in which meeting opportunities are structured by the educational system. That is, a given male agent and female agent are more likely to meet each other when they are in the same stage of their educational career (e.g. both are attending high school, both have left school already, etc.) than when they are in different stages.

Whenever two agents meet, they need to decide whether they want to start dating. Once they are dating, they can decide whether they would like to marry. The importance that agents attach to the education, earnings prospects, and age of prospective partners when making dating and marriage decision is governed by six parameters. The parameters w_s^m and w_s^f govern the importance that male and female agents attach to similarity in education (s) with prospective partners; w_y^m and w_y^f govern the importance that they attach to earnings prospects (y); w_a^m and w_a^f govern the importance that they attach to age (a). For each parameter, a larger value implies higher importance of the respective characteristic, but this increase in importance has a decreasing marginal effect. For example, increases in w_s^m initially lead to strong decreases in the willingness of male agents to date/marry female agents who do not have the same educational attainment as themselves. This decrease continues as w_s^m increases, but the marginal effect becomes lower at higher levels of w_s^m . At some point, male agents become so unwilling to date/marry female agents who are not a perfect educational match that increasing w_s^m any further has virtually no additional effect on their mating decisions.

The model simulates individual mate search over individuals' entire life course and includes simple assumptions about mortality and reproduction. This makes it possible to model mate search over successive cohorts. To generate plausible

agent cohorts in terms of educational attainment and earnings prospects, Grow and Van Bavel (2015) used data provided by the International Institute for Applied Systems Analysis/Vienna Institute for Demography (KC et al. 2010; Lutz et al. 2007) and data from the European Community Household Panel.² For validating model outputs, they used data from rounds 5 and 6 of the European Social Survey, collected in 2010 and 2012. The combination of these data sources enabled them to simulate mate search behaviour under realistic marriage market conditions among individuals born between 1921 and 2012, and to study patterns of EAM in 12 Western European countries.

In this chapter, I use the same input data and simulation settings as Grow and Van Bavel (2015) and focus on the pattern of EAM based on the input data for Belgium. My goal is to show how regression metamodels can be used to study the effects that the six focal model parameters w_s^m , w_s^f , w_y^m , w_y^f , w_a^m , and w_a^f have on the percentages of hypergamic, homogamic, and hypogamic couples (dating or married) among agents who are old enough to have attained their highest educational degree (i.e. in the ages between 24 and 79 years) at the end of a simulation run (i.e. in simulation years 2010 and 2012). I estimated one metamodel for each of the three outcomes. Note that all other parameters are based on the calibrated model described by Grow and Van Bavel (2015).

7.3.2 Sensitivity Analysis

7.3.2.1 Experimental Region

In the model of Grow and Van Bavel (2015), selecting the experimental region is complicated by the fact that each of the six parameters has a lower operational boundary at 0, but none of the parameters has an upper operational boundary. In their search for a parameter combination that could recreate observed patterns of EAM across Europe, Grow and Van Bavel (2015) considered values between 0 and 2 for w_s^m , w_s^f , w_y^m , and w_y^f , and between 0 and 20 for w_a^m and w_a^f . The reason was that based on their experience with the simulation model, they expected that within this region there might be a parameter combination that generates outputs that fit well with the observed patterns of EAM. In this chapter, I focus on a larger experimental region. More specifically, I focus on the region between the values 0 and 4 for w_s^m , w_s^f , w_y^m , and w_y^f , and between 0 and 40 for w_a^m and w_a^f , to be able to study the model's behaviour between parameter boundaries that could be considered 'extreme' from a substantive point of view. These boundaries can be considered extreme for the following reasons:

²Eurostat, European Commission and the national statistical offices collecting the data have no responsibility for the results and conclusions which were drawn in this paper on the basis of the European Community Household Panel data.

- When $w_s^m = w_s^f = 0$, agents do not care about the educational attainment of prospective partners. By contrast, when $w_s^m = w_s^f = 4$, the attraction that agents who belong to the lowest and the highest possible educational attainment category feel for each other approximates zero, regardless of all of other mating relevant characteristics.
- When $w_y^m = w_y^f = 0$, agents do not care about the earnings prospects of prospective partners. By contrast, when $w_y^m = w_y^f = 4$, the attraction that agents feel for somebody who has the lowest possible earnings prospects approximates zero, regardless of all of other mating relevant characteristics.
- When $w_a^m = w_a^f = 0$, agents do not care about the age of prospective partners. By contrast, when $w_a^m = w_a^f = 40$, the attraction that agents feel for somebody who is 10 or more years older or younger than the ideal age that they prefer approximates zero, regardless of all other mating relevant characteristics.

7.3.2.2 Tentative Regression Model

I expected that at least some of the parameters would interact in affecting the simulation output. Consider, for example, the possible interplay between female agents' preferences for education (w_s^f) and age (w_a^f). When female agents attach high importance to similarity in education (i.e. the value of w_s^f is high) but do not care much about the age of prospective partners (i.e. the value w_a^f is low), it might be relatively easy for them to find similarly educated mates, given that they can draw on all male members of the marriage market regardless of their age. Thus, an increase in w_s^f might lead to a strong increase in homogamy when the value of w_a^f is low. Yet, if female agents also have strong preference for partners who are slightly older than themselves (i.e. the value of w_a^f is high), the pool of potential partners shrinks drastically and might offer them fewer opportunities to find similarly educated men. Thus, an increase in w_s^f might lead to a weaker increase in homogamy when the value of w_a^f is high.

Second, as discussed Sect. 7.3.1, each of the six model parameters has a decreasing marginal effect on agents' willingness to date/marry somebody who is less than an ideal match in terms of their own preferences. I thus expected that these parameters might also have decreasing marginal effects on model outputs. If this is the case, we might expect some curvature in the observed relations.

Based on the foregoing considerations, I decided to use full second-order polynomials that included all possible two-factor interactions.

7.3.2.3 Experimental Design

Because of my choice of full second-order polynomials, I decided to use a central composite design in the simulation experiments. Yet, given that the lower boundary of the experimental region was constrained by the lower operational boundaries

Table 7.2 Parameter values in original and coded units for central composite design used in simulation experiment ($\alpha = 2.828$)

Parameter	Original units					Coded units				
	w_s^m	0	1.293	2	2.707	4	-1	-0.354	0	0.354
w_s^f	0	1.293	2	2.707	4	-1	-0.354	0	0.354	1
w_y^m	0	1.293	2	2.707	4	-1	-0.354	0	0.354	1
w_y^f	0	1.293	2	2.707	4	-1	-0.354	0	0.354	1
w_a^m	0	12.93	20	27.07	40	-1	-0.354	0	0.354	1
w_a^f	0	12.93	20	27.07	40	-1	-0.354	0	0.354	1

of the simulation parameters, I opted for an inscribed central composite design, to avoid that the axial points would be located outside these boundaries. Thus, I chose α as $\sqrt[4]{2^6} = 2.828$ and scaled the resulting design down to stay within the experimental region. Table 7.2 shows the different parameter values in both original and coded units. The complete design consisted of 84 design points and contained $n_c = 8$ centre runs for conducting a formal lack of fit test (i.e. only the runs at the centre of the design were repeated; see the supplementary material for the full experimental schedule). Given the stochastic nature of the simulation model, I used the average of 50 simulation runs per design point for the analyses, leading to a total of 4,200 simulation runs.

Conducting one simulation run took on average 36 s on an Intel Core i7-3770 processor with 3.40 GHz and eight cores, leading to a total computation time of about 42 h (i.e. about 5.25 h per core). To illustrate the efficiency of the approach selected here, it is helpful to note that a sensitivity analysis with the simple approach discussed in the introduction would have consisted of $5^6 \times 50 = 781,250$ runs, if there had been five levels per parameter as in the inscribed central composite design. Given an average computation time of 36 s, it would have taken more than 7,800 h of computation time to conduct these runs. Even if we had focused on only three levels per parameter, the experiment would have consisted of $3^6 \times 50 = 36,450$ runs which would have taken about 364 h of computation time. Furthermore, it is helpful to note that the parameter estimates that I present below can be used to instantaneously gain information about the model outcome for new parameter combinations within the experimental region, without having to invest the 30 min of computation time that it would take to run the model 50 times for the new parameter combination.

7.3.2.4 Model Validation

Table 7.3 shows the coefficients of the three regression metamodels. Before I could interpret these results, I needed to check whether the metamodels are valid representations of the observed relations between the parameters and outputs.

Table 7.3 Results of sensitivity analyses

Parameters	Hypergamic		Homogamic		Hypogamic	
	β	p	β	p	β	p
Intercept	5.83	**	84.91	**	9.25	**
Main effects						
w_s^m	-3.27	**	7.50	**	-4.24	**
w_s^f	-2.82	**	8.04	**	-5.22	**
w_y^m	-1.55	**	0.87	**	0.68	**
w_y^f	0.14		0.49	*	-0.64	**
w_a^m	-0.98	**	2.53	**	-1.55	**
w_a^f	-0.53	**	0.66	**	-0.12	
Interaction effects						
$w_s^m \times w_s^f$	0.91		-2.29	**	1.38	*
$w_s^m \times w_y^m$	1.32	*	-1.07		-0.25	
$w_s^m \times w_y^f$	-0.45		1.14		-0.69	
$w_s^m \times w_a^m$	-0.37		-0.12		0.49	
$w_s^m \times w_a^f$	-0.08		0.29		-0.21	
$w_s^f \times w_y^m$	1.59	**	-0.58		-1.02	
$w_s^f \times w_y^f$	-0.05		-0.92		0.97	
$w_s^f \times w_a^m$	0.53		-1.73	*	1.20	*
$w_s^f \times w_a^f$	0.10		0.03		-0.13	
$w_y^m \times w_y^f$	-0.11		0.35		-0.23	
$w_y^m \times w_a^m$	-0.80		1.77	*	-0.97	
$w_y^m \times w_a^f$	-0.42		-0.04		0.46	
$w_y^f \times w_a^m$	-0.42		0.69		-0.28	
$w_y^f \times w_a^f$	-0.26		2.19	**	-1.93	**
$w_a^m \times w_a^f$	-0.82		0.90		-0.08	
Quadratic effects						
$w_s^m \times w_s^m$	1.19	**	-1.70	**	0.52	
$w_s^f \times w_s^f$	-0.05		-1.86	**	1.91	**
$w_y^m \times w_y^m$	0.14		0.65		-0.79	
$w_y^f \times w_y^f$	0.02		-0.14		0.12	
$w_a^m \times w_a^m$	1.30	**	-2.15	**	0.85	
$w_a^f \times w_a^f$	-0.70		0.61		0.08	

Estimates are based on coded parameter units, * $p < 0.05$, ** $p < 0.01$

To assess whether the models fitted the data well and to assess whether the average of the approximation errors was zero for all parameter combinations, I conducted three formal lack of fit tests, one for each metamodel. Table 7.4 shows the results of these tests and indicates that none of them was significant. Additionally, I assessed whether each model explained a large and significant share of the variance in the respective output. Table 7.5 shows that the three coefficients of determination (R^2) and their adjusted versions ($R^2_{adjusted}$) were high and significant in all three cases.

Table 7.4 Lack of fit tests for main analyses

Source of error	DF	Value	F	p
Hypergamic				
SS_E	56	13.85		
SS_{LOF}	49	12.40	1.22	
SS_{PE}	7	1.45		
Homogamic				
SS_E	56	30.62		
SS_{LOF}	49	27.89	1.46	
SS_{PE}	7	2.73		
Hypogamic				
SS_E	56	18.96		
SS_{LOF}	49	15.60	0.66	
SS_{PE}	7	3.36		

DF degrees of freedom

* $p < 0.05$, ** $p < 0.01$

Table 7.5 Coefficients of determination for main analyses

Model	DF1, DF2	$R^2(R^2_{adjusted})$	F	p
Hypergamic	27, 56	0.94 (0.92)	35.50	**
Homogamic	27, 56	0.98 (0.97)	90.06	**
Hypogamic	27, 56	0.96 (0.95)	55.29	**

DF degrees of freedom

* $p < 0.05$, ** $p < 0.01$

To assess whether the distribution of the approximation errors was normal and whether their variance was constant for all parameter combinations, Fig. 7.4 shows quantile-quantile plots of the approximation errors and shows the relation between the model predictions and approximation errors. The figure suggests that the distribution of the residuals followed a normal distribution and that the spread was similar across all parameter combinations in all three regression metamodels.

I assessed the predictive adequacy of the three metamodels by randomly selecting 20 new parameter combinations from the experimental region and comparing the observed outcomes for each combination with the predictions of the metamodels. Table 7.6 shows the selected parameter combinations and shows the results for the case of hypergamic couples. Figure 7.5 plots the predicted values against the observed values for all three outcomes and the results suggest that the predictive adequacy of the three metamodels was high.

Finally, each simulation run used a different seed for initializing random numbers. This implies that the assumption of non-correlated residuals was also satisfied.

Taken together, the results suggest that each of the three metamodels was a valid representation of the observed associations.

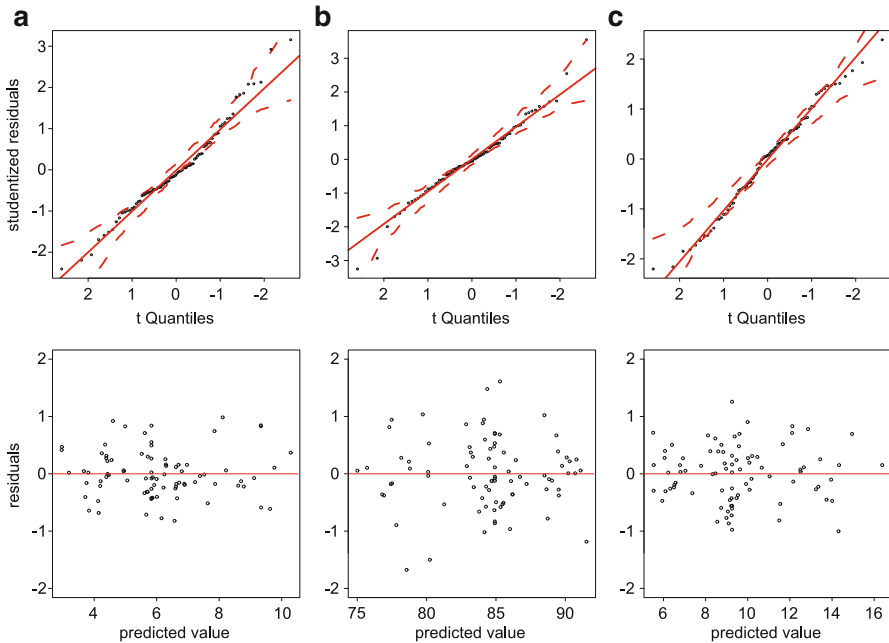


Fig. 7.4 Inspection of distribution of approximation errors and variance of approximation errors for main analyses, (a) hypergamic, (b) homogamic, and (c) hypogamic couples. The *first row* shows *quantile-quantile plots* and *dashed lines* show a 95 % confidence envelope (Fox 2008). The *second row* plots predicted values against the residuals

7.3.2.5 Simulation Results

The results shown in Table 7.3 suggest that many of the simulation parameters had significant main effects and that some of them affected simulation outputs in a nonlinear manner. Additionally, there were some significant two-way interactions between parameters. To aid the interpretation of the main effects and nonlinear effects, I plotted for each preference dimension (i.e. education, earnings prospects, and age) the male and female versions of the respective parameters against each other and inspected the predicted outputs for each of the three models (Figs. 7.6, 7.7, and 7.8); the contours at the bottom of each figure aim to facilitate the interpretation of the results. Subsequently, for illustrative purposes, I inspected some of the interaction effects that were not covered by these illustrations (Fig. 7.9). In all figures, all of the model parameters that are not shown are fixed at 0 in terms of coded units. Note that the three outcome measures are correlated with each other. For example, if the metamodels suggest that a given preference leads to an increase in hypergamy, this increase necessarily leads to a decrease in at least one of the two other outcomes.

Table 7.6 Excerpt of results of validation experiment

Design point	Parameter values in coded units						Hypergamic		
	w_s^m	w_s^f	w_y^m	w_y^f	w_a^m	w_a^f	z	\widehat{z}	$z - \widehat{z}$
1	-0.412	-0.344	0.896	0.428	-0.807	-0.526	8.88	8.71	0.17
2	-0.042	0.192	0.327	-0.266	-0.714	0.989	6.12	5.69	0.43
3	-0.084	0.343	-0.382	-0.574	-0.644	-0.184	6.16	6.05	0.11
4	0.164	0.319	-0.512	-0.224	-0.555	0.616	5.15	5.38	-0.22
5	-0.702	-0.170	-0.149	0.680	-0.228	0.313	9.23	10.14	-0.91
6	0.175	-0.172	0.494	0.964	-0.056	0.307	4.86	4.78	0.08
7	0.538	0.893	0.855	-0.486	-0.022	0.018	2.27	3.04	-0.77
8	-0.732	-0.586	0.308	0.529	0.051	-0.969	10.83	10.19	0.65
9	0.983	-0.888	-0.231	0.813	0.168	0.026	5.35	5.35	0.00
10	0.325	0.680	-0.156	-0.775	0.441	0.157	3.56	3.22	0.34
11	-0.137	-0.235	0.154	-0.239	0.456	0.026	6.41	6.38	0.03
12	0.523	-0.553	0.118	-0.345	0.664	-0.504	5.28	5.59	-0.31
13	0.813	-0.106	-0.774	0.225	0.683	0.078	5.27	4.72	0.55
14	-0.062	-0.357	-0.740	-0.114	0.724	-0.838	9.46	9.25	0.21
15	0.100	0.024	-0.467	0.151	0.726	0.451	5.45	5.75	-0.30
16	0.967	0.243	0.358	-0.276	0.733	0.344	2.38	2.64	-0.26
17	-0.235	0.113	-0.733	0.090	0.768	-0.286	8.50	8.39	0.11
18	0.036	0.302	0.811	0.130	0.896	-0.491	4.06	4.44	-0.38
19	0.035	-0.617	0.564	0.473	0.955	-0.150	5.70	5.67	0.03
20	0.118	-0.008	-0.035	-0.403	0.982	0.006	5.03	5.93	-0.90




Fig. 7.5 Predicted and observed outputs based on validation simulation experiment

Consider first the effects of agents’ preferences for similarly educated partners (Fig. 7.6). The results suggest that increasing both male and female agents’ preferences for similarly educated partners (i.e. w_s^m and w_s^f respectively) led to a decrease in the shares of hypergamic and hypogamic couples, and to an increase in the share of homogamic couples. Yet, these effects were subject to two-way interactions and nonlinearity. Figure 7.6 facilitates the understanding of these complex effects. Consider, for example, the case of homogamic couples. Increasing w_s^m or w_s^f each led to an increase in homogamy, but this increase became weaker if the respective




Fig. 7.6 Predicted outputs based on the interplay between preferences for similarly educated partners among male and female agents (w_s^m and w_s^f). All parameters that are not shown are fixed at the value 0 in terms of coded units




Fig. 7.7 Predicted outputs based on the interplay between preferences for high earnings prospects among male and female agents (w_y^m and w_y^f). All parameters that are not shown are fixed at the value 0 in terms of coded units




Fig. 7.8 Predicted outputs based on the interplay between preferences for age among male and female agents (w_a^m and w_a^f). All parameters that are not shown are fixed at the value 0 in terms of coded units

other parameter was at a high level. Substantively this means, for example, that even when female (male) agents are willing to date/marry someone who differs from them in educational attainment, heterogamy will still be low if male (female) agents are not also willing to do so. Furthermore, the decreasing marginal effect of each parameter can be attributed to the facts (1) that the parameters have decreasing marginal effects on agents' willingness to date/marry somebody who deviates from their ideals and (2) that there is a limit to the level of homogamy/heterogamy that




Fig. 7.9 Predicted outputs based on the interplay between preferences for age among male agents (w_a^m) and preferences for education among female agents (w_s^f). All parameters that are not shown are fixed at the value 0 in terms of coded units

can exist in an agent population (e.g. once all similarly educated individuals are partnered with each other, homogamy cannot increase anymore). The closer the agent population is already to these limits, the less a change in a given parameter might contribute to homogamy/heterogamy.

Consider next the effects of preferences for high earnings prospects (Fig. 7.7). The results suggest that an increase in the preferences for high earnings prospects among male agents (i.e. an increase in w_y^m) led to an increase in homogamic and hypogamic couples, and to a decrease in hypergamic couples. An increase in preferences for earnings prospects among female agents (i.e. an increase in w_y^f) also tended to increase homogamy, but tended to decrease hypogamy. Substantively this means that the correlation between education and earnings prospects leads preferences for earnings prospects to have an indirect effect on EAM. That is, the more importance male and female agents attach to the earnings prospects of their partners, the more likely they are to date/marry somebody who is equally or more educated, given that such agents tend to be more attractive in terms of earnings prospects than lower educated agents.

Consider now the effects of preferences for age (Fig. 7.8). The results suggest that the stronger male agents preferred partners who are in their mid-twenties (i.e. if w_a^m is high), the more likely they were to date/marry somebody who had a similar educational background, and the less likely they were to date/marry somebody with different (particularly with a higher) educational background. The preference of female agents for slightly older partners (i.e. the value of w_a^f) mattered much less for these outcomes, but also led to an increase in homogamy and a decrease in hypergamy. For interpreting these results it is important to remember that in older cohorts men tend to be more educated than women, whereas in younger cohorts women tend to be more educated than men. Thus, when male agents prefer partners who are in their mid-twenties, this implies that especially older and highly educated males look for partners in cohorts in which they are more likely to find somebody who is similarly educated. At the same time, young and highly educated female agents have better chances to find a similarly educated partner in older cohorts. These processes together lead to an increase in the share of homogamic couples and

lead to a decrease in the shares of heterogamic couples, in particular hypogamic couples.

Consider finally the interaction effects that existed between the importance that male agents attach to the age of prospective partners (w_a^m) and the importance that female agents attach to similarity in education (w_s^f) (Fig. 7.9). The results suggest that when male agents did not care much about the age of their partners (i.e. when w_a^m was low), an increase in female preferences for similarly educated partners (i.e. an increase in w_s^f) had a stronger positive effect on homogamy than when male agents cared strongly about age (i.e. when w_a^m was high). The opposite was the case for hypogamy. That is, the less male agents cared about the age of prospective partners, the more strongly an increase in female agents' preferences for similarly educated partners led to a decrease in hypogamy. From a substantive point of view, this underlines the role that age preferences play in combination with the composition of the marriage market in terms of educational attainment. When male agents prefer women who are in their mid-twenties, especially older, highly educated males are looking for partners in a segment of the mating market that is conducive to homogamy. This leads to a higher level of homogamy that decreases the effect of female agent's preferences for similarly educated partners. The reason for this latter effect is that there is a limit to the level of hypergamy, homogamy, and hypogamy that can exist in a given population, as explained earlier.

7.4 Discussion and Conclusion

In this chapter, I have demonstrated the use of regression metamodells for sensitivity analysis of computational simulation models. I hope that the applied example has illustrated the potential benefits for understanding the behaviour of complex agent-based models in computational demography. As indicated above, I believe that regression metamodells are an attractive tool for sensitivity analysis in the field of demography, because they are powerful and easily accessible for both model developers and audiences with a background in demography.

These advantages notwithstanding, regression metamodells also have their limitations. First, regression metamodells that are estimated with the ordinary least squares method are best suited for response surfaces that are smooth and 'well behaved', as was the case in the example shown here. For surfaces that cannot easily be represented by low-order polynomials, regression metamodells might be able to accurately describe the behaviour of the output over a small, local experimental region, but they might not be able to describe the behaviour of the output over the full ranges of all simulation parameters. For an example of such a situation see the chapter by Hilton and Bijak in this volume. In this case, we might benefit from choosing more complex metamodells, which make fewer assumptions about the data than regression metamodells (Kleijnen 2005). For example, Gaussian process emulators (Oakley and O'Hagan 2002) can deal with less regular surfaces, but require more statistical background knowledge to implement, which makes them

less accessible to wider audiences. Another approach to sensitivity analysis that has gained momentum in recent years is *model output variance decomposition* (e.g. Ligmann-Zielinska et al. 2014), in which the importance of the parameters of a simulation model is expressed in the amount of variation in the outcome that it accounts for. Compared to regression metamodels, this approach can more easily deal with nonlinearity in model behaviour, but it does not provide information about the direction of influence, given that variance indices range from 0 to 1.³

Furthermore, some of the social systems that have been studied with agent-based models are subject to discontinuous behaviour, in which changes in a given parameter up to a certain level might not have any noticeable effect on model outcomes, but increases beyond this level lead to drastic changes in the outcome. The seminal work of Schelling (1971) provides an example of such ‘tipping’ points in the case of residential segregation. Regression metamodels will have problems with describing such discontinuities, yet, the fact that a regression metamodel fits the data poorly might at least point to the existence of such discontinuity and might aid in finding the location of the discontinuity in the parameter space.⁴

Finally, when interpreting the coefficients of a regression metamodel based on coded data, it is important to keep the original scaling of the variables in mind, in particular when we investigate an experimental region that is smaller than the full parameter space. For example, the same linear effect for a given parameter might appear smaller if the coded values -1 and 1 represent the values 250 and 350 in original scaling, than when they represent 100 and 500 (Kleijnen 2008).

I aimed to acquaint the reader with some of the basic concepts of regression metamodels. There are a number of additional topics that I could not address here. For example, I have illustrated how regression metamodels can be used to predict the output for parameter combinations that are located within the experimental region. These predictions, in turn, can be used for optimization. In optimization, we search for a parameter combination that generates output values that are particularly high, low, or close to some predefined target. Grow and Van Bavel (2015) used this possibility for calibrating their simulation model with empirical data. That is, they determined a parameter combination that was most likely to generate outputs that were close to real-life patterns of EAM across Europe. The issue of optimization with regression metamodels is closely related to the response surface methodology, which is a set of tools that can be used to iteratively find optimal parameter combinations and is extensively described in Myers and Montgomery (1995).

Furthermore, I have shown how the predictive adequacy of a metamodel can be assessed by collecting simulation outputs for parameter combinations that were not included in the experimental design used for estimating the metamodel. This approach might not be feasible if the simulation model is computationally expensive; Kleijnen (2008) therefore describes a cross-validation approach that does not require additional simulation runs. With this approach, some of observations that

³I thank an anonymous reviewer for pointing this out.

⁴I thank an anonymous reviewer for pointing this out.

were used to estimate the metamodel are dropped and the model is re-estimated. The metamodel is valid if the parameter estimates are not sensitive to such omissions.

To conclude, I believe that regression metamodels have merit for use in agent-based computational demography. Given their simplicity, they are potentially a good first choice for learning about the behaviour of complex simulation models. If the simulation model turns out to produce irregular outcome behaviour, analysts might consider employing more complex metamodels. However, given the long and successful track record that regression metamodels have in the area of computational simulation at large, I expect that in many cases they will provide detailed and accurate insights into the behaviour of complex agent-based models in demographic research.

Acknowledgments The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement no. 312290 for the GENDERBALL project. I thank Jan Van Bavel, Hideko Matsuo, and two anonymous reviewers for helpful comments on earlier versions of this chapter.

References

- Antony, J. (2003). *Design of experiments for engineers and scientists*. Oxford: Butterworth-Heinemann.
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. D. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the United Kingdom with Gaussian process emulators. *Demographic Research*, 29, 729–766.
- Billari, F. C., Ongaro, F., & Prskawetz, A. (2003). Introduction: Agent-based computational demography. In F. C. Billari & A. Prskawetz (Eds.), *Agent-based computational demography: Using simulation to improve our understanding of demographic behaviour* (pp. 1–17). Heidelberg: Physica Verlag.
- Billari, F. C., Prskawetz, A., Aparicio Diaz, B., & Fent, T. (2007). The “Wedding-Ring”: An agent-based marriage model based on social interaction. *Demographic Research*, 17, 59–82.
- Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl. 3), 7280–7287.
- Coutts, S. R., & Yokomizo, H. (2014). Meta-models as a straightforward approach to the sensitivity analysis of complex models. *Population Ecology*, 56(1), 7–19.
- De Hauw, Y., Grow, A., Van Bavel, J. (2015). The shifting gender balance in higher education and assortative mating in Europe. Paper presented at the 2015 Annual Meeting of the Population Association of America, San Diego, USA.
- De Mulder, W., Grow, A., Molenberghs, G., & Verbeke, G. (2015). Application of statistical emulation to an agent-based model: Assortative mating and the reversal of gender inequality in education in Belgium. In H. Friedl & H. Wagner (Eds.), *Proceedings of the 30th international workshop on statistical modelling*, vol. 1 (pp. 139–144). Linz.
- England, P., & McClintock, E. A. (2009). The gendered double standard of aging in US marriage markets. *Population and Development Review*, 35(4), 797–816.
- Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity*, 4(5), 41–60.

- Esteve, A., García-Román, J., & Permanyer, I. (2012). The gender-gap reversal in education and its effect on union formation: The end of hypergamy? *Population and Development Review*, 38(3), 535–546.
- Fox, J. (2008). *Applied regression analysis and generalized linear models* (2nd ed.). Thousand Oaks: Sage.
- Grow, A., & Van Bavel, J. (2015). Assortative mating and the reversal of gender inequality in education in Europe: An agent-based model. *PLoS ONE*, 10(6), e0127806.
- Kalmijn, M. (1994). Assortative mating by cultural and economic occupational status. *American Journal of Sociology*, 100(2), 422–452.
- KC, S., Barakat, B., Goujon, A., Skirbekk, V., Sanderson, W. C., & Lutz, W. (2010). Projection of populations by level of educational attainment, age, and sex for 120 countries for 2005–2050. *Demographic Research*, 22, 383–472.
- Kleijnen, J. P. C. (1992). Sensitivity analysis of simulation experiments: Regression analysis and statistical design. *Mathematics and Computers in Simulation*, 34(3–4), 297–315.
- Kleijnen, J. P. C. (2005). An overview of the design and analysis of simulation experiments for sensitivity analysis. *European Journal of Operational Research*, 164(2), 287–300.
- Kleijnen, J. P. C. (2008). *Design and analysis of simulation experiments*. Tilburg: Springer.
- Kleijnen, J. P. C., & Sargent, R. G. (2000). A methodology for fitting and validating metamodels in simulation. *European Journal of Operational Research*, 120(1), 14–29.
- Kleijnen, J. P. C., van Ham, G., & Rotmans, J. (1992). Techniques for sensitivity analysis of simulation models: A case study of the CO2 greenhouse effect. *Simulation*, 58(6), 410–417.
- Kleijnen, J. P. C., Sanchez, S. M., Lucas, T. W., & Cioppa, T. M. (2005). A user's guide to the brave new world of designing simulation experiments. *INFORMS Journal on Computing*, 17(3), 263–289.
- Law, A. M., & Kelton, D. W. (2000). *Simulation modeling and analysis* (3rd ed.). Boston: McGraw-Hill.
- Lenth, R. V. (2009). Response-surface methods in R, using rsm. *Journal of Statistical Software*, 32(7), 1–17.
- Ligmann-Zielinska, A., Kramer, D. B., Cheruvelil, K. S., & Soranno, P. A. (2014). Using uncertainty and sensitivity analysis in sociological agent-based models to improve their analytical performance and policy relevance. *PLoS ONE*, 9(10), e109779.
- Lorscheid, I., Heine, B.-O., & Meyer, M. (2012). Opening the “Black Box” of simulations: Increased transparency and effective communication through the systematic design of experiments. *Computational and Mathematical Organization Theory*, 18(1), 22–62.
- Lutz, W., Goujon, A., KC, S., & Sanderson, W. C. (2007). Reconstruction of populations by age, sex and level of educational attainment for 120 countries for 1970–2000. *Vienna Yearbook of Population Research*, 5, 193–235.
- Macy, M. W., & Flache, A. (2009). Social dynamics from the bottom up: Agent-based models of social interaction. In P. Hedström & P. Bearman (Eds.), *The oxford handbook of analytical sociology* (pp. 245–268). Oxford: Oxford University Press.
- Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28, 143–166.
- McCarthy, M. A., Burgman, M. A., & Ferson, S. (1995). Sensitivity analysis for models of population viability. *Biological Conservation*, 73(2), 93–100.
- Myers, R. H., & Montgomery, D. C. (1995). *Response surface methodology: Process and product optimization using designed experiments*. New York: Wiley.
- Oakley, J., & O'Hagan, A. (2002). Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*, 89(4), 769–784.
- Oppenheimer, V. K. (1988). A theory of marriage timing. *American Journal of Sociology*, 94(3), 563–591.
- R Core Team. (2014). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing.
- Rao, C. R. (1959). Some problems involving linear hypothesis in multivariate analysis. *Biometrika*, 46(1/2), 49–58.

- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2), 143–186.
- Schwartz, C. R. (2013). Trends and variation in assortative mating: Causes and consequences. *Annual Review of Sociology*, 39, 451–470.
- Silverman, E., Bijak, J., Hilton, J., Cao, V. D., & Noble, J. (2013). When demography met social simulation: A tale of two modelling approaches. *Journal of Artificial Societies and Social Simulation*, 16(4), 9.
- Simão, J., & Todd, P. M. (2002). Modeling mate choice in monogamous mating systems with courtship. *Adaptive Behavior*, 10(2), 113–136.
- Simpson, T. W., Peplinski, J. D., Koch, P. N., & Allen, J. K. (2001). Metamodels for computer-based engineering design: Survey and recommendations. *Engineering with Computers*, 17(2), 129–150.
- Smaldino, P. E., & Schank, J. C. (2011). Human mate choice is a complex system. *Complexity*, 17(5), 11–22.
- Thiele, J. C., Kurth, W., & Grimm, V. (2014). Facilitating parameter estimation and sensitivity analysis of agent-based models: A cookbook using NetLogo and R. *Journal of Artificial Societies and Social Simulation*, 17(3), 11.
- Todd, P. M., Hills, T., & Hendrickson, A. (2013). Modeling reproductive decisions with simple heuristics. *Demographic Research*, 29, 641–662.
- Tunali, S., & Batzman, I. (2000). Dealing with the least squares regression assumptions in simulation metamodeling. *Computers and Industrial Engineering*, 38(2), 307–320.
- Van Bavel, J. (2012). The reversal of gender inequality in education, union formation and fertility in Europe. *Vienna Yearbook of Population Research*, 10, 127–154.

Chapter 8

Design and Analysis of Demographic Simulations

Jason Hilton and Jakub Bijak

8.1 Introduction

As the many novel contributions to this volume show, Agent-Based Models (ABMs) offer exciting possibilities for including explanatory mechanisms, such as behavioural rules governing individual behaviour, in the analysis of demographic phenomena. Knowledge about the abstract statistical individual (Courgeau 2012) derived from empirical data can in this way be augmented by rule-based explanations, giving demography much-needed theoretical foundations (Billari et al. 2003).

As ABMs gain more traction in demography, they will inevitably become more sophisticated, and, as a consequence, more complicated (Grazzini and Richiardi 2013). As demographers explore the possibilities of the methodology, they may attempt to make their agent-based simulations match reality more closely; to model more fields of social life; to pay attention to the effect of institutions and policy; and to enrich their models with more data to attempt to bring them in line with what is observed (Silverman et al. 2011; Squazzoni 2012).

This progress towards greater sophistication in agent-based approaches introduces additional sources of uncertainty to the modelling process, which need

Electronic supplementary material The online version of this chapter (doi: [10.1007/978-3-319-32283-4_8](https://doi.org/10.1007/978-3-319-32283-4_8)) contains supplementary material, which is available to authorized users.

J. Hilton (✉)
University of Southampton, Southampton, UK
e-mail: jdh4g10@soton.ac.uk

J. Bijak
Department of Social Statistics and Demography, University of Southampton, Southampton, UK
e-mail: j.bijak@soton.ac.uk

to be acknowledged (Kennedy and O’Hagan 2001). Furthermore, it creates a need for a more considered approach to the design and analysis of agent-based computer *experiments*. This chapter discusses how existing techniques from the computational experiment literature might be fruitfully applied to overcome some of these difficulties, and specifically focuses on the work of the *Managing Uncertainty in Complex Models* community (MUCM 2011) and how it might be applied to demographic models. In particular, three specific questions are addressed:

1. At what input values should the simulation be run?
2. How does the simulation respond to variations in inputs?
3. How can the simulation be calibrated so that it matches observed quantities?

This chapter will discuss each of these points in turn. It is argued that the first question can be considered one of experimental design, and the use of space-filling Latin hyper-cube sampling provides an efficient default response to this problem. Next, the use of Gaussian Process Emulators is proposed as a way of analysing the behaviour of ABMs. The emulation approach also provides a framework for the calibration of such simulations. Finally some of the limitations of the approach proposed are discussed, and a brief pointer is given to some potential extensions to the basic approach described here that may benefit the analysis of demographic ABMs. This chapter tackles similar issues to those discussed by Grow elsewhere in this volume, although the presented methods are different. Our approach is rooted in the Bayesian statistical tradition, and thus combines the analysis of different sources of uncertainty in a joint probabilistic model, making inferences about the underlying complex population dynamics on that basis.

8.2 Design of Computer Experiments

Epstein and Axtell’s seminal book on ‘Growing Artificial Societies’ (1996) famously considered ABMs as an analogue of physical experiments for social scientists. Modellers, they state, could grow experimental scenarios *in silico*, enabling them to examine the effects of manipulating various inputs. To take this claim seriously, and to maximise what we learn from our quasi-experiments, borrowing from the literature on the design of experiments can be instructive.

Computational experiments differ from physical experiments in several important respects (Santner et al. 2003). Firstly, computational experiments tend to be cheaper, and so can be run more times and at more points. Secondly, the modeller has complete control of the experimental conditions, and results are therefore not subject to unobserved nuisance factors that may cloud inference. Thirdly, greater freedom is possible in the specification of the experiments to be run. For instance, in a simulated environment, one could re-run Galileo’s famous experiment regarding the speed of falling objects under different gravitational conditions, which would be extremely difficult to do in earth-bound physical experiments (Epstein and Axtell 1996).

One consequence of this freedom, however, is that simulations tend to have a greater number of free parameters than their physical equivalents

(Santner et al. 2003). For a given simulation run, values of each of these parameters (inputs) must be set in order to get an output value from the simulation, and each parameter can thus be considered as a dimension across which individual simulation runs differ (Montgomery 2013). For two parameters, we have a two-dimensional input space, and a single simulation run is a point in this space, with a location defined by the values of its two parameters. Our problem is to place points (that is, select combinations of parameters) in this space at which to run our simulation in order to get as much information as possible about the output variable of interest, acknowledging that for continuous inputs we can never run the simulation at *all* points.

The choice of input points, or more correctly, our *experimental design*, depends very much on both the number of experiments we are able to conduct, and our assumptions about how the simulation is expected to behave (Kleijnen 2008). Consider a one-dimensional simulation where the response of the output to changes in the input is consistently linear. For deterministic simulations, only two points are required to work out the slope of this linear response, and thus to have a good estimate of the simulation for any value of the input (ibid).

ABMs, however, cannot be assumed to have simple relationships between inputs and outputs. Because such simulations are by definition caused by the interaction of many autonomous units, the system as a whole can be defined as *complex*. In this context the word has a technical meaning, and complex systems tend to be characterised by tipping points, non-linearities, and other such features (Mitchell 2009). This suggests we must be agnostic about the behaviour of the simulation in question when choosing our design (see Santner et al. 2003, section 5), and often means that a large number of runs are required to get a handle on the behaviour of the simulation over the entire parameter space (Grimm and Railsback 2005).

An important consideration is that ABMs are almost always stochastic, in the sense that repetitions of a simulation run at the same parameter values will give a different outcome due to the use of pseudo-random numbers in driving various elements of the model (Grimm and Railsback 2005). This may mean that multiple simulation runs will need to be conducted at a single parameter combination in order to understand the distribution of outcomes at that point (Law 2007).¹

8.2.1 Factorial Designs

Let us consider an experiment, whether simulated or physical, as a mapping of some input \mathbf{x} to output y . We denote this mapping as a function $f(\cdot)$. Generally \mathbf{x} will be multidimensional, in that there are k parameters to the model $\mathbf{x} = \{x_1, x_2, \dots, x_k\}$

¹In some cases, for instance, when the frequency of rare events are of interest, very large numbers of repetitions may be required to infer about the quantities of interest. Different approaches from those advocated here would likely be required for such problems, one of which might be to apply the analysis and calibration methods discussed in later sections to understand the behaviour of a different, more frequently observed output measure first, simplifying the problem of analysing the rare event.

and we restrict attention here to real inputs, so that $X \in \mathbb{R}^k$. The design problem is to choose a vector of input points D so that we can learn as much as possible about how $f(\cdot)$ responds to \mathbf{x} .

The larger the number of dimensions, the harder this problem becomes. The instinctive response is simply to pick a set of values (levels) for each parameter, and run the model at all combinations of these parameters. This sort of grid-based approach is commonly used in agent-based modelling (e.g. Aparicio Diaz et al. 2011; Fent et al. 2013; Klabunde 2014), and it is a sensible default response for simple models. However, even for relatively small numbers of levels and parameters, this can quickly become prohibitively time consuming, unless the simulation in question is extremely fast. For a simulation with 6 input parameters, $k = 6$, each of which we want to run at 5 levels, we need a total of 5^6 , or 15,625 runs. If our simulation takes only 1 min to run, this would require a total of 260 h of runtime, or some multiple of this number if we wish to repeat observations at each point. Of course, computing power is relatively cheap at present, and multi-processor clusters or cloud computing resources can easily reduce this time to a few hours, or even less. However, as ABMs begin to simulate more agents and involve more complicated decision making, run-times are likely to increase. As a result, it is prudent to consider more efficient experimental designs.

A grid is a particular case of a more general set of designs known as *factorial* designs, which are commonly used for physical experiments (Montgomery 2013), and have also been heavily used in conducting experiments with Operational Research simulations such as queueing models (amongst many other types) (Kleijnen 2008). The nature of the factorial design chosen generally depends upon the expected nature of the response of the simulation outputs to inputs; the most common factorial designs assume that the relationship can be approximated by low order polynomials and possibly two-way interactions. A *full factorial* design is an analogue of the grid design discussed above; for a two-level full factorial design, each factor (input) is considered to have two levels (values), and design points are obtained for every possible combination of levels of the distinct factors, giving 2^k points (Montgomery 2013). A two-level full factorial design assumes linear relationship between variables, and allows for two-way interactions to be identified. If quadratic effects are suspected, *central composite* designs add additional points in the centre and at the extremities of the design space, while *fractional factorial* designs can be used to reduce the number of runs required, effectively by assuming some two-way interactions are equal to zero (ibid). The chapter by Grow elsewhere in this volume describes the use of such designs in the context of demographic ABMs.

8.2.2 Latin Hyper-Cube Sample Designs

The key limitation with grid designs is that when projected or ‘collapsed’ onto one dimension, many design points are replicated, and thus wasted (Urban and Fricker 2010). To put it another way, factorial designs enforce a strong relationship between

the number of dimensions k and the number of runs required n , for a fixed number of levels (Kleijnen 2008). Latin hyper-cube samples avoid these problems by ensuring no two design points share values of any parameter, thus reducing the dependence of n on k (Urban and Fricker 2010). The principle is simple – to create a Latin hyper-cube sample, divide the input space equally into g sections along each axis, so that there are g^k cells in total. Then, choose g of these cells such that there is only one cell in each section (or column) for every axis. To complete the process pick a point randomly within each chosen cell, resulting in a sample of size $n = g$ (Santner et al. 2003).

Latin hyper-cube samples are not guaranteed to fill the entire parameter space, so some further criteria are needed to ensure that all parameter combinations are explored adequately (Santner et al. 2003). Generating several candidate samples, and picking the one with the highest minimum distance between points will in general suffice (O’Hagan 2006). The R package `lhs` has a number of functions for producing such samples for arbitrary dimensions very easily.² These will then need to be scaled up from the existing $[0, 1]$ range to reflect the input ranges required by any particular simulation.

The key advantage of Latin hyper-cube sample designs is the scaling in high dimensions. For the example above, rather than requiring thousands of simulations to explore a six-dimensional space, samples of around 60 points may be sufficient. Loeppky et al. (2009), for example, investigate the relationship between sample size and meta-model predictive adequacy, and find that the established rule of thumb of $n = 10k$ is generally reasonable, but that this number will vary dependent on whether all or only some of the inputs strongly affect the output. Repetition of simulations at individual design points may also be desirable in order to account for stochasticity in simulation outputs. These issues are discussed in Kleijnen (2008), Ankenman et al. (2010), and Boukouvalas (2010). Other advantages of the Latin hyper-cube sample are discussed in Santner et al. (2003).

8.3 Analysis of Computer Experiments

Once a design has been settled on, and the simulation has been run at the design points, the next step is the analysis of the simulation results. In high dimensions, understanding the relationships between inputs and outputs simply from the raw results is often difficult. The method of analysis chosen to analyse the simulation should reflect prior expectations about its behaviour, and is closely related to the choice of design; some methods require particular designs, and work better for simulators with certain properties.

²Other methods of obtaining Latin hyper-cube samples are available. For instance, @Risk (www.palisade.com/risk) is an add-on for Excel which provides this functionality, as does the statistics and machine learning tool-kit (uk.mathworks.com/help/stats/lhsdesign.html) of the Matlab mathematical programming software. However, both of these are proprietary packages and not freely available.

ABMs, as suggested above, are generally *complex*, and thus may be highly non-linear. Furthermore, given that ABMs are often built in order to explain or explore some particular poorly-understood phenomena, we may not have any clear idea about how we expect the simulation to behave. This suggests that analysis methods (and designs) that make fewer assumptions and are able to capture many different types of relationships between inputs and outputs are preferable (Santner et al. 2003).

From a statistical point of view, modellers must also ensure that they account for the various sources of uncertainty inherent in the analysis of simulations (Kennedy and O'Hagan 2001). Sources of such uncertainty include:

- **Uncertainty due to simulation stochasticity.** This occurs when running the code twice at the same parameters gives different results, because of the use of random number generation in the simulator itself. This means even if we have the result of one trial at a point, we cannot predict with certainty the value of another such trial. Random number generators are used to represent *aleatory* uncertainty in the real world phenomena; that is, uncertainty due to inherent randomness (O'Hagan 2006).
- **Uncertainty about the output at new points.** In a continuous parameter space, we can never run the simulation at every conceivable point. Instead, we must estimate at points we have yet to run, which we do with some error.
- **Input uncertainty.** In many cases, we do not know what the 'correct' value of any given parameter is. We may, however, have a reasonable range or probability distribution that characterises our beliefs as to where the 'true' value lies. This uncertainty about inputs clouds our knowledge about outputs.
- **Model discrepancy.** The model is unlikely to be a perfect representation of reality. The mechanisms simulated will differ from what takes place in the world in appreciable but uncertain ways. Thus, our lack of knowledge about the ways and extent to which our model is wrong is another source of uncertainty.
- **Measurement error.** Comparing simulated results to real results may add an additional source of uncertainty, as real world quantities are subject to errors in measurements (Kennedy and O'Hagan 2001).

The last four sources of uncertainty are largely *epistemic*, in that the uncertainty is the result of our lack of knowledge about the quantities of interest (although the last one is also partly *aleatory*) (O'Hagan and Oakley 2004). Failing to take these sources of uncertainty, whether aleatory or epistemic, into account can lead to faulty inferences (O'Hagan 2006). This can be particularly problematic if policy advice is the goal of the simulation; representing uncertainty about the phenomena in question is vital if potential risks are to be mitigated (Bijak 2011).

There are a number of ways in which the analysis of simulation outputs can be approached. Firstly, a 'brute-force' Monte-Carlo approach can be considered. Sampling from distributions representing the above sources of uncertainty many times and obtaining simulation results for each sample would allow for a coherent accounting. However, this requires a *large* number of replications (O'Hagan 2006), and so has computation time implications, and thus other approaches are preferred.

Several other approaches to the problem have been proposed. Bayesian melding (Poole and Raftery 2000) involves reconciling prior knowledge about outputs and inputs with observations and simulation outputs. This technique is robust in its accounting for most sources of uncertainty, and incorporating prior knowledge, but also requires many replications to build up posterior distributions. Ševčíková et al. (2007) used this approach in a stochastic urban simulation that modelled household behaviour with interesting results.

Finally, our preferred approach is the use statistical *emulators* to approximate the simulation, so called because they emulate the behaviour of the simulation: they are *meta-models* of the underlying models. Meta-models can take a variety of forms, from regression models involving low-order polynomials, to more complex regression tree approaches, to neural networks (Santner et al. 2003; Kleijnen 2008). The chapter by Grow elsewhere in this volume adopts a second-order polynomial meta-model in order to examine the sensitivity of an ABM of marriage markets to changes in parameters. However, given our desire to make few assumptions about the functional form of the relationship between simulator inputs and outputs, Gaussian Process Emulators (often called *kriging* models) are our preferred approach (Kennedy and O’Hagan 2001). The need to incorporate the many different sources of uncertainty discussed above suggests a Bayesian framework, in which both epistemic and aleatory uncertainties can naturally be represented as distributions, and included in output predictions through Bayes’ rule (O’Hagan and Oakley 2004; Oakley and O’Hagan 2004). Equivalent frequentist approaches do exist; see, for example Kleijnen (2008), Forrester et al. (2008), and Ankenman et al. (2010).

Gaussian Process Emulators (and meta-models in general) introduce another layer of uncertainty as they only provide an approximation of the simulator output, but are flexible and less computationally expensive than obtaining results at all points of interest, and also provide other benefits in terms of ease of analysis, as discussed below. However, they do require the modeller to make two main assumptions about the simulator. Firstly, it is assumed that the relationship between inputs and output is *smooth* to some degree, although the extent of this smoothness is estimated from the data (O’Hagan 2006). Secondly, the process is assumed to be second-order stationary over the parameter space, effectively meaning that the degree of smoothness remains constant across the parameter space (Santner et al. 2003). These assumptions may not always hold for ABMs, but it is argued that they are less restrictive assumptions than are required for many other meta-models, and further, when they fail, Gaussian Process Emulators can still give useful information about the general behaviour of the simulation. An introduction to Gaussian Process Emulators is provided below, and detailed information can be found on the website of the research community Managing Uncertainty in Complex Models (MUCM 2011).

8.4 Gaussian Process Emulators: A Primer

Gaussian processes are extremely flexible statistical tools as they make few assumptions about the form of the function they are used to represent. Given what we have said about the complex and non-linear nature of agent-based simulations,

this makes them well suited to our purpose. The underlying premise is that outputs near to each other in the parameter space are also nearby in the output space. A Gaussian process represents this idea by insisting that outputs at any collection of input points are *joint multivariate normal* (Kennedy and O’Hagan 2001). A more formal treatment follows.

8.4.1 General Premises

Let $f(\cdot)$ be the base computational model of interest. We focus on a vector of k inputs to this simulation, $\mathbf{x} \in X \subset \mathbb{R}^k$, and a single output, $y \in Y \subset \mathbb{R}$, such that $y = f(\mathbf{x})$. X does not have to exhaust the whole parameter space, but rather relate to those inputs which are considered important from the point of view of the output studied. Following Oakley and O’Hagan (2002, p. 771) and Kennedy (2004, p. 2), we define a Gaussian Process Emulator, conditionally on its parameters, as a multivariate Normal distribution for N realisations of f , $y_1 = f(\mathbf{x}_1), \dots, y_N = f(\mathbf{x}_N)$, denoted jointly as f (ibid):

$$f(\cdot) \mid \beta, \sigma, \omega \sim N[m(\cdot), \sigma^2 c(\cdot, \cdot)] \quad (8.1)$$

A number of options are possible for the mean of the process. Often, it is chosen to be a constant so that $m(\cdot) = \beta_0$. In other contexts, it is modelled through a vector linear regression function of \mathbf{x} , $h(\mathbf{x})$, with coefficients β , such that for every output $f(\mathbf{x})$, $m(\cdot) = h(\cdot)^T \beta$. Throughout this chapter, we use the latter, and choose $h(\mathbf{x})$ to be a simple function of the inputs, so that $m(\mathbf{x}) = \beta_0 + \beta_1 x_1 + \dots + \beta_k x_k$, with k the number of input dimensions.³ The number of basis functions (in this context, the number of additive terms) in the mean function is denoted by q , equalling one in the case of a constant mean function and $k + 1$ in the simple linear regression case. The covariance matrix is determined by correlation function $c(\cdot, \cdot)$, which determines how quickly nearby points become uncorrelated, and the variance parameter σ , which determines the extent of deviation from the mean function. Several forms are possible for the function $c(\cdot, \cdot)$, the most common of which is the squared exponential or Gaussian function (Rasmussen and Williams 2006):

$$c(\mathbf{x}_i, \mathbf{x}_j) = \exp \left\{ - \sum_{l=1}^k \omega_l (x_{il} - x_{jl})^2 \right\} \quad (8.2)$$

The Ω parameter vector $\Omega = \{\omega_1, \dots, \omega_k\}$ is composed of roughness parameters (or ‘correlation lengths’), which indicate how strongly the emulator

³Setting the mean function to a constant is often called ‘Ordinary Kriging’, while using a regression model is referred to as ‘Universal Kriging’ (Kleijnen 2008).

responds to particular inputs (Kennedy and O’Hagan 2001, pp. 432–433; O’Hagan 2006). In order to incorporate any inherent simulator stochasticity into the emulator, an additional variance term τ^2 (referred to as a nugget) can be added to the covariance function $v(\mathbf{x}_i, \mathbf{x}_j) = \sigma^2 c(\mathbf{x}_i, \mathbf{x}_j)$ when $i = j$, so that (MUCM 2011):

$$v(\mathbf{x}_i, \mathbf{x}_j) = \sigma^2 \exp \left\{ - \sum_{l=1}^k \omega_l (x_{il} - x_{jl})^2 \right\} + I_{i=j} \tau^2 \quad (8.3)$$

where $I_{i=j}$ is an indicator variable that equals 1 if $i = j$ and 0 otherwise.

8.4.2 Estimation

In order to estimate the parameters of the emulator, a set of simulation data $f(D) = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]$ (the *training set*) is required for n experimental points $D = \mathbf{x}_1, \dots, \mathbf{x}_n$, where $\mathbf{x} \in D \subset X$ (Kennedy 2004, p. 2). Conditional on this training data and the values of the Gaussian process parameters β, σ^2 , and Ω , the distribution of simulator outputs at new points \mathbf{x}' is joint multivariate normal. Taking non-informative priors on β and σ^2 such that $p(\beta, \sigma^2) \propto \sigma^{-2}$, it is possible to marginalise β and σ^2 , obtaining a multivariate t-distribution for outputs at new points, and the following likelihood for the roughness parameters (Andrianakis and Challenor 2011):

$$p(f(D)|\Omega) \propto |A|^{-1/2} |H^T A^{-1} H|^{-1} (\hat{\sigma}^2)^{\frac{n-q}{2}} \quad (8.4)$$

where H is the matrix of basis function generated by $h(D)$ and A is the correlation matrix for the training set, defined by $c(D, D)$. The values that maximise this likelihood can then be found and can be used as ‘plug-in’ posterior mode estimates of the values for Ω (Kennedy and O’Hagan 2001; Oakley 1999).

Given Ω , conditional expressions follow for estimates of β and σ^2 (Andrianakis and Challenor 2011):

$$\hat{\beta} = (H^T A^{-1} H)^{-1} H^T A^{-1} f(D) \quad (8.5)$$

$$\hat{\sigma}^2 = \frac{1}{n - q - 2} (f(D) - H \hat{\beta})^T A^{-1} (f(D) - H \hat{\beta}) \quad (8.6)$$

Although this approach neglects the uncertainty around values of Ω , it is suggested that this uncertainty is not significant compared to that for other quantities. Full details and examples can be found in Andrianakis and Challenor (2011) and on the MUCM website (MUCM 2011).

Alternatively, full MCMC sampling approaches can be used to estimate the posterior distributions of the unknown hyper-parameters (Gramacy 2005). Direct maximisation of the multivariate normal likelihood for all parameters is also often

used, particularly in machine learning contexts (Rasmussen and Williams 2006; Boukouvalas 2010). To estimate the nugget parameter τ^2 , following Roustant et al. (2012, Appendix A2) we introduce an additional parameter α which determines the proportion of the total variance $v^2 = \sigma^2 + \tau^2$ that is explained by the inputs. The covariance function thus becomes (ibid):

$$v(\mathbf{x}_i, \mathbf{x}_j) = \hat{v}^2 \left\{ \alpha \exp \left\{ - \sum_{l=1}^k \omega_l (x_{il} - x_{jl})^2 \right\} + (1 - \alpha) I_{i=j} \right\} \quad (8.7)$$

where, as before, $I_{i=j}$ is an indicator variable that equals 1 if $i = j$ and 0 otherwise. Note that now $\alpha v^2 = \sigma^2$, while $(1 - \alpha)v^2 = \tau^2$. The proportion α can be estimated by including it in the set of parameters estimated by maximising the distribution in Eq. 8.4.

Repeated runs at each design point can help gain better estimates of α , although the potential for doing so is limited by a desire to minimise the size of the correlation matrix A , and thus computation time. Repeated points are treated in exactly the same way as any others in the design; correlations between repeats will take the maximum possible value of $1 - \alpha$ (with a solitary run's correlation with itself being 1 by definition). Using only single points can lead to difficulty in disentangling the stochastic variance and input-related variance, and multi-modal likelihoods can ensue (Andrianakis and Challenor 2012). However, the validation procedures described below can help choose between competing modes. Alternatively, the approaches of Ankenman et al. (2010) and Boukouvalas (2010) provide a more robust method for including stochastic variance in emulators by fitting Gaussian processes to the mean and variance separately, estimating these moments at each point from a repeated sample.

8.4.3 Predicting New Quantities

One immediate advantage of the Gaussian process approach is that once the parameters are estimated and the posterior distribution of the function f is obtained, new estimates of simulator outputs are very easy to obtain, a particular advantage if the simulation is slow to run. As discussed, the marginalisation of the σ^2 and β parameters mean that the predictive distribution at any collection of points is a multivariate T distribution with $n - q$ degrees of freedom. Conditional on the training sample and the hyper-parameter estimates, the posterior mean of this distribution for simulator outputs at the new point \mathbf{x} is just the result of matrix multiplication (Oakley 1999):

$$m^*(\mathbf{x}) = h(\mathbf{x})\beta + t(\mathbf{x})^T A^{-1} \mathbf{e} \quad (8.8)$$

where $m^*(\mathbf{x})$ denotes the posterior mean function; $t(\mathbf{x})$ the correlation between the new point \mathbf{x} and the elements of the training set D ; and \mathbf{e} is the difference between simulator outputs $f(D)$ and the mean prediction $h(\mathbf{x})^T \beta$. This allows the analyst or

modeller to get a complete picture of the parameter space very easily. Furthermore, the uncertainty induced by using the emulator as an estimate of the simulator is readily evaluated as well, so this source of uncertainty is not lost (Oakley 1999):

$$v^*(\mathbf{x}_i, \mathbf{x}_j) = \hat{\sigma}^2 \{ c(\mathbf{x}_i, \mathbf{x}_j) - t(\mathbf{x}_i)^T A^{-1} t(\mathbf{x}_j) \\ + (h(\mathbf{x}_i)^T - t(\mathbf{x}_i)^T A^{-1} H)(H^T A^{-1} H)^{-1} (h(\mathbf{x}_j)^T - t(\mathbf{x}_j)^T A^{-1} H)^T \} \quad (8.9)$$

where $v^*(\mathbf{x}_i, \mathbf{x}_j)$ is the posterior covariance between points \mathbf{x}_i and \mathbf{x}_j and H is the matrix of basis function generated by $h(D)$. Whilst the variance function above is a complicated-looking function, it is simple to evaluate as again it only requires linear algebra.

8.4.4 Uncertainty Analysis

The emulator, once built, can also be used for an uncertainty analysis, which looks at how much uncertainty in the output is being induced by the set of inputs X (Oakley 1999). This is particularly important in predictive, real world applications of ABMs where we might wish our simulation to inform decision making. Some model inputs might be based on noisy estimates from real world data, others may be given priors that reflect our subjective assessment of their probable values (Werker and Brenner 2004). In either case, we would like to quantify this lack of knowledge by treating these inputs as random variables with some assumed probability distributions. The uncertainty analysis propagates this uncertainty through the emulator, and takes it into account in providing estimates of the simulator's mean and variance (Oakley 1999).

An orthodox Monte Carlo approach to this problem would be to repeatedly sample from the input distributions, run the simulator at each point, and examine the resulting distribution on the output (Saltelli et al. 2004). However, this is computationally expensive, as many simulation runs are required to get a good approximation of the output distribution (O'Hagan 2006). An alternative approach is to use the fitted emulator to conduct the Monte Carlo analysis instead, as it is many orders of magnitude faster in generating predictions (MUCM 2011). Even better, however, is that for inputs with normally distributed priors and squared exponential covariance functions, the work of Haylock (1997) and Oakley and O'Hagan (2004) provides analytical expressions for the relevant integrals of the emulator output over the input uncertainty, allowing easy computation of the posterior expectation of the simulation output, the variance of this estimate, and the expectation of the simulator variance.⁴

⁴Barton et al. (2014) and Xie et al. (2014) also suggest approaches whereby input uncertainty can be propagated using meta-models in order to obtain output distributions.

To summarise, assuming that $G(\cdot)$ is the distribution function of the random input variables X , then the mean $E[f(X)]$ and variance $V = \text{Var}[f(X)]$ of the distribution of output $f(\mathbf{x})$ taking into account the input uncertainty are (MUCM 2011):

$$\begin{aligned} E[f(X)] &= \int_X f(\mathbf{x}) dG(\mathbf{x}) \\ \text{Var}[f(X)] &= E[f(X)^2] - \{E[f(X)]\}^2 \\ E[f(X)^2] &= \int_X f(\mathbf{x})^2 dG(\mathbf{x}) \end{aligned} \quad (8.10)$$

The analytical expressions for these integrals are given in MUCM (2011).

8.4.5 Sensitivity Analysis

The purpose of a sensitivity analysis is to understand how the model output responds to changes in inputs. Historically, these have been conducted by assessing the change in output for small changes in input at some specified point of interest (a local sensitivity analysis) (Saltelli et al. 2004). The partial derivatives of the function in question approximated at this point are often used for this purpose (ibid). This is problematic in the case where the whole input space is potentially of interest, particularly if the model is non-linear, in which case the derivatives at one point are not representative of the rest of the input space (Saltelli et al. 2008). Thus global measures of model sensitivity that summarise the behaviour of the outputs across the input space are to be preferred (ibid).

Although there are various methods for conducting a sensitivity analysis (Saltelli et al. 2004, 2008), variance-based methods provide a way to utilise emulators to maximise efficiency (Oakley and O'Hagan 2004). Sensitivity analysis is defined by Saltelli et al. (2004, p. 45) as “the study of how uncertainty in the output of a model (numerical or otherwise) can be apportioned to different sources of uncertainty within the model input”. This definition provides a link to the uncertainty analysis described in the previous section: a method for finding the expectation of the total uncertainty due to inputs in our model – the variance $\text{Var}\{f(\mathbf{x})\}$. The sensitivity analysis methods described below aim to partition this uncertainty between inputs (Oakley and O'Hagan 2004).

The principal variance-based measure used in this chapter is the *sensitivity variance* V_w , where w here identifies the input or collection of inputs which we are interested in apportioning variance to MUCM (2011). This measures the reduction in overall variance that would result from knowing the value(s) of x_w (ibid):

$$\begin{aligned} V_w &= \text{Var}\{f(\mathbf{x})\} - E\{\text{Var}(f(\mathbf{x})|\mathbf{x}_w)\} \\ &= \text{Var}\{E(f(\mathbf{x})|\mathbf{x}_w)\} \\ &= E\{E(f(\mathbf{x})|\mathbf{x}_w)^2\} - E\{f(\mathbf{x})\}^2 \end{aligned} \quad (8.11)$$

Dividing by the total variance induced by the inputs $V = \text{Var}\{f(\mathbf{x})\}$, obtained from the uncertainty analysis, we get a scale-invariant *sensitivity index* $S_w = V_w/V$ (Oakley and O’Hagan 2004). When the set of inputs w contains only one input, we obtain the main effect for that input. Reductions obtained for combinations of inputs reflect the joint (interaction) effects (Oakley and O’Hagan 2004). All of these sensitivity measures can be estimated through Monte-Carlo or Quasi-Monte-Carlo sampling of the output (Sobol 2001), but the job is made considerably easier by the use of emulators. The expectations of the conditional variances for input subsets are given in closed form given the emulator parameters in MUCM (2011), again assuming normal priors on the inputs and squared exponential correlation functions.

Sensitivity analysis can be an extremely useful tool in analysing ABMs and assessing their robustness. Firstly, just knowing which inputs are important in a simulation and which are not is helpful in understanding the processes involved – if a simulation is not sensitive to a given parameter, then that parameter can safely be ignored (Grimm and Railsback 2005). Secondly, given that ABMs may require assumptions regarding the values of some parameters due to lack of data, finding that outputs are not that sensitive to changes in such parameters helps justify these modelling choices (ibid). Thirdly, understanding which inputs are contributing most to uncertainty helps target where to gather more information in order to increase the precision with which we can estimate outputs of interest (Oakley and O’Hagan 2004).

8.5 Worked Example

Building on the work conducted in Bijak et al. (2013), we now present a brief example of the use of emulators to examine the behaviour of a simple ABM of partnership formation. The model in question is a re-implementation of the Wedding Ring model of Billari et al. (2007), with the addition of demographic data for the UK. In particular, we introduce fertility and mortality data from 1950–2011, together with a starting population taken from UK census data (see Bijak et al. 2013, for data sources). The model itself aims to show how aggregate age-at-marriage patterns can be built up from the effect of social pressure on individual partner search intensities (Billari et al. 2007).

8.5.1 Model Description

The focus of this section is on explaining the use of the emulator, but a brief description of the model follows in order to aid understanding. A fuller exposition can be found in Bijak et al. (2013), as well as in the original paper by Billari

et al. (2007), and the model code is available at <https://www.openabm.org/model/3549/version/2/view>. Individuals within the simulation reside on a ring and have a number of ‘relevant others’ who form their social network. The proportion of these others who are married enters a function that determines the radius within which an individual searches for partners. The sigmoid shape of this function is controlled by two parameters a and b .⁵ An additional parameter ‘spatial distance’, or sd controls the distance from within which individuals can draw social network members.

We analyse here a single output quantity, the average proportion of agents married over the course of the simulation. The results presented here differ slightly from those presented in previous work, as a larger starting population is used, and a Latin hyper-cube sample rather than a grid sampling design defined the set of input points. As the output data is a proportion, transformations might have been considered to ensure the predicted values remain bounded between $[0, 1]$. However, given that the output data does not approach either bound, the data was left untransformed (Gelman et al. 2014).

8.5.2 *Input Design*

Firstly, a training set was obtained by generating a Latin hyper-cube sample of 200 design points, each consisting of three values, one for each simulation input. Design points were not repeated in this case, although doing so may improve the estimation of the α parameter. An additional 50 points were obtained for the purposes of validating the emulator. Following the recommendations in MUCM (2011) and the discussion in Challenor (2013), these consisted of 25 additional space filling points, chosen to maximise distance from the existing points, and 25 points relatively close to the original sample. Such choices increase the information gained from the validation sample, as they better test the estimated values of both the correlation and variance parameters (Challenor 2013). The simulation was then run at all of these points, obtaining 200 training set input and output pairs, and a further 50 validation pairs. Note that the Latin hyper-cube sample is generated to lie between $[0, 1]$, and so must be scaled for purposes of input to the simulation. First, a range of possible input values must be specified for each parameter, representing our best guess (prior knowledge) of what the most extreme reasonable values for these parameter might be. Then the following transformation is applied to each Latin hyper-cube sample point to get to the required scale:

$$b_i = x_i(\text{high}_i - \text{low}_i) + \text{low}_i \quad (8.12)$$

⁵In Bijak et al. (2013) and Billari et al. (2007), the respective parameters were α and β , however, to avoid confusion with the emulator mean coefficients and correlation parameters, a and b are used here.

Table 8.1 Parameter estimates

Parameter	Estimates			
$\beta = (\beta_0, \beta_a, \beta_b, \beta_{sd})$	(0.707	-0.322	-0.231	0.005)
$\Omega = (\omega_a, \omega_b, \omega_{sd})$	(31.090	46.153	0.368)	
σ^2	0.00436			
τ^2	0.00018			

where b_i is the i th simulation input; x_i the corresponding Latin hyper-cube sample input in the range $[0, 1]$; and $high_i$ and low_i represent the relevant range endpoints for that input.

8.5.3 Emulator Fit

A Gaussian Process Emulator was subsequently fitted to the training data. Note that the original $[0, 1]$ scale input design was used in fitting the emulator, in order to ensure the roughness parameters could be estimated accurately and compared easily. The difference in scale must be taken into account later when interpreting the parameters. The R Statistical Computing Language was used for all estimation (R Development Core Team 2015), and the code used is also provided on the website of this book.⁶ To estimate the emulator hyper-parameters, the mode of the joint marginal likelihood of the roughness parameters and the hyper-parameter α (Eq. 8.4) was first found numerically using the built-in R function `optim`. Several starting points were trialled to avoid a local maximum being chosen. Values of β , σ^2 and τ^2 follow given these hyper-parameters, and the full emulator is obtained. The values of the fitted parameters are given in Table 8.1. The four β parameters refer respectively to the intercept and the coefficients of the mean function for each of the input dimensions. The three ω parameters refer to the roughness of the Gaussian process across each dimension; the high values for the a and b parameters indicate that the output becomes uncorrelated (changes) quickly for small changes in the inputs for those parameters.

8.5.4 Validation

Before we can be confident that our emulator accurately represents our simulator, we should attempt to check its predictions against the validation dataset. Bastos and O'Hagan (2009) propose several metrics to assess emulator validity, two of

⁶Hankin (2005) and Roustant et al. (2012) have produced R-based toolkits to fit Gaussian processes that have influenced the code produced for this chapter. The former only deals with deterministic simulations, however.




Fig. 8.1 Plot of standardised residuals at validation points

which are reported below. Emulator predictions and variances are obtained for the validation input points, and standardised residuals calculated by dividing the difference between the mean emulator prediction and the observed simulation value by the standard deviation predicted by the emulator: these are displayed in Fig. 8.1.

The standardised residuals appear to be relatively reasonable; most lie within the range $[-2, 2]$. The two that perform less well lie on the ‘phase transition’, which is often difficult to accurately represent. Metrics such as standardised residuals that consider each error independently do not take into account the correlation between residuals implied by the Gaussian covariance structure. This can be problematic if two output points which the emulator considers to be highly correlated have errors in opposite directions; individually the errors may appear OK, but when considered together they represent a mismatch between the emulator and the simulator. The Mahalanobis distance can better represent emulator validity taking this into account, and is calculated through the formula (Bastos and O’Hagan 2009):

$$MD = (y_{cv} - E(f(X_{cv})))^T (V(f(X_{cv})))^{-1} (y_{cv} - E(f(X_{cv}))) \tag{8.13}$$

where y_{cv} indicates the outputs for the m validation points, and $E(f(X_{cv}))$ and $V(f(X_{cv}))$ represent the emulator mean and variance estimates at these points, calculated from Eqs. 8.8 and 8.9 respectively. This value, multiplied by $(n - q)/(m(n - q - 2))$, can be compared to the quantiles of an F-distribution with $m = 50$ and $n - q = 200 - 4 = 196$ degrees of freedom (Bastos and O’Hagan 2009). Small values indicate under-confident predictions, in that the predictive distributions are too wide given the actual differences between simulator and emulator, while high values indicate the opposite. For this emulator, the 95% interval for the relevant scaled F-distribution are $[30.9, 74.9]$, and the calculated Mahalanobis distance is 38.6, suggesting the emulator is reasonably accurate in quantifying its




Fig. 8.2 Plot of output by a with 90 % predictive interval. Parameters b and sd held constant at the centre of their ranges

uncertainty about unknown points. In practice, high values for Mahalanobis distance are sometimes seen, suggesting a poorly fitting emulator that is overconfident about its ability to represent the simulator. In such cases, outputs at more design points could be collected in order to attempt to obtain a better fit (Bastos and O’Hagan 2009).

8.5.5 Predictions, Uncertainty, and Sensitivity Analyses

Using the fitted emulator, predictions are obtained for a range of values for the first two parameters, using the formulas in Eqs. 8.8 and 8.9. These predictions are displayed in Figs. 8.2 and 8.3 with the corresponding 90 % predictive intervals resulting from uncertainty due to simulation stochasticity and emulator uncertainty. The bivariate predictions are displayed in Fig. 8.4 – the fitted emulator allows many predictions to be generated easily for such plots. Looking at the shape of the function, it is suspected that quite high-order polynomials would be needed to approximate it to a reasonable degree of accuracy, in part justifying the decision to use Gaussian Process Emulators rather than simpler meta-models.

To conduct uncertainty and sensitivity analyses, assumptions about the distribution of the inputs must be made. For convenience, we assume normal distributions around the midpoint of the input ranges, with variances chosen to assign positive probability over the input range but only small probabilities beyond this, so that




Fig. 8.3 Plot of output by b with 90% predictive interval. Parameters a and sd held constant at the centre of their ranges




Fig. 8.4 Plot of output by a and b . Parameter sd held constant at the centre point of its range

Table 8.2 Sensitivity analysis

	a	b	sd
Main effect	0.357	0.461	0.001
Interaction with a		0.176	0.001
Interaction with b	0.176		0.002
Interaction with sd	0.001	0.002	

$x_i \sim N(0.5, 0.02)$. In other applications, these distributions could reflect substantive prior knowledge. In such a case, the uncertainty analysis allows us to infer how input uncertainty feeds through to uncertainty about outputs. Given the chosen distributions, then, the predicted mean simulator output was 0.33508, and the variance of this estimator was close to 0. The expectation of the overall simulator variance was 0.00874. Similarly, a sensitivity analysis is conducted using the methods discussed in Sect. 8.4.5, to examine how sensitive the simulator is to changes in the various inputs, given the probability distributions of these inputs. The findings are summarised in Table 8.2.

As previously reported in Bijak et al. (2013), the parameters a and b controlling the way marriage search intensity responds to social pressure are most significant in causing changes in outputs. The numbers in the table can be interpreted as proportions of total output variance (excluding the stochastic variance associated with the nugget) explained by each input or combination of inputs. The first row refers to variance associated with each input in isolation, while the subsequent rows refer to interactions. The figures above the diagonal may not sum to one, as a small amount of residual variance is found in the three way interaction, and not reported.

8.6 Extensions

Emulators are most useful in more complicated situations than the simple example described above. Most notably, emulators can assist with problems with higher-dimensional input spaces, and when simulations take a long time to run. Furthermore, the basic framework sketched above can be extended in a number of ways.

8.6.1 Multidimensional Output

Only single-output emulators have been discussed above, but it is possible to extend emulation output into multiple dimensions as well (Kleijnen and Mehdad 2014). One option is simply to assume independence between simulator outputs, and construct separate emulators for different outputs (Vernon et al. 2010), but it is also possible to include correlation structures between different outputs (MUCM

2011). However, an experiment by Kleijnen and Mehdad (2014) suggests that the extra difficulty of constructing multivariate emulators does not have a pay-off in terms of better output predictions. On the contrary, in fact, they find that multivariate emulators tend to perform less well than multiple univariate equivalents.

Alternatively, it is possible to collapse a series of output values of the same type into one single output with an additional input dimension. For instance, multiple population size outputs over time can be reduced to one output by considering time as an additional input dimension. However, this can be problematic, as by doing so assumptions are made about the correlation structure across time that may be inappropriate (MUCM 2011, page labelled AltMultipleOutputsApproach).

8.6.2 Probabilistic Calibration

Another promising application of emulators for demographic ABM is in the area of calibration of simulation parameters. Two techniques for doing this are described below, although it should be noted that calibrating a model does not guarantee that it is correct or accurate (Oreskes et al. 1994). Rather, it merely makes statements about likely values of inputs given the structure of the simulation. The structure of the simulation itself still needs to be justified, either by recourse to theory or by micro-level validation (or, if possible, both) (Rossiter et al. 2010). Full details on the estimation are rather involved, and are not described here. Instead, the aim of the following is to make the reader aware of what can be achieved and the possible utility of such methods for demographic ABMs.

The idea behind calibration is that we can learn about ‘true’ values of unobserved inputs by examining what values of these inputs result in simulation outputs that match observations (Kennedy and O’Hagan 2001). This requires several additional steps. Firstly, input parameters must be divided into two groups. The first are termed ‘control’ or ‘location’ parameters, and these are known for every empirical data point we collect. Examples of such parameters are physical coordinates, which are generally known for every empirical measurement taken. In a demographic context, age and time might be considered location parameters, but care must be taken when modelling time in this way, as the Gaussian correlation structure may not be appropriate in this context (MUCM 2011; Rasmussen and Williams 2006).

The other set of parameters are those to be calibrated. These are not observed in reality, but are assumed to have true values for which we would like obtain a probability distribution. For ABMs in demography, such parameters might govern agents’ decision-making processes: for example, one parameter might control aversion to risk.

Calibration proceeds by relating simulation outputs, as approximated by the emulator, to empirical values by means of a calibration equation (Kennedy and O’Hagan 2001):

$$z(\mathbf{x}_{loc}) = f(\mathbf{x}_{loc}, \theta) + d(\mathbf{x}_{loc}) + \varepsilon(\mathbf{x}_{loc}) \quad (8.14)$$

where \mathbf{x}_{loc} represents a point in the location parameter space, $z(\mathbf{x})$ denotes the observation of the process being simulated at \mathbf{x} , $f(\mathbf{x})$ indicates the simulator output, ε the observation error, and θ the ‘true values’ of the calibration parameters. The other element of the calibration equation is the model discrepancy term $d(\mathbf{x}_{loc})$. This represents the mismatch between the simulator and reality *given that the simulator is run at the ‘true’ values of the calibration parameters*. This captures the idea that the simulator is a simplification of reality and may not match observed values exactly even if it were fully calibrated. This discrepancy function over the location inputs is often modelled as another Gaussian process, priors for the parameters of which must be elicited from the relevant modellers and domain specialists (Oakley 2002).

Given the above framework, the vector of simulation outputs and observations can be modelled as a function of the emulator and discrepancy function, and MCMC methods can be employed to obtain a probability distribution for the ‘true values’ of the calibration parameters (Kennedy and O’Hagan 2001; Qian and Wu 2008; MUCM 2011). This further allows calibrated prediction of reality at new points, taking into account the estimated distribution for the calibration inputs (Kennedy and O’Hagan 2001).

8.6.3 History Matching

A related and slightly simpler approach to restricting the range of input parameters is *history matching* (Vernon et al. 2010). This technique makes use of an ‘implausibility’ metric that gives values to input points that reflects how likely these points are to have generated the observed empirical output, given our uncertainty about simulation outputs, model discrepancy, and measurement error (ibid). A fitted emulator is used to calculate this quantity for a large range of possible calibration parameter values. Formally, the implausibility is defined (ibid) as

$$I(\mathbf{x}) = \frac{(z(\mathbf{x}) - E(f(\mathbf{x})))^2}{\text{Var}(f(\mathbf{x})) + \text{Var}(d(\mathbf{x})) + \text{Var}(\varepsilon(\mathbf{x}))} \quad (8.15)$$

Next, any values that fall beyond a reasonable cut-off point (Vernon et al. (2010) suggest values of $I(\mathbf{x}) > 3$), are rejected as implausible. This generally greatly reduces the ‘non-implausible’ area of the input space. An additional ‘wave’ of simulation runs from the reduced space can then be taken, and a new emulator built. These steps can be repeated until a plausible subset of the input space is identified. This process of *iterative refocusing* can act to calibrate the simulation, although unlike in the previous step, a distribution over the calibration parameters is not obtained. This approach seems well suited to demographic ABM applications, as it is relatively intuitive, and it has been used to good effect in stochastic traffic simulations (Boukouvalas et al. 2014).

8.7 Discussion

This chapter has discussed both the design and analysis of agent-based demographic models. By treating ABMs as computational experiments and therefore choosing efficient designs appropriate to the complex nature of ABMs, an understanding of the simulator's behaviour can be obtained with relatively few runs (Urban and Fricker 2010). The adoption of relatively simple and easily generated Latin hyper-cube samples is recommended as a default, as is standard in much of the numerical computation literature (Santner et al. 2003). Gaussian Process Emulators may assist with the process of examining an agent-based demographic model, and with assessing sensitivity to various parameters as a check of robustness, as well as an aid to understanding. Input uncertainty can be propagated through the emulator to understand its effect on outputs through uncertainty analysis. Various calibration techniques enabling the provision of probability distributions for unknown parameters may be used to perform calibrated prediction, taking into account all possible sources of model uncertainty.

Gaussian Process Emulators have some restrictions that must be borne in mind. Firstly, they assume that both stochastic simulation uncertainty and uncertainty about untried inputs can be represented by Gaussian distributions, conditional on the data and hyper-parameters (MUCM 2011). Transformations may help to address these problems, but these make interpreting the output distribution more problematic (*ibid*). Secondly, a related implication of this assumption is that the output is smooth and not discontinuous. If a model features a very sharp discontinuity in the parameter space, it may struggle to be fully captured by the Gaussian Process Emulator (Gramacy 2005). As touched upon, the complex nature of ABMs means that discontinuities may occur. However, because Gaussian processes are constrained to lie close to the observations, they will at least show the location of sharp changes in outputs, as can be observed from the predictive plots in Figs. 8.2 and 8.3, even if these may be over-smoothed. Analytical expressions for the derivatives of Gaussian processes, given in Oakley (1999), may help to quickly identify the location of discontinuities and phase transitions, as these will be characterised by high gradients (*cf.* Luke 2007). Sequential Experimental Designs may also help understand behaviour around discontinuities. Improvement in predictive performance in these areas may be sought by rebuilding an emulator using an additional round of runs at locations chosen by some criteria that favour areas with high gradients.

Another limitation is that the emulators described above assume that stochastic variance inherent to the model is homoskedastic. This may not be a suitable assumption for some ABMs. By following the framework set out by Kersting et al. (2007), Boukouvalas (2010), and Ankenman et al. (2010) it is possible to relax this assumption by using paired emulators, one representing the simulator mean and the second the simulator variance. Additionally, Rasmussen and Williams (2006) describe the use of different link functions that generalise the Gaussian Process Emulator approach used in this paper in the spirit of the Generalised Linear Models

framework (Nelder and McCullagh 1972). Use of these extensions to the base model can further advance the utility of the emulation approach described above.

Despite the limitations discussed above, Gaussian Process Emulators allow modellers to understand the behaviour of their simulator and the uncertainties relating to it in an efficient and coherent manner, and provide tools for sensitivity analysis and calibration. The balance of flexibility, uncertainty quantification and interpretability give Gaussian processes advantages over both more flexible but opaque models such as neural networks, as well simpler but more rigid polynomial based meta-models – although much depends on the nature of the simulation under study.

Acknowledgements The authors gratefully acknowledge the Engineering and Physical Sciences Research Council (EPSRC) grant EP/H021698/1 “Care Life Cycle” and the support of the EPSRC Doctoral Training Centre grant (EP/G03690X/1). We are grateful to Jonathan Forster and two anonymous reviewers for insightful comments and suggestions. Any errors remain exclusively ours.

References

- Andrianakis, Y., & Challenor, P. G. (2011). Parameter estimation for Gaussian process emulators. Technical report, Managing Uncertainty in Complex Models. <http://www.mucm.ac.uk/Pages/Downloads/TechnicalReports/>
- Andrianakis, I., & Challenor, P. G. (2012). The effect of the nugget on Gaussian process emulators of computer models. *Computational Statistics & Data Analysis*, 56(12), 4215–4228.
- Ankenman, B., Nelson, B. L., & Staum, J. (2010). Stochastic kriging for simulation metamodeling. *Operations Research*, 58(2), 371–382.
- Aparicio Diaz, B., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography*, 48(2), 559–79.
- Barton, R. R., Nelson, B. L., & Xie, W. (2014). Quantifying input uncertainty via simulation confidence intervals. *INFORMS Journal on Computing*, 26(1), 74–87.
- Bastos, L. S., & O’Hagan, A. (2009). Diagnostics for Gaussian process emulators. *Technometrics*, 51(4), 425–438.
- Bijak, J. (2011). *Forecasting international migration in Europe: A Bayesian view*. Dordrecht: Springer.
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. D. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the United Kingdom with Gaussian process emulators. *Demographic Research*, 29(27), 729–766.
- Billari, F., Ongaro, F., & Prskawetz, A. (2003). Introduction: Agent-based computational demography. In F. Billari & A. Prskawetz (Eds.), *Agent-based computational demography: Using simulation to improve our understanding of demographic behaviour*, (pp. 1–17). Heidelberg: Physica-Verlag.
- Billari, F., Fent, T., Prskawetz, A., & Aparicio Diaz, B. (2007). The “Wedding-Ring”: An agent-based marriage model based on social interactions. *Demographic Research*, 17, 59–82.
- Boukouvalas, A. (2010). Emulation of random output simulators. Ph.D. thesis, Aston University. <http://eprints.aston.ac.uk/15776/>
- Boukouvalas, A., Sykes, P., Cornford, D., & Maruri-Aguilar, H. (2014). Bayesian precalibration of a large stochastic microsimulation model. *IEEE Transactions on Intelligent Transportation Systems*, 15(3), 1337–1347.

- Challenor, P. (2013). Experimental design for the validation of kriging metamodels in computer experiments. *Journal of Simulation*, 7(4), 290–296.
- Courgeau, D. (2012). *Probability and social science: Methodological relationships between the two approaches*. Dordrecht: Springer.
- Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Washington, DC/Cambridge: Brookings Institution Press/MIT Press.
- Fent, T., Aparicio Diaz, B., & Prskawetz, A. (2013). Family policies in the context of low fertility and social structure. *Demographic Research*, 29, 963–998.
- Forrester, A., Sobester, A., & Keane, A. (2008). *Engineering design via surrogate modelling: A practical guide*. Chichester: Wiley.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (3rd ed.). Abingdon: CRC.
- Gramacy, R. B. (2005). Bayesian treed Gaussian process models. Ph.D. thesis, University of California, Santa Cruz.
- Grazzini, J., & Richiardi, M. G. (2013). Consistent estimation of agent-based models by simulated minimum distance. Technical report 130, Laboratorio Riccardo Revelli.
- Grimm, V., & Railsback, S. F. (2005). *Individual-based modelling and ecology*. Princeton: Princeton University Press.
- Hankin, R. (2005). Introducing BACCO, an R bundle for Bayesian analysis of computer code output. *Journal of Statistical Software*, 14(16).
- Haylock, R. (1997). Bayesian inference about outputs of computationally expensive algorithms with uncertainty on the inputs. Ph.D. thesis, University of Nottingham. <http://eprints.nottingham.ac.uk/13193/1/338522.pdf>
- Kennedy, M. (2004). Descriptions of the Gaussian process model used in gem-sa. <http://www.tonyohagan.co.uk/academic/GEM/index.html>
- Kennedy, M., & O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society*, 63(3), 425–464.
- Kersting, K., Plagemann, C., Pfaff, P., & Burgard, W. (2007). Most likely heteroscedastic Gaussian process regression. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07* (pp. 393–400). Corvallis/New York: ACM. doi: [10.1145/1273496.1273546](https://doi.org/10.1145/1273496.1273546). <http://doi.acm.org/10.1145/1273496.1273546>. ISBN 978-1-59593-793-3.
- Klabunde, A. (2014). Computational economic modeling of migration. Technical Report 471, Ruhr University Bochum, Bochum.
- Kleijnen, J. P. (2008). *Design and analysis of simulation experiments*. New York: Springer.
- Kleijnen, J. P. C., & Mehdad, E. (2014). Multivariate versus univariate kriging metamodels for multi-response simulation models. *European Journal of Operational Research*, 236(2), 573–582.
- Law, A. M. (2007). *Simulation modeling and analysis* (4th ed.). New York: McGraw-Hill.
- Loepky, J. L., Sacks, J., & Welch, W. J. (2009). Choosing the sample size of a computer experiment: A practical guide. *Technometrics*, 51(4), 366–376.
- Luke, S. (2007). Finding interesting things: Population-based adaptive parameter sweeping. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation* (pp. 86–93), London.
- Mitchell, M. (2009). *Complexity: A guided tour*. Oxford: Oxford University Press.
- Montgomery, D. (2013). *Design and analysis of experiments* (8th ed.). Hoboken: Wiley.
- MUCM (2011). Managing uncertainty in complex models toolkit. mucm.aston.ac.uk
- Nelder, J., & McCullagh, P. (1972). Generalized linear models. *Journal of the Royal Statistical Society, Series A*, 135(3), 370–384.
- Oakley, J. (1999). Bayesian uncertainty analysis for complex computer codes. Ph.D. thesis, University of Sheffield.
- Oakley, J. (2002). Eliciting Gaussian process priors for complex computer codes. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 51(1), 81–97.
- Oakley, J., & O'Hagan, A. (2002). Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*, 89(4), 769–784.

- Oakley, J., & O'Hagan, A. (2004). Probabilistic sensitivity analysis of complex models: A Bayesian approach. *Journal of the Royal Statistical Society: Series B*, 66(3), 751–769.
- O'Hagan, A. (2006). Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering and System Safety*, 91(10–11), 1290–1300.
- O'Hagan, A., & Oakley, J. E. (2004). Probability is perfect, but we can't elicit it perfectly. *Reliability Engineering and System Safety*, 85(1–3), 239–248.
- Oreskes, N., Shrader-Frechette, K., & Belitz, K. (1994). Verification, validation, and confirmation of numerical models in the earth sciences. *Science*, 263, 641–646.
- Poole, D., & Raftery, A. E. (2000). Inference for deterministic simulation models: The Bayesian melding approach. *Journal of the American Statistical Association*, 95(452), 1244.
- Qian, P. Z. G., & Wu, C. F. J. (2008). Bayesian hierarchical modeling for integrating low-accuracy and high-accuracy experiments. *Technometrics*, 50(2), 192–204.
- R Development Core Team (2015). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna. <http://www.r-project.org/>
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. Cambridge: MIT. <http://www.ncbi.nlm.nih.gov/pubmed/15112367>
- Rossiter, S., Noble, J., & Bell, K. (2010). Social simulations: Improving interdisciplinary understanding of scientific positioning and validity. *Journal of Artificial Societies and Social Simulation*, 13(1), 10.
- Roustant, O., Ginsbourger, D., & Deville, Y. (2012). DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodelling and optimisation. *Journal of Statistical Software* 51(1).
- Saltelli, A., Tarantola, S., Campolongo, F., & Ratto, M. (2004). *Sensitivity analysis in practice*. Chichester: Wiley.
- Saltelli, A., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global sensitivity analysis. The primer*. Chichester: Wiley.
- Santner, T. J., Williams, B. J., & Notz, W. I. (2003). *The design and analysis of computer experiments*. New York: Springer.
- Ševčíková, H., Raftery, A. E., & Waddell, P. A. (2007). Assessing uncertainty in urban simulations using Bayesian melding. *Transportation Research Part B: Methodological*, 41(6), 652–669.
- Silverman, E., Bijak, J., & Noble, J. (2011). Feeding the beast: Can computational demographic models free us from the tyranny of data? In *Advances in Artificial Life, ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, Paris (pp. 747–757). MIT Press. <http://eprints.ecs.soton.ac.uk/22839/>
- Sobol, I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1–3), 271–280.
- Squazzoni, F. (2012). *Agent-based computational sociology* (1st ed.). Hoboken: Wiley.
- Urban, N. M., & Fricker, T. E. (2010). A comparison of Latin hypercube and grid ensemble designs for the multivariate emulation of an earth system model. *Computers & Geosciences*, 36(6), 746–755.
- Vernon, I., Goldsteiny, & M., Bowerz, R. G. (2010). Galaxy formation: A Bayesian uncertainty analysis. *Bayesian Analysis*, 5(4), 619–670.
- Werker, C., & Brenner, T. (2004). Empirical calibration of simulation models. Technical report, Max Planck Institute of Economics. <https://www.econstor.eu/handle/10419/88244/>
- Xie, W., Nelson, B. L., Barton, R. R., Nelson, B. L., & Barton, R. R. (2014). A Bayesian framework for quantifying uncertainty in stochastic simulation. *Operations Research*, 62(6), 1439–1452.

Chapter 9

How to Describe Agent-Based Models in Population Studies?

Jürgen Groeneveld, Anna Klabunde, Michelle L. O'Brien, and André Grow

9.1 Introduction

Agent-based modelling is a powerful method to investigate fundamental research questions that would be very hard to address by more established frameworks such as statistical models, equation-based models or reduced-form microsimulations (Epstein 2006). Whenever local interactions, adaptation or individual variability are considered to be important, agent-based simulation models seem to be a very useful addition to the methodological toolbox (Grimm 1999). Over the last years, such factors have increasingly moved into focus of population and migration research (cf. Courgeau et al., Chap. 2, in this volume). As a consequence, agent-based models are increasingly employed in these areas (Aparicio Diaz et al. 2011; Billari et al. 2007; Biondo et al. 2013; Espindola et al. 2006; Fent et al. 2013; Filho et al. 2011; García-Díaz and Moreno-Monroy 2012; Grow and Van Bavel 2015; Hassani-Mahmooei

J. Groeneveld (✉)

Department of Ecological Modelling, Helmholtz Centre for Environmental Research – UFZ,
Permoserstr. 15, 04318 Leipzig, Germany
e-mail: juergen.groeneveld@ufz.de

A. Klabunde

Max Planck Institute for Demographic Research, Konrad-Zuse-Straße 1, 18057 Rostock,
Germany

M.L. O'Brien

Department of Sociology, University of Washington, Seattle, WA, USA

Center for Studies in Demography and Ecology, University of Washington, Seattle, WA, USA

e-mail: shanmli@uw.edu

A. Grow

Centre for Sociological Research, University of Leuven (KU Leuven), Parkstraat 45,
3000 Leuven, Belgium

© Springer International Publishing Switzerland 2017

A. Grow, J. Van Bavel (eds.), *Agent-Based Modelling in Population Studies*,
The Springer Series on Demographic Methods and Population Analysis 41,
DOI 10.1007/978-3-319-32283-4_9

237

and Parris 2012; Hills and Todd 2008; Klabunde 2016; Kniveton et al. 2011, 2012; Rehm 2012; Reichlová 2005; Silveira et al. 2006; Smith 2014; Todd et al. 2005; Walsh et al. 2013).

Despite the advantages of the method, it has been discussed in the literature that progress and theory building in agent-based modelling has been hampered by the ad hoc nature of many agent-based simulations and incomplete and intransparent model descriptions (Müller et al. 2014; Richiardi et al. 2006; Waldherr and Wijermans 2013). Reproducibility is frequently not granted based on the published model description, which precludes the efficient build-up of a pool of baseline models about topics such as migration, fertility, or marriage.

To address these issues, guidelines, techniques and protocols have been suggested to describe agent-based and individual-based models (Grimm et al. 2006, 2010; Müller et al. 2013; Richiardi et al. 2006). The ODD (Overview, Design, and Details) protocol is an example for such a prescriptive verbal model description which originated in ecology (Grimm et al. 2006, 2010). Prescriptive verbal model descriptions fulfil a range of purposes such as understanding, communication and model comparison. The choice of the kind of model description is purpose dependent (Müller et al. 2014), that is, the purposes of communication to stakeholders on the one hand and to researchers aiming at quantitative model replication on the other hand require different kinds of model descriptions. For policy makers, agent-based models are a promising addition to their toolkit because of the possibility to perform policy experiments within the model. However, for the model results to seem trustworthy to a non-modeller, it is not as much provision of the code that matters, but a comprehensible and clear verbal model description.

It is important to understand what the current standard in model description is and which means of communicating the model structure and processes are currently being used in demography. This can yield best practice for future model description, which will eventually improve acceptance and usage of agent-based models, and possibly even the models themselves. In this chapter, we contribute to this aim by summarizing how agent-based models for human migration and population studies are currently described. Subsequently, we introduce the reader to the ODD + D protocol as suggested by Müller et al. (2013) and illustrate its usefulness for addressing some of the shortcomings of existing model descriptions by using an example of one agent-based model on circular migration (Klabunde 2016). We highlight that the use of such standardized descriptions can aid the understanding and increase reproducibility of agent-based models by readers. We have chosen the ODD + D (Overview, Design, Details + Decisions) protocol because the original ODD protocol became a standard in the description of agent-based ecological models and the ODD + D seemed to be a suitable extension for agent-based models where human decisions are explicitly modelled. Both the model developer and the first author of this paper prepared an ODD + D description based solely on the published natural language description of the model. This exercise helped us identify common pitfalls in model description and revealed how the published model description could have been improved upon had the ODD + D been used in the first place.

The remainder of this chapter is structured as follows: in Sect. 9.2 we reflect on the particularities of ABMs in population studies as compared to other disciplines and the requirements concerning the model description which arise from these. In Sect. 9.3 we review the descriptions of all agent-based models on human migration which have been published until 2014, as well as other prominent examples of agent-based demography models in order to identify whether a common standard has emerged. Section 9.4 contains a description of the main features of the ODD + D protocol and our experiment using this protocol. We discuss our findings in Sect. 9.5.

9.2 Requirements Concerning the Model Description for Demographic ABMs

To make models most useful for the target audience, it is important that model developers are specific and transparent about their modelling aims, the modelling approach, and the efforts taken to verify and validate the simulation model (Waldherr and Wijermans 2013). While this is true in general, we want to discuss the requirements for the model description in population studies in more detail. Population studies take a special place within the social sciences. Much more than other fields of social inquiry (e.g., sociology or social psychology), population research is data-driven and inherently empirical (Bijak et al. 2014; Morgan and Lynch 2001; Courgeau et al., Chap. 2, in this volume). This is also reflected in the use of agent-based modelling in this area. That is, in many areas of social research agent-based models are mostly used for conducting abstract thought experiments and researchers often rely on ‘stylized facts’, given a paucity of relevant data for initializing and validating simulation models (Boreo and Squazzoni 2005). Models in population studies, by contrast, increasingly draw on existing, large scale datasets (e.g., census data and social surveys) to create semi-artificial populations and use this data also for validating outcomes (e.g., Bijak et al. 2013; Grow and Van Bavel 2015; Williams et al., Chap. 6, in this volume). The question of how much data should be used in calibrating and validating a simulation model is therefore often of greater importance in population studies than it is in other areas of social inquiry (Klabunde and Willekens 2016; Silverman et al. 2011). A model description should thus make it transparent where and which data have been used in model development, calibration and validation.

The three main subject matters of demography are fertility, mortality and migration because they together determine the size and composition of a population (Courgeau and Franck 2007). However, fertility at the population level as well as migration flows and stocks of migrants are the result of decisions at the individual level, namely the decisions to migrate and to have a child. Mortality is indirectly influenced by individual behaviour through decisions that have an impact on health. Examples are the decisions to smoke or to exercise. These individual decisions are

often difficult to model with more traditional methods than agent-based modelling, in particular when decisions involve thresholds, if-then rules, and path-dependence (Bonabeau 2002). The flexibility of agent-based models allows implementing different types of decision making which have been shown to be empirically good predictors of behaviour. Here an important distinction between demographic ABMs and those in economics becomes evident: whereas the behavioural assumption in economic ABMs is usually utility maximization, in line with the current paradigm of the discipline, the variety of different decision rules in demographic ABMs is far greater (for a recent review, see Klabunde and Willekens 2016). Because of this variety in options to choose for modelling decision-making it is important to dedicate sufficient attention to describing the decision-making rules and to provide reasons why a particular rule was chosen.

Another important property of modern population studies is the recognition that demographic phenomena are highly interdependent (Courgeau and Franck 2007; Courgeau et al., Chap. 2, in this volume). One reason for this interdependence is that an individual takes demographically relevant decisions not independently from other decisions taken previously or from those planned for the future. Moreover, individuals are influenced by other individuals, such as the members of their own cohort or the members of the social networks they are part of. Migration and childbearing are very prominent examples of decisions which are regularly influenced strongly by the behaviour of others in the social network. Information as well as social and financial capital is transmitted through migrant networks. Migrants help new migrants and potential migrants during job search and provide assistance when first settling in. The prospect of this financial and emotional support decreases the perceived cost of migration and can thus convince a person that migration is a feasible option (see e.g., Haug 2008; Munshi 2003). Many demographic ABMs are built with the explicit aim of understanding the nature of such social influence processes (e.g., Aparicio Diaz et al. 2011; Billari et al. 2007; González-Bailón and Murphy 2013). Thus, demographic ABMs must be described in such a way that the kind of social influence modelled, e.g. whether and what kind of social network was used, becomes clear.

9.3 Model Descriptions in Agent-Based Simulation Models of Population Studies

As indicated above, modelling decision-making in demographic contexts with agent-based models is a relatively new emerging field. To assess the current practice in describing agent-based models in this context, we selected models which (i) have been published, (ii) are located in the field of population studies and/or migration and (iii) are clearly agent-based. We employ the concept of “agent” as defined by Macal and North (2010): agents are discrete entities which are capable of making autonomous decisions. They have goals, at least implicitly, and their

behaviour is rule-based. There is some kind of explicit interaction. There are no equations which govern the overall social structure on the macro level. This definition is in line with many other definitions in the literature (e.g., Epstein 2006; Tesfatsion 2006). Our focus on population studies encompasses models that centre on demographic outcomes such as births, deaths, marriages, and migration (cf. Demopædia 2015). Furthermore, we define migration as relocations which involve a change of residence beyond an administrative boundary. Thus, models of residential mobility are excluded, i.e. mobility initiated by a desire for better housing.

For a large subset of models fulfilling all three inclusion criteria in the area of migration research (Filho et al. 2011; Biondo et al. 2013; Espindola et al. 2006; García-Díaz and Moreno-Monroy 2012; Hassani-Mahmooei and Parris 2012; Ichinose et al. 2013; Klabunde 2016; Kniveton et al. 2011, 2012; Rehm 2012; Reichlová 2005; Silveira et al. 2006; Smith 2014; Walsh et al. 2013; Williams et al., Chap. 6, in this volume) we review the model description (Table 9.1). To our knowledge this selection comprises all agent-based migration models published until 2014 and two more recent papers by two authors of this chapter. In addition, we have included some prominent examples of agent-based population models (Aparicio Diaz et al. 2011; Bijak et al. 2013; Billari et al. 2007; Fent et al. 2013; González-Bailón and Murphy 2013; Hills and Todd 2008; Noble et al. 2012; Todd et al. 2005) that are widely cited (Table 9.2). Especially, we checked whether the authors have provided a graphical overview of their model, whether they have followed a specific protocol for model description, whether they have provided a table of the used reference parameter values and/or equations to describe decision rules, whether they have provided pseudo code or even published the complete source code and/or additional information in (online) appendices. Graphical model overviews, for example Unified Modelling Language (UML) charts (Rumbaugh et al. 1999), are often helpful for illustrating the scheduling of simulation runs and thereby can facilitate the re-implementation of a simulation model. However, especially when the scheduling deviates from a linear flow, providing pseudo code that details the exact order in which different procedures are executed can facilitate understanding of the exact modelling processes. Furthermore, providing a list with parameter values and/or equations facilitates understanding the models' basic processes (Grimm et al. 2006).

Our survey revealed that some authors use graphical model representations which differ strongly in style, ranging from technical UML class diagrams to illustrative simulation snapshots. Simple diagrams are used to illustrate behavioural assumptions, as in Kniveton et al. (2011) or Filho et al. (2011). More complicated UML diagrams are mostly used to illustrate processes as in Hassani-Mahmooei and Parris (2012) and Walsh et al. (2013), or as class diagrams for an overview of the model structure as in Kniveton et al. (2011). Examples for illustrative snapshots are found frequently, such as in Reichlova (2005), Noble et al. (2012) or Rehm (2012). They are usually used in spatially explicit models. Snapshots can be supplemented by illustrative components to show, for example, possible agent movements on a grid, as illustrated in Ichinose et al. (2013). An example of an animated simulation run can be found in Fent et al. (2013).

Table 9.1 Overview of means of description in agent-based migration models

	Graphical model overview	Simulation snap-shot	Using protocol	Table of parameters	Pseudo code	Equations used to describe decision rules	Narrative approach	Open source code	(Online) appendix
Biondo et al. (2013)	No	Yes	No	Yes (subset, no references)	No	Yes (social capital)	Yes	Yes (available at OpenABM)	Yes (short model description that accompanies model code)
Espindola et al. (2006)	No	No	No	No	No	No	No	No	No
Filho et al. (2011)	Yes (decision making, conceptual)	No	No	Yes (subset, no references)	Yes	Yes	No	No	No
García-Díaz and Moreno-Monroy (2012)	No	Yes	No	Yes (no references)	No	Yes	No	No	No
Hassani-Mahmooei and Parris (2012)	Yes (UML flow chart)	No	Yes (ODD)	Yes (subset, only references, no values)	No	Yes (some)	No	No	No
Ichinose et al. (2013)	No	Yes	No	No	No	No	Yes	No	Yes (movies of simulation runs and additional results)
Klabunde (2016)	No	Yes	Yes	Yes	No	Yes	No	Yes	Yes (input files, ODD + D description, additional files for model)

Kniveton et al. (2011)	Yes (conceptual model, UML class diagram)	No	No	No	No	No	No	No	No	No
Kniveton et al. (2012)	No	No	No	No	Yes (some)	No	No	No	Yes (information on empirical data, additional results, and model details)	No
Rehm (2012)	No	Yes	No	Yes (subset, no references)	Yes	No	No	No	No	No
Reichlová (2005)	No	Yes	No	No	Yes	No	No	No	No	No
Silveira et al. (2006)	No	Yes	No	No	Yes	No	No	No	No	No
Smith (2014)	Yes	No	No	Yes (subset)	Yes (some)	No	No	No	No	No
Walsh et al. (2013)	Yes	No	No	No	Yes	No	Yes	No	No	No
Williams et al. (Chap. 6, this volume)	Yes	No	Yes	No	Yes	No	Yes	No	Yes (ODD description, decision flow charts, variable lists, additional information)	No

In only one paper authors followed a particular protocol, namely Hassani-Mahmoei and Parris (2012) following the ODD protocol. Some authors provide additional information on the web, especially the ODD (+ D) protocol, including Klabunde (2016)¹ and Williams et al. (Chap. 6, in this volume).² In all other studies authors came up with a case specific tailored model description structure. In one paper algorithms were introduced by pseudo code (Filho et al. 2011), and four papers (Bijak et al. 2013; Biondo et al. 2013; Klabunde 2016¹; Noble et al. 2012) provided the model code online.³ One study provided an overview of decision rules in a table together with the main reference, which worked very well in aiding the understanding of the model (Hassani-Mahmoei and Parris 2012). Some authors (e.g., Biondo et al. 2013; García-Díaz and Moreno-Monroy 2012) provided tables with the parameter values or ranges used as baseline for the model. Ideally, this should be supplemented by the data sources used (if any) for determining those parameter values or ranges. Smith (2014) does this.

However, apart from very simple models usually not all parameter values are provided but only those considered important or meaningful by the author(s). This becomes evident when reading the verbal description carefully, and at the latest when trying to replicate a model. This can be problematic because the workings of the model can depend crucially on parameters which do not carry a lot of meaning in terms of model content, but which can alter the results dramatically. A sensitivity analysis (for examples see Aparicio Diaz et al. 2011; Billari et al. 2007 and also the chapters by Grow (Chap. 7) and by Hilton and Bijak (Chap. 8) in this volume) should be performed to determine the parameters that do have a large impact on model results (Thiele et al. 2014). If such an analysis was performed and the reported parameters chosen based on such an analysis, the author should say so. There may be cases where the large number of parameters does not allow for a systematic representation and sensitivity analyses. This problem can be mitigated if the code and all necessary files to run the model are provided in e.g. an online appendix or on a website or repository such as OpenABM.⁴

We found it very helpful when equations were used to describe the decision rules, as in Rehm (2012) or Bijak et al. (2013). Developers may find this superfluous, but in fact they are the only unambiguous way to communicate what the agents actually do. This can be supplemented by graphs of important functions. Smith (2014) illustrates the nonlinear way that the migration probability is assumed to depend on rainfalls by plotting the function.

Sometimes narratives might be helpful additions to better understand the model, for example to narrate the actions of a specific agent in the model similar to commenting actions of a specific actor in a sport event (Millington et al. 2012).

¹At <https://www.openabm.org/model/3893/version/3/view>

²At www.bit.ly/NepalABM

³We have not approached any of the authors for provision of the source code, which they might provide upon request.

⁴www.openabm.org

This allows authors to switch between the micro and the macro level of the model and helps to better understand the origins of the model dynamics. Biondo et al. (2013), Fent et al. (2013), Noble et al. (2012) and partly also Ichinose et al. (2013) are examples of such a narrative approach.

From this collection of ABM models it becomes clear that inclusion of graphical summaries of the model increases readability (see also Schwabish 2014). Furthermore comparability of models would be significantly increased if authors applied the same or at least a similar structure of model description, which would enable the reader to find relevant information in an efficient manner and to allow for replication. In the next section, we discuss the ODD + D protocol as one way of describing agent-based models in a standardized manner.

9.4 Example Application of the ODD + D Protocol

9.4.1 *The ODD + D Protocol*

The ODD + D is a protocol for prescriptive verbal model descriptions (full details can be found in Müller et al. 2013). ODD + D starts with the overview section, which should inform the reader about the purpose (e.g., what is the main research question that the modellers want to address with their model?), the entities (e.g., what types of agents are in the model?), the temporal and spatial scales (e.g., how do simulation steps map onto real-life time?), the process overview and the scheduling of the model (e.g., in the form of pseudo code or UML charts) in a concise manner. The overview section is followed by the design concepts sections where the author should report on the following ten design concepts that are ordered from general to more detailed information:

1. ‘Theoretical and empirical background’: The aim here is to put the work into context with existing theories, concepts and data.
2. ‘Individual decision making’: Here the author should provide details on the decision making submodel such as the object of decision making, whether agents are able to adapt their behaviour, or whether social norms or cultural values play a role in the decision making process.
3. ‘Learning’: Here it should be reported whether learning is considered in the model and briefly described.
4. ‘Sensing’: Includes the information that the agents can sense and therefore have available for their decision making. Also information whether the sensing process is erroneous and what costs are associated with sensing should be provided here.
5. ‘Prediction’: In the prediction design concept the authors should briefly describe if and what an agent can predict and in addition whether these predictions are systematically biased.

6. ‘Interactions’: This concept reports on the kind of interactions between agents.
7. ‘Collectives’: The collectives concept reports on the potential formation of collectives in the model.
8. ‘Heterogeneity’: Heterogeneity reports on the variability of agents, e.g. if there are several types of agents that behave in different ways.
9. ‘Stochasticity’: Mechanistic simulation models often contain stochastic processes. In this concept it should be reported which processes are affected by stochastic processes.
10. ‘Observation’: Finally the author should summarize what kind of information is collected and stored during the simulations that are used for the analysis.

Not all ten design concepts will be applicable for all models and therefore non-applicable concepts do not have to be addressed in the individual ODD + Ds. The main contribution of the more recent ODD + D protocol, above and beyond the older and perhaps better-known ODD (Grimm et al. 2006), is the addition of the ‘individual decision making’ design concept and the corresponding specific guiding questions that help describe the assumptions underlying decisions that the agents make. These questions are: On what assumptions is/are the agents’ decision model(s) based? Why is/are certain decision model(s) chosen? If the decision model is based on empirical data, where do the data come from? At which level of aggregation were the data available? What are the subjects and objects of the decision-making? Are multiple levels of decision making included? Do agents pursue an explicit objective or have other success criteria? How do agents make their decisions? Do the agents adapt their behaviour to changing endogenous and exogenous state variables? Do social norms or cultural values play a role in the decision-making process? Do spatial aspects play a role in the decision process? Do temporal aspects play a role in the decision process? To which extent and how is uncertainty included in the agents’ decision rules? As outlined above, one of the major aspects that distinguishes demographic agent-based models from models in, e.g., ecology or even economics is that researchers have a large pool of decision-making theories to choose from (see Klabunde and Willekens 2016 for a discussion). Moreover, temporal aspects are likely to play a part, since demography is often concerned with the timing of events, such as birth or marriage. Spatial aspects may also be important, especially in models of social influence. The results of models often crucially depend on the assumptions made with respect to decision-making. Therefore, in our applied example in the next section we use the ODD + D instead of the ODD.

The ODD + D finishes with the “Details” section where information on the initialisation of the model, the input data and submodels should be provided. The submodels should be described in such detail that allows replication. Since the details sections can be very long it is often feasible to include only the overview and design concepts part in the main text and provide the details in an online appendix. Guiding questions are provided for each entry in the ODD + D to support the user to compile the model description (see Table 1 in Müller et al. 2013).

9.4.2 *Applied Example*

To illustrate the usefulness of protocols of prescribed verbal language descriptions of agent-based simulation models, we have described a simulation model of circular human migration (Klabunde 2016) following the ODD + D protocol. In particular, we have described the model from two perspectives: from the modeller perspective who, of course, knows her intentions, and the reader's perspective who only had the published material at hand (without consulting the source code). The modeller's version of the ODD + D can be found on the platform OpenABM at <https://www.openabm.org/model/3893/version/3/view>, along with the model code and all files necessary to run the model.⁵ Klabunde (2016) is a model of circular labour migration of Mexican migrants to the US. Migration flows are largely determined by the structure of a network evolving over time. The model can replicate the distribution of migrants across cities in the US and the distribution of numbers of trips of migrants and thus offers an explanation for the observed patterns.

We compared the two ODD + Ds to investigate how the non-modeller comprehended the model based solely on the narrative (non-ODD + D) description provided in Klabunde (2016). We aimed to explain why the differences in comprehension between the modeller and reader arose. The experiment also served to identify some pitfalls in model description in general, and with the ODD + D in particular.

The experiment suggested that letting someone who has not originally implemented the model write the model description in a formal way helps to identify redundancies, or details that the programmer forgot to report since she may have the perception that these details are obvious and self-explanatory. For an independent reviewer it is sometimes easier to write a clear model description since she is not burdened with the history of the project and the technical difficulties. Thus, in joint projects, it might be sensible to have a person different from the programmer write the model description. In this particular case, the modeller sometimes did not stick to a consistent terminology in her non-ODD + D description (e.g. salary and wage). In verbal descriptions in scientific papers it is tempting to vary the wording in order to not have a lot of repetitions and make the text sound better. This should be avoided, because it comes at the cost of comprehensibility of the model description and confused the first author of this paper. Furthermore, some of the essential procedures were not explicitly mentioned in the ODD + D of the programmer, because technically they were subroutines of other procedures.

While using the ODD + D instead of a non-formal verbal description did improve completeness and clarity of model description considerably, it is important to mentally separate the meaning of different procedures and their implementation in the code. Seemingly simple processes may require many lines of code, whereas on the other hand there are smart algorithms and existing libraries which can render

⁵A working paper version of the verbal description is available at <http://www.rwi-essen.de/publikationen/ruhr-economic-papers/603/>

seemingly complicated processes lean and short. This should not influence model description, which should be guided by logically separable processes only. In a similar vein, in the ODD + D attributes should be assigned to the entities that they logically belong to, which is not necessarily implemented in exactly the same way in the code. In our example, “wage” was an attribute of firms in the code, not of workers. However, the ODD + D description became clearer when “wage” became an attribute of workers as well to better explain what happens in the model. This problem can be avoided if the ODD + D is not written *after* the model has been implemented, but *before*, as a starting point after conceptual model development and just before implementation. Protocols such as the ODD + D serve as a check list and the structured documentation of the model may help to organize the architecture of the entities and the whole implementation. This is still helpful when the model has to be adjusted after the first implementation. Updating the ODD + D will be simple and this process will not outweigh the benefits from having the ODD + D prior to the first implementation as a checklist and blueprint.

9.5 Discussion

We think that the ODD and the ODD + D are good starting points, however the structure of these protocols will be part of the scientific discourse and additions will be suggested. The current form of the ODD + D protocol will likely benefit from some rearrangements regarding the order in which things are reported. For example, sensing should be reported before decision making to be in line with the order how these processes are usually implemented. Furthermore, there should be more room given to networks; so far there is only one minor point on collective networks. Instead, there could be more specific guiding questions about the role and structure of the network, which would greatly improve the usability of the ODD + D for demographic applications. Finally, resources of the empirical information used in the model should not exclusively be asked for in respect to decision making, but for all empirical information used. Also, the filling in of the different sections should not be mechanical, but should always occur with the reader in mind. The reader cannot be expected to read the description several times, so one should make sure that every section is entirely clear given only the previous sections. Therefore, specific information might go into the overview section to enable the reader to fully understand the subsequent paragraphs.

In Sect. 9.2 we identified requirements that a model description should fulfil in particular for demographic ABMs: the description should be clear with respect to the usage of data, the implementation of decision-making, social influence and networks, and it should be written keeping in mind that non-modellers are not readily able to read source code. The ODD + D is helpful in addressing most of these special needs.

The subsection “Input Data” in the “Details” section provides room for empirical underpinnings of the model. Additionally, we find it very helpful when authors

provide data sources for parameters and time series in the form of overview tables. Regarding the description of decision-making, the ODD + D protocol provides clear and extensive guiding questions about the way decision-making was modelled and thus fulfils this requirement. As we already pointed out above, the ODD + D could be improved if more specific guiding questions regarding the type of network used were included. As long as this is not the case the modeller has to remember to do so, since “linked lives” and social influence are a feature of almost all ABMs in population studies and their exact implementation is a crucial prerequisite for model replication (see also Bijak et al. 2013). If a network is used, graphical representations of the network or of stylized parts of it can aid understanding considerably.

Finally, the ODD + D is a verbal description which nevertheless follows a rigorous structure and is thus a good middle path for researchers who need more information than just a quick verbal overview, but who do not want to or are not able to read source code.

The ODD and ODD + D protocols have been criticized as being ‘overdone’ for simple ABMs (Grimm et al. 2010) and our own experience from discussions at workshops and presentations is that researchers are often concerned that such descriptions are too long for standard publication in journals. However, we agree with Grimm et al. (2010) that the benefits of standardized descriptions also hold for simple models and there is always the possibility to publish the ODD + D as an online supplement, either on a personal webpage, or, preferably, on a platform such as OpenABM. This way of using the ODD + D is particularly useful in the case when several publications make use of the same model. The researcher can then refer to the same online location of the code description, and can keep the description in the actual paper concise. Changes to previous model versions can be pointed out in the paper and in updated versions of the ODD + D. We suggest this will facilitate to realize that the same model was used in several publications, which is much more efficient than having to go through a lengthy model description in each paper, only to realize that the model is already known. Many journals offer the possibility to provide an online appendix as well, which allows including the ODD + D as an appendix to each publication associated with one particular model. Of course the researcher should also be explicit about whether or not she is using the same model in different publications. Additionally, authors of increasingly complex models can use these online repositories, such as OpenABM or Github⁶, in order to streamline manuscripts for publication. Online repositories may hold all equations, graphical representations, code, and protocol for a model, while the manuscript may discuss a few examples and highlight general processes. This allows for a more flexible and legible text, while also providing all the information necessary for replication.

However, detailed prescriptive verbal model description results in additional workload. Whether this additional workload will be beneficial in contrast to well-

⁶<https://github.com>

documented source code will strongly depend on the audience and their willingness to consult source code. For example, the ODD + D protocol, especially the overview part together with a graphical representation such as activity diagrams, may offer an attractive tool for regulatory purposes which are not in place yet. For formal regulation processes, such as the authorisation process of new plant protection products, the responsible panel at the European Food Safety Authority (EFSA) published an opinion paper where they demand sufficient model documentation for mechanistic effect models and suggest but do not prescribe the use of standard protocols (EFSA PPR Panel 2014). For researchers, an online appendix with the ODD + D, UML diagrams and pseudo code greatly facilitates reproducibility. Furthermore, all parameter values used and their sources should be provided. In essence, if models should be used by other researchers or stakeholders they have to be trusted and made easily available. In general, the details provided in the “Details” section of the ODD + D protocol cannot be too detailed. Whatever can be expressed in a formal way should be, to ease reimplementing and to avoid confusion. Of course the source code should be published whenever this is possible as part of good modelling practice. We recommend writing the ODD + D description before starting to program because it helps separate the logical model structure from the structure of the code.

A further criticism that might be raised is that standardized descriptions are not flexible enough and might require the addition of new sections specific to the model at hand.⁷ The necessity of adding such sections reduces the level of standardization of the description but we do not think that this negates the benefits from the standardized description of the remaining model parts. In fact, the occurrence of such modifications in the context of demographic research might help to further improve the standard and to tailor it to the requirements of demographic simulation models. The explicit inclusion of decision processes in the original ODD protocol (leading to the ODD + D standard) is an example of such an extension based on existing research experience. Above, we have indicated similar extensions that might be necessary to further enhance the usefulness of the ODD + D protocol in the context of agent-based computational demography.

In this chapter, we have discussed the current practice in describing demographic agent-based models and have argued that standardized descriptions – in particular the ODD + D protocol – have the potential to help making model descriptions more transparent and to facilitate their reimplementing. Yet, model description is just one part of the modelling process or cycle (Schmolke et al. 2010). There are also promising initiatives (Grimm et al. 2014; Richiardi et al. 2006) to account for that and to provide frameworks to guide modellers beyond pure model description through the whole modelling process including model building, implementation, testing, simulation experiments, analysis, and validation. Discussing such comprehensive guidelines is out of the scope of this chapter, but we hope that our work

⁷We thank an anonymous reviewer for pointing this out.

contributes to furthering the development of common standards in the description of demographic agent-based models.

Acknowledgments We thank two anonymous reviewers for very valuable comments on the manuscript. Michelle L. O'Brien received partial support for this research from a *Shanahan Endowment Fellowship* and a *Eunice Kennedy Shriver* National Institute of Child Health and Human Development training grant, T32 HD007543, to the Center for Studies in Demography & Ecology at the University of Washington. André Grow has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement no. 312290 for the GENDERBALL project.

References

- Aparicio Diaz, B., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography*, *48*(2), 559–579.
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. D. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the United Kingdom with Gaussian process emulators. *Demographic Research*, *29*, 729–766.
- Bijak, J., Courgeau, D., Silverman, E., & Franck, R. (2014). Quantifying paradigm change in demography. *Demographic Research*, *30*, 911–924.
- Billari, F. C., Prskawetz, A., Aparicio Diaz, B., & Fent, T. (2007). The “Wedding-Ring”: An agent-based marriage model based on social interaction. *Demographic Research*, *17*, 59–82.
- Biondo, A. E., Pluchino, A., & Rapisarda, A. (2013). Return migration after brain drain: A simulation approach. *Journal of Artificial Societies and Social Simulation*, *16*(2), 11.
- Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(Suppl. 3), 7280–7287.
- Boreo, R., & Squazzoni, F. (2005). Does empirical embeddedness matter? Methodological issues on agent-based models for analytical social science. *Journal of Artificial Societies and Social Simulation*, *8*(4), 6.
- Courgeau, D., & Franck, R. (2007). Demography, a fully formed science or a science in the making? An outline programme. *Population-E*, *62*(1), 39–45.
- Demopædia. (2015). *Multilingual demographic dictionary*, Unified Second Edition, English volume, 2nd edition. http://en-ii.demopaedia.org/wiki/Main_Page
- EFSA PPR Panel. (2014). Scientific opinion on good modelling practice in the context of mechanistic effect models for risk assessment of plant protection products. *The EFSA Journal*, *12*(3), 3589.
- Epstein, J. M. (2006). Agent-based computational models and generative social science. In J. M. Epstein (Ed.), *Generative social science: Studies in agent-based computational modeling* (pp. 4–46). Princeton: Princeton University Press.
- Espindola, A. L., Silveira, J. J., & Penna, T. J. P. (2006). A Harris-Todaro agent-based model to rural urban migration. *Brazilian Journal of Physics*, *36*(3), 603–609.
- Fent, T., Aparicio Diaz, B., & Prskawetz, A. (2013). Family policies in the context of low fertility and social structure. *Demographic Research*, *29*, 963–998.
- Filho, H. S. B., de Lima Neto, F. B., & Fusco, W. (2011). Migration and social networks – An explanatory multi-evolutionary agent-based model. In *IEEE symposium on intelligent agent* (pp. 1–7). Paris.
- García-Díaz, C., & Moreno-Monroy, A. I. (2012). Social influence, agent heterogeneity and the emergence of the urban informal sector. *Physica A: Statistical Mechanics and its Applications*, *391*(4), 1563–1574.

- González-Bailón, S., & Murphy, T. E. (2013). The effects of social interactions on fertility decline in nineteenth-century France: An agent-based simulation experiment. *Population Studies*, 67(2), 135–155.
- Grimm, V. (1999). Ten years of individual-based modelling in ecology: What have we learned and what could we learn in the future? *Ecological Modelling*, 115(2–3), 129–148.
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., et al. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2), 115–126.
- Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., & Railsback, S. F. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, 221(23), 2760–2768.
- Grimm, V., Augusiak, J., Focks, A., Frank, B. M., Gabsi, F., Johnston, A. S. A., et al. (2014). Towards better modelling and decision support: Documenting model development, testing, and analysis using TRACE. *Ecological Modelling*, 280, 129–139.
- Grow, A., & Van Bavel, J. (2015). Assortative mating and the reversal of gender inequality in education in Europe: An agent-based model. *PLoS ONE*, 10(6), e0127806.
- Hassani-Mahmoei, B., & Parris, B. W. (2012). Climate change and internal migration patterns in Bangladesh: An agent-based model. *Environment and Development Economics*, 17(6), 763–780.
- Haug, S. (2008). Migration networks and migration decision-making. *Journal of Ethnic and Migration Studies*, 34(4), 585–605.
- Hills, T., & Todd, P. M. (2008). Population heterogeneity and individual differences in an assortative agent-based Marriage and Divorce Model (MADAM) using search with relaxing expectations. *Journal of Artificial Societies and Social Simulation*, 11(4), 5.
- Ichinose, G., Saito, M., Sayama, H., & Wilson, D. S. (2013). Adaptive long-range migration promotes cooperation under tempting conditions. *Scientific Reports*, 3, 2509.
- Klabunde, A. (2016). Computational economic modelling of migration. In S.-H. Chen, M. Kaboudan & Y.-R. Du (Eds.), *Handbook on computational economics and finance* (pp. 536–568). Oxford: Oxford University Press (forthcoming).
- Klabunde, A., & Willekens, F. (2016). Decision-making in agent-based models of migration – State of the art and challenges. *European Journal of Population*, 32(1), 73–97.
- Kniveton, D. R., Smith, C., & Wood, S. (2011). Agent-based model simulations of future changes in migration flows for Burkina Faso. *Global Environmental Change*, 21(Sup.1), S34–S40.
- Kniveton, D. R., Smith, C. D., & Black, R. (2012). Emerging migration flows in a changing climate in dryland Africa. *Nature Climate Change*, 2, 444–447.
- Macal, C. M., & North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of Simulation*, 4, 151–162.
- Millington, J. D. A., O’Sullivan, D., & Perry, G. L. W. (2012). Model histories: Narrative explanation in generative simulation modelling. *Geoforum*, 43(6), 1025–1034.
- Morgan, S. P., & Lynch, S. M. (2001). Success and future of demography: The role of data and methods. *Annals of the New York Academy of Sciences*, 954, 35–51.
- Müller, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., et al. (2013). Describing human decisions in agent-based models – ODD + D, an extension of the ODD protocol. *Environmental Modelling & Software*, 48, 37–48.
- Müller, B., Balbi, S., Buchmann, C. M., de Sousa, L., Dressler, G., Groeneveld, J., et al. (2014). Standardised and transparent model descriptions for agent-based models – Current status and prospects. *Environmental Modelling & Software*, 55, 156–163.
- Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the U.S. labour market. *The Quarterly Journal of Economics*, 118(2), 549–599.
- Noble, J., Silverman, E., Bijak, J., Rossiter, S., Evandrou, M., Bullock, S., & Vlachantoni, A Falkingham, J. (2012). Linked lives: The utility of an agent-based approach to modeling partnership and household formation in the context of social care. In *Proceedings of the winter simulation conference 2012 (WSC2012)* (pp. 1–12). Berlin: IEEE.
- Rehm, M. (2012). *Migration and remittances. An agent-based model*. Ph.D. thesis, New School of Social Research, New York.

- Reichlová, N. (2005). *Can the theory of motivation explain migration decisions?* Charles University Prague, Faculty of Social Sciences, Institute of Economic Studies, IES Working Papers UK FSV-IES No.97, Prague.
- Richiardi, M., Leombruni, R., Saam, N., & Sonnessa, M. (2006). A common protocol for agent-based social simulation. *Journal of Artificial Societies and Social Simulation*, 9(1), 15.
- Rumbaugh, J., Jacobson, I., & Booch, G. (1999). *The unified modeling language reference manual*. Boston: Addison-Wesley.
- Schmolke, A., Thorbek, P., DeAngelis, D. L., & Grimm, V. (2010). Ecological models supporting environmental decision making: A strategy for the future. *Trends in Ecology & Evolution*, 25(8), 479–486.
- Schwabish, J. A. (2014). An economist's guide to visualizing data. *The Journal of Economic Perspectives*, 28(1), 209–233.
- Silveira, J. J., Espindola, A. L., & Penna, T. J. P. (2006). Agent-based model to rural-urban migration analysis. *Physica A: Statistical Mechanics and its Applications*, 364, 445–456.
- Silverman, E., Bijak, J., & Noble, J. (2011). Feeding the beast: Can computational demographic models free us from the tyranny of data? In T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, M. Dorigo, & R. Doursat (Eds.), *Advances in artificial life, ECAL 2011: Proceedings of the eleventh European conference on the synthesis and simulation of living systems*.
- Smith, C. D. (2014). Modelling migration futures: Development and testing of the rainfalls agent-based migration model – Tanzania. *Climate and Development*, 6(1), 77–91.
- Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In L. Tesfatsion & K. L. Judd (Eds.), *Handbooks in economics 2* (pp. 831–880). Oxford: North-Holland.
- Thiele, J. C., Kurth, W., & Grimm, V. (2014). Facilitating parameter estimation and sensitivity analysis of agent-based models: A cookbook using NetLogo and R. *Journal of Artificial Societies and Social Simulation*, 17(3), 11.
- Todd, P. M., Billari, F. C., & Simão, J. (2005). Aggregate age-at-marriage patterns from individual mate-search heuristics. *Demography*, 42(3), 559–574.
- Waldherr, A., & Wijermans, N. (2013). Communicating social simulation to sceptical minds. *Journal of Artificial Societies and Social Simulation*, 16(4), 13.
- Walsh, S. J., Malanson, G. P., Entwisle, B., Rindfuss, R. R., Mucha, P. J., Heumann, B. W., et al. (2013). Design of an agent-based model to examine population–environment interactions in Nang Rong District, Thailand. *Applied Geography*, 39, 183–198.

Part III
Modelling Decision Processes

Chapter 10

The Decision to Emigrate: A Simulation Model Based on the Theory of Planned Behaviour

Frans Willekens

10.1 Introduction

Many people desire to emigrate but few do leave their country, resulting in relatively low levels of international migration (e.g. Esipova et al. 2011; Bilsborrow 2012; van Dalen and Henkens 2012; Abel and Sander 2014; Moses 2015). That observation motivated this paper. I present a decision model that reproduces the observation. A combination of the theory of planned behaviour (Ajzen 1985; Fishbein and Ajzen 2010) and the process character of the emigration decision offers an explanation for the discrepancy between the desire to emigrate and emigration. To show that, the theory of planned behaviour (TPB) is extended to a process theory of planned behaviour and the process theory is applied to model the emigration decision. The model is validated by assessing its ability to predict stylized facts of international migration.

The theory of planned behaviour states that intentions predict behaviour. Intentions are shaped by the subjective evaluation of the outcomes of the behaviour (behavioural belief), the individual's perception of normative pressures (normative beliefs), and the individual's perception of facilitators and obstacles that influence the performance of the behaviour (control beliefs). Actual access to resources moderates the effect of intention on behaviour. Fishbein developed the theory in the 1970s as the theory of reasoned action and Ajzen extended the theory and called

F. Willekens (✉)

Netherlands Interdisciplinary Demographic Institute (NIDI), The Hague, The Netherlands
e-mail: Willekens@nidi.nl

it the theory of planned behaviour.¹ The theory has inspired migration researchers for decades (e.g. de Jong and Fawcett 1981). Today one observes a growing interest in computational models of decision making that implement the theory of planned behaviour (see e.g. Kniveton et al. 2011, 2012; Orr et al. 2013; Willekens 2013; Klabunde 2014; Klabunde and Willekens 2016).

The focus of the paper is on the decision to emigrate. Emigration is a relocation of the usual residence to another country. Two features of the decision-making are highlighted: the sequential nature of decision-making and the uncertainties involved. The sequential nature of the decision process has two implications. First, the decision should be modelled as a process consisting of several stages. Second, the theory of planned behaviour should incorporate time in order to account for the time it takes to form behavioural, normative and control beliefs. That calls for an extension to a process theory of planned behaviour. Uncertainty is an important factor in the decision to emigrate because many events and conditions that affect the decision process and its outcome cannot be predicted with certainty. To properly account for the uncertainties, the decision process should be modelled as a stochastic process.

The model proposed in this paper is a *multistage stochastic process* model. In each stage, an individual has two options. He or she may continue to the next stage or may decide to drop out of the decision process, which implies a decision to stay in the country. Dropout is an important part of the model and a necessary condition to reproduce empirical regularities in emigration. Both systematic factors (e.g. personal attributes and context) and random factors affect the decision process and its outcome. Discrete (binary) choice models are used to account for the uncertainty. The duration in a stage at time of continuation or dropout is random too. Possible values of the duration are given by a probability distribution, known as waiting time distribution. A common waiting time distribution is the exponential distribution, which follows from the assumption that the rate of leaving a stage is constant. The exponential distribution is used in this paper, but the normal distribution and the beta distribution are used too. The parameters of the waiting time distribution depend on systematic factors that may vary. In the model presented in this paper, parameters do not vary with age as in most models of migration. The age profile of emigration is an *outcome* of the model instead. The age at which an individual emigrates depends on how long an individual stays in each of the stages of the decision process.

The proposed model is referred to as a simulation model and not an agent-based model. The strength of the model is the operationalization of the theory of planned behaviour into a stochastic process model of action. The emphasis is on the stages of the process, the time in each stage, and the random factors involved. Individuals have the capacity to act independently and make their own choices. The influence of others on the action is through social norms and support. The social interactions that generate norms and support are not considered explicitly. Hence, the model is not

¹See Fishbein and Ajzen (2010, p. 18 ff) for a historical perspective on the theory of reasoned action and theory of planned behaviour.

an agent-based model. An extension of the model to an agent-based model requires that social norms and support emerge as outcomes of interactions between agents. Viewing the model as a simulation model is also pragmatic. A discussion of whether the model meets the criteria of an agent-based model would shift the attention away from the strength of the model. The contribution of this paper to agent-based modelling is the operationalization of an established theory of action into a model of action that goes beyond earlier models that have used the same theory of action. The model proposed in this paper extends earlier work by recognizing that actions are outcomes of random (decision) processes.

In this paper, a single systematic factor is considered for illustrative purposes: skill level. An individual has one of two skill levels: low/medium or high. The precise definition of each skill level is not important to illustrate the model. The cut-off point could be at completed secondary education. Individuals with not more than completed secondary education have a low/medium skill level. A high skill level requires a completed tertiary education. The simulation starts with a cohort of 15-year olds. The skill level at age 15 is the skill level an individual ever develops.

Some parameters of the model are derived directly from data on international migration, more particularly from the Gallup World Poll 2005 on the desire to emigrate. Most parameters are however plausible guesstimates; they are not estimated from data because data do not exist. The validity of the model is determined by its ability to reproduce stylized facts on international migration. Three facts are singled out, two relate to the level of emigration and one to the age pattern. The first stylized fact is the annual emigration rate, recently estimated by Abel and Sander (2014). The second is the proportion of the world population that is living in a country different from the country of birth. That figure is published in the World Migration Report issued by the International Organization for Migration. The third is the typical age profile of migration. That profile was documented extensively in the literature (see e.g. Rogers and Castro 1981; Raymer and Rogers 2008). If the model is a plausible description of the emigration decision process and reproduces these facts, the model passes the test of validation.

The structure of the paper is as follows. The TPB is reviewed in Sect. 10.2. The theory is extended into a process theory of decision-making in Sect. 10.3. The process theory of planned behaviour has much in common with other theories that view decision making as a process with stages. To place the process TPB in context, Sect. 10.3 includes brief descriptions of other process theories of decision-making. The main part of Sect. 10.3 is the multistage process model that implements the TPB. The model is a multistate event history model. States are stages and events are transitions between stages. In each stage of the decision process, an individual may decide to continue to the next stage or decide to stay in the country and hence abandon the decision process. Continuation and drop-out are competing transitions to which the theory of competing risks apply (Marley and Colonius 1992). The parameters of the model are based on the few observations on international migration available, augmented by guesstimates. The data used in the paper are presented in Sect. 10.4. The parameters of the process model are presented

in Sect. 10.5. Section 10.6 presents the outcomes of the model and compares the outcomes with observations on emigration. A discussion of model development and findings is presented in Sect. 10.7.

10.2 The Theory of Planned Behaviour

The TPB was developed by Ajzen as an extension of Fishbein's theory of reasoned action (Fishbein and Ajzen 2010). The theory of reasoned action states that the best predictor of behaviour is the individual's attitude toward the behaviour along with the social norms that influence the likelihood of performing the behaviour. The key determinants of attitudes are the individual's subjective expectations about salient consequences of the behaviour. In this section, the theory is reviewed, with a focus on emigration as the relevant behaviour.

Ajzen (1985) added perceived and actual control to the theory of reasoned action, in order to capture the influence of perceived obstacles and constraints that might prevent behaviour from occurring. The theory of planned behaviour is summarized in Fig. 10.1. Three beliefs determine an intention to act, i.e. emigrate. The first is the subjective belief that emigration is beneficial to one's future well-being. It is the basis for a positive attitude (ATT) towards emigration. The subjective belief may be an outcome of a conscious calculus, but may also be a result of limited and biased information received from others (emigrants and non-migrants). The second is the subjective belief that significant others approve of the emigration. That normative belief determines the subjective norm (SN). The third is the subjective belief that one has the capabilities to remove obstacles and to make emigration a success. That control belief determines the perceived behavioural control (PBC). ATT, SN and PBC are predictors of the intention to perform a behaviour. The stronger ATT, SN and PBC, the stronger the intention.

Many individuals who intend to emigrate do not leave their country because they lack the actual capability to remove barriers and take advantage of opportunities. In the remainder of this section, the beliefs are discussed in more detail. The discussion is more in-depth than is required for the modelling in the next section of the paper. In the last section of the paper, I will draw on the discussion to recommend directions of research aimed at a theory-driven simulation model of emigration. In the theory of planned behaviour, attitudes, subjective norms and perceived behavioural control act independently on intentions. They do not interact.

Background factors influence the formation of beliefs, and, indirectly, intentions and behaviour. They include personal characteristics and societal factors (Fig. 10.1). They determine differences in behavioural, normative and control beliefs.

For illustrative purposes, one background factor is included in the model presented in this paper: skill level. The skill level is an outcome of education and training.

Attitudes towards emigration, subjective norms and perceived and actual behavioural control are now discussed in more detail.

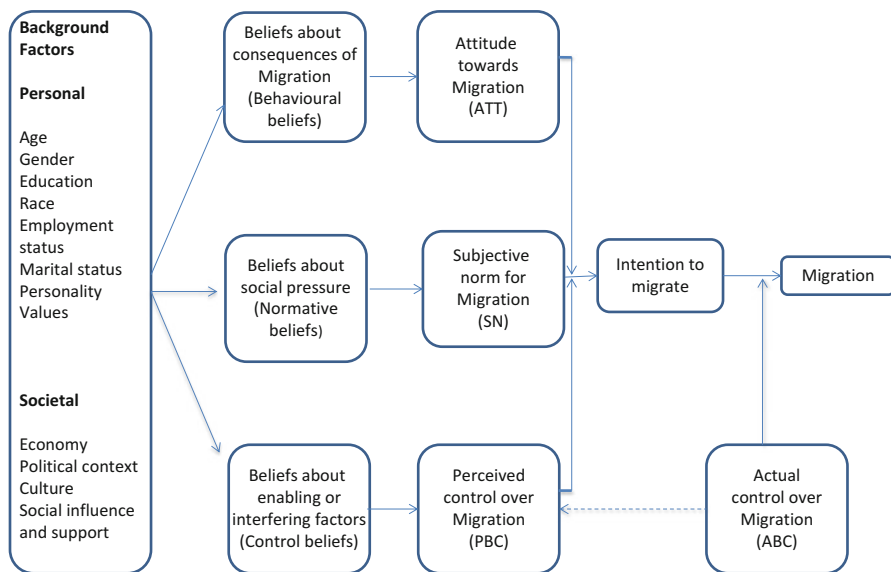


Fig. 10.1 Schematic presentation of the theory of planned behaviour (Source: Adapted from Fishbein and Ajzen (2010, p. 22))

(a) Attitude towards emigration

Emigration may produce several outcomes. Let i denote an outcome of emigration. Individuals associate various qualities or characteristics with an outcome, some are good and some are bad. By associating qualities with outcomes, individuals form a belief about the likely consequences of emigration and these beliefs determine the attitude towards emigration. The belief that emigration results in outcome i is the subjective probability that emigration produces outcome i . It is denoted by p_i . The strength of a belief in outcome i is expressed as a likelihood: the probability that, according to the individual, emigration produces outcome i . The subjective value of outcome i to an individual is denoted by e_i . Beliefs and subjective values determine the attitude towards emigration. The behavioural attitude (BA) is the sum of beliefs that emigration produces specific benefits and costs, weighted by the subjective values the individual attaches to these benefits and costs. It is the *expected value* of emigration:

$$BA = \sum_i p_i e_i$$

The belief that emigration produces outcome i may depend on the time frame an individual considers. In the short term, costs are likely to exceed benefits, but in the long run, perceived benefits may outweigh costs. Townsend and Oomen (2015) found that emigrants take short-term risks to secure long-term benefits.

The evaluation of an outcome i depends on inborn needs and on values persons hold, i.e. what they consider important in life. De Jong and Fawcett (1981) distinguish eight values that influence the decision to migrate.

Behavioural beliefs are related to reasons or motives for or against emigration. Reasons against emigration received considerably less attention than reasons in favour of emigration (push and pull factors). Beliefs are formed as a result of direct observation, education, word of mouth, media attention, and other factors. Personality and emotions are background factors that influence how people go about to form beliefs in the likelihood of particular outcomes and the valuation of outcomes. Van Dalen et al. (2005) found that the intentions to emigrate out of Africa are largely driven by optimism surrounding the net benefits of emigration.

In the process model presented in this paper, the effects of the factors that determine BA are summarized in a binary score. The score summarizes effects of many factors that influence the benefits and costs of emigration. If the value is one, the individual develops a positive attitude towards emigration. If the value is zero, the individual does not consider emigration to be beneficial. The model can be extended easily if data on these factors are available. The factors may be included in a binary logit model or a binomial logistic regression model, which produces a BA score between 0 and 1. The rationale for using a binary score in this paper is to specify a model that is as simple as possible and as complicated as necessary.

(b) Subjective norm with respect to emigration

A positive attitude towards emigration, i.e. a belief that emigration is beneficial, is a necessary condition for developing an intention to emigrate. It is not a sufficient condition, however. Many people expect to benefit from emigration but few intend to emigrate. Individuals who consider emigration may be sensitive to group norms, social pressure and social approval. Let i denote an important referent (individual or institution). Important normative referent individuals for the decision to emigrate are one's partner, family members, and friends. A partner who supports the emigration decision can act as an important stimulus (Van Dalen and Henkens 2012). Opinion leaders may be important too. Ajzen (2006) provides guidelines for eliciting salient normative referent individuals.

The influence of referent i on the individual depends on the individual's perception of what i wants, e.g. the normative belief, and on the individual's motivation to comply. The subjective norm (SN) with respect to emigration is the sum of the subjective evaluations of different social norms regarding how to behave, weighted by the individual's motivation to comply:

$$SN = \sum_i c_i n_i$$

where n_i is the belief of what referent i wants and c_i is the motivation to comply. Subjective norms change over time. As people grow older, referents change and individuals are more likely to comply with some social norms and reject other

norms. Subjective norms also change because of the diffusion of values, ideas and norms in society.

The formation of normative beliefs is a learning process and different forms of learning contribute to it. Following Montgomery and Casterline (1996) normative beliefs are acquired by learning from others and the influence through which some individuals exert control over others, by virtue of their power or authority. Ajzen and Klobas (2013) distinguish between normative beliefs based on what others say and those based on what others do. Some societies develop a culture of emigration (see e.g. Cohen and Sirkeci 2011; Kōu and Bailey 2014). In these societies, individual normative beliefs evolve to shared (collective) normative beliefs and become institutionalized. The Philippines, Mongolia and Ireland developed a culture of emigration. Some agricultural societies developed a culture of migration to maintain farm size. Stark et al. (2009) propose the idea that a culture of migration (*programmed migration*) has an evolutionary edge, i.e. that some populations might develop a genetic disposition to migrate.

In this paper, the effects of the factors that determine SN are summarized in an SN score between 0 and 1. If data on the determinants of SN are available, a binary logit model or logistic regression can be used to produce an SN score between 0 and 1.

(c) Perceived control over emigration

Individuals who consider emigration beneficial and who experience the social pressure to emigrate will not develop an intention to emigrate unless they believe that they have the resources to remove the obstacles to emigration and to make emigration a success. The perceived behavioural control (PBC) is the extent to which people believe that they are capable of performing a given action. It takes into account the availability of skills, opportunities, constraints and resources required to perform the action. The concept is closely related to Bandura's (1977) concept of self-efficacy and the sociological concept of agency. Ajzen (2002) indicates that PBC can be viewed as the combined influence of two components: self-efficacy (a person's judgment about being able to perform a particular action) and controllability (the extent to which the performance of the action is up to the actor) (see also Fishbein and Ajzen 2010, pp. 165 ff). Self-efficacy depends on available resources and the belief that barriers can be removed, while controllability depends on the presence of obstacles. Resources include financial means, but also human capital, social capital and cultural capital. Obstacles include distance (physical and cultural distance), institutional barriers (visa requirements, lack of portability of pensions and health insurance, lack of recognition of professional qualifications, etc.) and cultural barriers (differences in language, religion, etc.). Belot and Ederveen (2012) consider cultural barriers and their effect on migration between OECD countries. They find that cultural barriers do a much better job in explaining the pattern of migration flows between developed countries than traditional economic variables such as income and unemployment differentials. Adsera and Pytlikoa (2012) study the role of language differences and language diversity in shaping international migration. They develop an indicator of language

distance and find that migration rates increase with linguistic proximity. People who lack the necessary language skills are less likely to emigrate. Isphording and Otten (2014) develop measures of linguistic distance to study the impact of linguistic distance on language acquisition of immigrants. The lower the distance, the easier it is for immigrants to acquire the language of the destination country.

Let i represent a control factor that facilitates or inhibits emigration and let o_i denote the power of control factor i to facilitate or inhibit emigration. The control belief q_i is the subjective probability or belief that control factor i is present. The perceived behavioural control (PBC) is the sum of control beliefs, weighted by their perceived power:

$$PBC = \sum_i q_i o_i$$

Consider an example. To an individual who considers emigration beneficial, visa and residence permit are control factors. If an individual does not want to emigrate unless he or she has a valid visa and does not want to emigrate to a country unless he or she has a valid residence permit, then o_i is large. If an individual intends to enter a country illegally (without a visa) or intends to overstay a tourist visa (without a residence permit), the factor is not important, i.e. o_i is low. The perceived control over visa requirement and residence permit depends on the individual's subjective belief that he/she will get it, which is denoted by q_i .

In this paper, the perceived behavioural control is summarized in a PBC score that can be negative or positive. A logit model is used to convert the score to a value between 0 and 1.

(d) Actual control over emigration

Emigration intentions are good predictors of emigration if perceived behavioural control is matched by actual behavioural control (ABC). Whether intentions predict behaviour depends in part on factors beyond the individual's control (Ajzen 2011). People who overstate their capabilities to overcome barriers and to take advantage of opportunities and facilitators of migration are likely to remain in the intention stage. A necessary condition for intentions to predict actions is that individuals have actual control over their behaviour, i.e. they are able to behave as intended. Using longitudinal data on individuals in the Netherlands who expressed an intention to emigrate, Van Dalen and Henkens (2013) found that emigration intention is a good predictor of emigration. About one third (34%) of respondents who stated an intention to emigrate actually emigrated within a 5-year follow-up period. De Groot et al. (2011) found, in a study of residential mobility, that people with a strong intention to move are almost four times as likely to move than people with a less strong intention to move. De Jong (1994) reviewed several studies on internal migration in different countries and found that people who intend to migrate are three to four times as likely to migrate in a specified time frame than people who intend to stay. If intentions predict behaviour, we have an effective prediction method. Intentions are often not good predictors, however. De Jong (1994) gives

several reasons for the inconsistencies between intentions and behaviour. The difference between perceived and actual control is one factor. Another is the insufficient detail in measuring intentions in a survey, a reason for inconsistency also stressed by Ajzen (2011). A third reason is changing intentions. If the interval between measurement of intentions and recording of behaviour is large, intentions may have changed. Intentions and behaviour should therefore be measured in a relatively narrow period of time (see also Ajzen 2011).

The theory of planned behaviour indicates that the reason for a weak performance of intentions as predictors of behaviour is the discrepancy between perceived behavioural control and actual behavioural control. The difference is expected to be smaller when PBC is measured closer to emigration. The closer to the emigration, the more accurate an individual's perception of self-efficacy, barriers, resources and support is likely to be (Sheeran et al. 2003). Ajzen uses PBC as a proxy measure for actual behavioural control, and notes that it can substitute for control when an individual's perceptions are realistic (for a discussion, see Darnton 2008).

In this paper, the ABC score is the PBC score plus a random factor. The random factor measures the uncertainty individuals with a given PBC score have about the actual resources they need to be able to emigrate.

10.3 The Process Theory of Planned Behaviour

In this section the theory of planned behaviour is extended to a process theory of decision making. A process theory of decision making emphasizes the temporal dimension and stresses that decision making involves a progression through a number of stages. A number of process theories of human decision making exist. Two theories are briefly reviewed. The process theory of planned behaviour is presented next. In the process theory of planned behaviour, attitudes, intentions and behaviour are treated as stages of a decision process. The theory is operationalized in a stochastic process model with several stages and transitions between the stages. Stages do not need to follow a fixed sequence. For instance, if an individual considers emigration because significant others expect him or her to emigrate (e.g. to follow an education, to get a job, or to join a partner for marriage), then SN triggers the interest in emigration. In this paper, I consider a fixed sequence.

10.3.1 Process Theories of Decision Making: A Brief Review

The two process theories reviewed are the Rubicon model and the 'horse race' model. Process theories not covered include Janis and Mann (1977), the transtheoretical model of action (Prochaska et al. 1992) and the dynamic model of job search (McCall 1970).

Heckhausen (1991) presents a phase model of action, known as the Rubicon model. It originated in developmental psychology. The model postulates that individuals pursue development goals to produce the life course they want and mobilize cognitive and other resources to achieve the goals. Developmental goals are anticipated end states. They motivate an individual to act in a particular way. The process of action consists of several stages. It begins with the awakening of a wish to achieve a goal and ends after the goal has been accomplished. The initial Rubicon model (Heckhausen 1991) distinguishes four phases: the predecisional phase, the postdecisional but pro-actional phase, the actional phase and the postactional phase. Transitions between the phases are discrete shifts rather than gradual changes (hence the reference to Rubicon). Later the Rubicon model was extended to a theory of motivation that covers the entire life span (Heckhausen et al. 2010; Heckhausen and Heckhausen 2010). Kley (2011) adopted the Rubicon model to study the migration decision process. Coulter (2013) used Kley's process model to study the abandonment of desires to relocate in the context of residential mobility. He is one of the few authors who stress the need to study the decision to stay, which is the abandonment of the desire to relocate. Abandonment is as much an expression of agency as the decision to move. Kōu and Bailey (2014), in a study of the emigration of highly skilled Indians to the Netherlands and the UK, embed the phase model of action in the life course and show how individuals and families mobilize different types of resources and access different networks to assure that emigration produces the desired outcome.

The 'horse race' model is an offspring of random utility theory. Random utility models account for the stochastic variability underlying choices due to differences between individuals, between the object of choice, and changes in choice situation. The random utility discrete choice model predicts the probability of a choice between a limited number of alternatives. It does not consider the time it takes to reach a decision and it gives no insight into the cognitive process that underlies decision making. Marley and Colonius (1992) extended the random utility model by including the time individuals take to accumulate and process evidence in favour of an alternative. The time, known as response time, deliberation time and decision time, is random and follows a response time distribution, which is a waiting time distribution. The factors that influence the choice affect the choice probability as well as the time it takes to make a decision. The evidence accumulation model is a simple description of the cognitive process that underlies decision-making. In psychology, there is considerable support for the thesis that evidence accumulation drives decision making (see e.g. Rodriguez et al. 2014; Usher et al. 2013). A particularly useful observation, made by Marley and Colonius (1992) is the relation between the evidence accumulation model and the theory of competing risks. The challenge is to determine the joint likelihood of a decision (deliberation choice) and the time it takes to make a decision (deliberation time) (see also Colonius 2001; Hawkins et al. 2014). Some decisions are taken quickly, while other decisions take a lot of deliberation, which requires time (Kahneman 2011). For a brief and general overview of models that account for effects of deliberation times on choice probabilities, see Busemeyer and Rieskamp (2014). Early attempts to extend the

discrete choice model to integrate choice probabilities and waiting times to the decision/action include Pudney (1989) in economics. The model Pudney proposed is a competing risk model too. The competing risk model and the theory of competing risks have untapped potential in choice modelling.

Hybrid choice models (Ben-Akiva et al. 2002, 2012) are extensions of discrete choice models. They introduce elements of social psychology in economic choice models.

The above process theories distinguish several stages or phases in the process of decision making. They originated in different disciplines and therefore seem to differ considerably. They have important elements in common, however. First, time is explicit and time matters. Second, stages are similar. Successive stages imply an increased commitment to the intended action. Third, the benefit or utility of the action is uncertain. Hence decisions are made under uncertainty. Fourth, valuations of alternatives are subjective. They depend on inborn characteristics and one's values, preferences and goals. They also depend on the incomplete information available to the individual at a point in time. Fifth, process theories seem to be converging to a transdisciplinary theory of action; they increasingly incorporate elements of other theories and disciplines. Several of the elements are also included in the process model of planned behaviour, presented in the next section.

10.3.2 Process Model of Planned Behaviour

The process model originates from the theory of planned behaviour and has features in common with process models reviewed in Sect. 10.3.1, in particular the 'horse race' random utility model, the rubicon model and the transtheoretical model of action. To be consistent with the theory of planned behaviour, the process model should distinguish at least four stages. A person in the first stage never considered emigration. The person leaves the state when he/she develops an interest in emigration as a viable option or decides that emigration is not a viable option. In the second stage, the person develops behavioural beliefs, normative beliefs and control beliefs. These beliefs determine ATT, SN and PBC. ATT, SN and PBC determine the intention to emigrate. Attitude (ATT) is a latent disposition or tendency to favour or disfavour an action (Fishbein and Ajzen 2010, p. 76). Fishbein and Ajzen use the term attitude to refer to the evaluation of a behaviour along a dimension of favour or disfavour, good or bad, like or dislike, approval or disapproval, advantageous or disadvantageous. A person who has developed an intention to emigrate moves to the next stage and starts planning and preparation. During this stage, the person needs to mobilize resources, to overcome barriers and to take advantage of opportunities that may arise. Planning and preparation will be successful if the person is capable of dealing adequately with control factors. In case the actual behavioural control is deficient, the person is likely to stay. Persons who leave the country enter the fourth and final stage of the decision process. The first stage is denoted by 'n'; the second stage by 'a', the third stage by 'i' and the fourth




Fig. 10.2 Stages of the emigration decision: the process theory of planned behaviour

stage by ‘e’. The model is shown in Fig. 10.2. The model is programmed in R. The source code is available from the author.

In the second stage, ATT, SN and PBC may act independently on intention or they may interact. In the theory of planned behaviour, ATT, SN and PBC are independent. Ajzen recognizes the possibility that PBC moderates the effect of ATT and SN on intention, but that interaction effect is not a formal part of the TPB (Ajzen 2002; Fishbein and Ajzen 2010, p. 181). The reason Ajzen has given is that being capable of performing an action does not imply an intention to perform that action. In the literature, PBC interactions did not receive much attention because, in statistical models, the interaction is often not significant. Yser (2012) argues that the limited attention to PBC interactions is a missed opportunity for advancing our understanding of intention formation (for a discussion, see Boudewyns 2013). Fife-Schaw et al. (2007) study the moderating effect of SN on the effect of ATT on intention. In the process model proposed in this paper, SN and PBC are intervening factors in the transition from attitude to intention. They moderate the effects of attitude on intention, i.e. a positive attitude leads to an intention only if the SN and the PBC are supportive (positive). An individual who considers emigration beneficial may or may not develop an intention to emigrate depending on (1) the perceptions of what significant others want and (2) the perceived available resources.

Individuals in a given stage may be thought of as collecting information and accumulating evidence to make a decision whether to continue the emigration decision process or to drop out and to stay in the country (at least for the foreseeable future). The time it takes to reach a decision depends on (a) the stage in the decision process, and (b) individual attributes including personality traits, and contextual (societal) factors. Continuation and discontinuation (dropout, attrition) are competing risks. They compete to be the reason for exit from the current stage. Simple rules govern the choice. Exceeding a threshold, as in the ‘horse race’ model is one such rule. A similar rule is used in the process model of the TPB (see below).

The decision process depends on personal attributes. In this paper, I consider a single attribute: skill level. As stated in the introduction, the precise definition of each skill level is not important to illustrate the model. Skill levels are assigned to individuals. Since we have no data on skill levels, a random draw from a theoretical probability distribution determines the skill level of an individual.

Skill level is a binary variable, because it has two possible values: low/medium and high. A low or medium skill level is coded 0 and a high skill level is 1. The theoretical probability distribution with two possible values is the Bernoulli distribution. For each individual in the virtual population, a random number is drawn from a Bernoulli distribution. The distribution has a single parameter p (usually denoted as probability of success). Random draws from a Bernoulli distribution produce 0's and 1's. If the number is 0, the individual is allocated a low/medium skill level. If it is 1, the skill level is high. When the virtual population is sufficiently large, the proportion of highly skilled is equal to the theoretical value p .

In the process model of the TPB presented in this paper, the probability that an individual develops an interest in emigration, i.e. considers emigration, is a parameter of the model. The age at which an individual considers emigration as a viable option and effectively starts the emigration decision process is the age at which he or she starts assessing the benefits and costs of emigration. The age at which an individual enters the assessment stage is a random variable denoted by X_a , the age at transition from stage n to stage a . The possible values of X_a and the likelihood of each of the possible values are described by a probability distribution. If data are available on ages at which individuals start reflecting on the advantages and disadvantages of emigration, the distribution can be determined empirically and the data should be used. In the absence of empirical evidence, the ages may be inferred from a theoretical distribution. The distribution of ages at which individuals consider emigration is essentially a waiting time distribution. Waiting time distributions are common in survival analysis and event history analysis (see e.g. Steele 2005; Aalen et al. 2008). They play a central role in the process model of planned behaviour.

The waiting time distribution is a probability distribution and, as any other probability distribution, it has three related specifications. The probability density function gives the probability that a transition occurs at a given exact age. The distribution function gives the probability that a transition occurs before a given age. The survival function gives the probability that a transition does not occur before a given age. Several waiting time distributions are documented in the literature. The normal distribution and the logistic distribution are among them. The probability densities of these distributions are symmetric. The normal distribution is often justified as a limit of sums, including sums of waiting times (Central Limit Theorem). For instance, if an individual accumulates several pieces of evidence to determine the costs and benefits of emigration, and the accumulation of each piece of evidence takes time, then the sum of the distributions of the durations approximates a normal distribution. In the stochastic process model presented in this paper, it is assumed that the age (X_a) at which individuals consider emigration, i.e. develop an interest in emigration, and effectively start the emigration decision process follows a truncated normal distribution (Burkardt 2014; Pudney 1989, pp. 302 ff). A truncated distribution is used to prevent that individuals at very young age or even negative age are selected. The age distribution implies that the rate of exit from the first stage (never considered emigration – n) and entry into the second stage (assessment – a) increases with age.

Individuals in the second stage develop beliefs leading to ATT, SN and PBC. First, consider the development of behavioural beliefs and an attitude towards emigration. It takes time to accumulate evidence to determine whether emigration is beneficial. That time, which will be denoted by T_{a+} , is a random variable with possible values described by a waiting time distribution. I assume a simple exponential waiting time distribution, which implies a constant exit rate (by skill level). A constant rate for individuals with similar skill levels does not mean that the rate at which individuals complete the assessment of the benefits and costs of emigration is constant. The rate may decline because of selection. If individuals with skill level A accumulate evidence faster or need less evidence than individuals with skill level B, then the share of people with skill level B increases in the population at risk and the average rate of completing the assessment declines with time. In the process model presented in this paper, the time a given individual k takes to determine whether emigration has a net benefit is obtained by a random draw from an exponential waiting time distribution with constant rate.

The direction of exit (continuation or abandonment of decision process) depends on the probability that the benefits of emigration exceeds the costs. In the model presented in this paper, the probability is fixed exogenously. The benefit of emigration does not depend on cost-benefit calculations. Klabunde (2014) presents a model in which the benefit of emigration depends on the expected future income and the benefit derived from family reunification.

If emigration is perceived as beneficial, beliefs are developed about social pressure to emigrate or stay and about the availability of resources. The age at which the formation of these beliefs starts is the sum of the age at which the individual first considers emigration and the time it takes to determine whether emigration is beneficial. The likelihood that an individual continues to the intention stage depends on the outcome of that belief formation and the strength of the perceived behavioural control (PBC) and the subjective norm (SN).² The time it takes to proceed to the intention stage depends on PBC and SN too. Individuals with a strong belief in their resources, strong normative belief and a strong motivation to comply, and individuals with a strong belief in the absence of resources and with a strong normative belief that he/she should not emigrate need less time to determine whether they intend to emigrate or stay than individuals with moderate levels of resources and normative beliefs. An individual with a high degree of self-efficacy, i.e. who believes that he or she is able to mobilize resources and remove barriers, is more likely to develop an intention sooner than individuals who doubt about their ability to mobilize resources and support. Individuals with a strong SN to stay or a weak SN to emigrate, and individuals with low PBC are not likely to develop an intention to emigrate. They are more likely to abandon the decision process, and to abandon it sooner the weaker the SN and the lower the PBC. Similarly to the ‘horse race’

²Recall that this perspective differs from that in the theory of planned behaviour. In the TPB, attitude (considering emigration beneficial), SN and PBC act independently on intentions; the effect of attitude on intention is not moderated by SN and PBC.

model, the process model of planned behaviour distinguishes between likelihood of ever developing an intention to emigrate and the timing of that intention.

Suppose we record, for each individual considering emigration beneficial, the skill level and whether and when that individual develops an intention or decides to discontinue the decision process. Suppose we also have SN and PBC scores. In empirical studies, SN is usually measured on a 7-point scale from *I should not to I should*. PBC is usually measured on a 7-point scale from *no control* and *full control*. For this paper on international migration, I have no data on skill level and SN and PBC values. The distribution of skill level in the population and the distribution of SN and PBC scores by skill level are generated by random draws from theoretical distributions. In other words, each individual is given a skill level and a SN and PBC value by drawing a random number from probability distributions. I assume that SN is a continuous random variable with few people having a very low SN, the majority having a relatively low to medium SN and a sizable minority having a high SN. The beta distribution is used because it is a flexible distribution, which can produce different shapes. The distribution is defined on the interval from 0 to 1, with 0 representing the complete absence of a subjective norm and 1 a very strong subjective norm and a high willingness to comply. The beta distribution has two positive shape parameters and is able to describe different shapes. The distribution of PBC scores in the population follows a normal distribution, with mean and variance that vary by skill level. PBC scores can be large positive and negative values. These values are transformed to values between 0 and 1 using the cumulative distribution of PBC values.

In principle, the distributions of SN and PBC in the population depend on several factors. The SN score may depend on the presence of emigrants in one's social network and on the level of remittances received from these emigrants. The PBC score may depend on the presence of opportunities (e.g. job offer; admission to college) and barriers (e.g. border enforcement, cost of emigration), and on one's assessment of his or her access to resources to take advantage of opportunities and remove barriers. These factors are not explicit in this paper. For a model that incorporates these factors explicitly, see Klabunde (2014).

Individuals with very high SN and PBC are likely to develop an intention to emigrate soon after they consider emigration beneficial. Individuals with a desire to emigrate but with very low SN and PBC scores are likely to drop out of the decision process and to drop out soon after developing a desire to emigrate. Individuals with low (high) SN and low (high) PBC or with average values of SN and PBC take more (less) time to decide whether to intend to emigrate or to abandon the decision process. These are the persons who doubt about social support and financial and other resources. SN and PBC scores determine the probability of choice and the time it takes to make a choice, as in the accumulation model.

The effects of SN and PBC are combined into a single score to determine their effect on the intention to emigrate. The effects are assumed to be multiplicative, not additive as in the original theory of planned behaviour. A very low level of either SN or PBC results in a low score. To obtain a high score, both SN and PBC should be high. The score is a value between 0 and 1. SN and PBC may be weighted

differently in computing the score. In a society where self-reliance is valued highly, PBC receives a high weight. In societies where conformity is valued highly, SN receives a high weight. A Cobb-Douglas utility function incorporates these desired features (Pindyck and Rubinfeld 2013). The function is

$$V(SN, PBC) = \gamma SN^\alpha PBC^\beta \quad (10.1)$$

where V is the score and α , β and γ are parameters to be fixed in simulation. V is a random variable because SN and PBC are random variables. The parameter γ is a scaling factor. The score an individual receives depends on his or her SN and PBC levels. The parameters α and β are elasticities, with α the percentage increase in score V resulting from a 1% increase in SN and β is the percentage increase in V resulting from a 1% increase in PBC. The model Kniveton et al. (2012) developed is a special case of the Cobb-Douglas score model. If $\alpha = \beta = 1$ and γ is the proportion of individuals that believe they can mobilize resources and remove barriers, then the model is that of Kniveton et al. (2012).

The effect of skill level (and other covariates) on V is indirect. The skill level of an individual influences the SN and PBC scores. Some latent characteristics may influence the score directly. To accommodate that effect, a random effect is added to V . Addition of a random factor results in a model similar to the random utility model, used in discrete choice models (see e.g. Hess and Daly 2014). I do not use discrete choice theory in the model, although one could approach the choice between developing an intention to emigrate and discontinuation of the decision process as a binary choice with a random utility affecting V . In the *current* version of the model, an individual develops an intention to emigrate if his/her score V is equal to or exceeds a threshold value to be estimated from data or fixed in the simulation. Let V_H denote the threshold value and V_k the score of individual k . If V_k is less than the threshold value, individual k drops out of the decision process. Hence, individual k 's intention to emigrate is 1 if $V_k \geq V_H$ and 0 if $V_k < V_H$. The threshold value determines the proportion of people with a desire to emigrate that develops an intention to emigrate. Because the distribution of V in the sample population is known, the proportion developing an intention changes when V_H changes. If the threshold value is not known, but the proportion of people with a desire to emigrate that develops an intention (P_H) is observed empirically or fixed in simulation, then V_H is the value of V for which $P(V \geq V_H) = P_H$. It is a quantile of the distribution of V .

In the model, an individual develops an intention to emigrate if his/her V score is equal to or exceeds a threshold value V_H . The threshold value is determined by the exogenously given proportion developing an intention. Klabunde (2014) follows a different approach. Instead of V , she uses an intention score that varies from minus infinity to infinity. A linear model relates a person's migration intention score to SN and PBC scores. The intention score is transformed to a probability, using the logistic distribution (logit model). A random draw from $U[0,1]$ determines whether an individual moves to the next decision stage.

The time it takes to develop normative and control beliefs and SN and PBC is a random variable, which will be denoted by T_{v+} . The possible values are given by a waiting time distribution. The value of T_{v+} depends on V and, indirectly on SN and PBC. Small and large values of V result in small waiting times and values of V close to the median value produce large waiting times. The following function is used to derive the waiting time from the value of V :

$$T_{v+} = -\theta \ln [abs(V - median(V))] \quad (10.2)$$

where $median(V)$ is the median of all values of V in the population, $abs()$ denotes the absolute value and θ is a scaling factor to be estimated from data or fixed in simulation. The factor determines the location of the waiting time distribution on the time axis. The time an individual needs to develop normative and control beliefs and SN and PBC is a realization of T_{v+} . In the model presented in this paper, the distribution of T_{v+} does not follow a theoretical distribution. Hence, the ‘empirical’ distribution, which is the distribution in the virtual population, is used. The waiting time is equal to θ if $1 = -\ln [abs(V - median(V))]$ or $\exp(-1) = abs[V - median(V)]$. In Klabunde (2014), the distribution of T_{v+} is an exponential distribution, with the rate of transition from the attitude stage to the intention stage depending on the intention score.

The waiting time between developing an interest in emigration and developing an intention or drop-out is the time it takes to develop normative and control beliefs (T_{av}) plus the time it takes to develop an intention (T_{vi}).

$$T_{ai} = T_{av} + T_{vi}$$

The probability that an individual develops an intention to emigrate is determined by the probability of considering emigration, the probability that an individual who considers emigration considers emigration beneficial, and the probability that an individual who considers emigration beneficial has sufficiently high levels of SN and PBC. The age at which an individual develops an intention to emigrate is determined by the age at considering emigration and the time it takes to assess the benefits and costs of emigration and to develop normative and control beliefs and SN and PBC. It is equal to: $X_i = X_a + T_{av} + T_{vi}$, with X_i the age at developing an intention to emigrate (transition from v to i). The age X_i in the (virtual) population is the sum of three random variables, X_a , T_{av} and T_{vi} . The distribution of that sum is a convolution of three distributions. The distribution of X_a is a normal distribution and the distribution of T_{av} is an exponential distribution. The distribution of T_{vi} does not follow a theoretical distribution. The probability that an individual of a given age did not yet develop an intention to emigrate is the *empirical* survival function of X_i . The non-parametric method (Kaplan-Meier estimator) is used to estimate the empirical survival function. The probability density of X_i gives the distribution of ages at which individuals develop an intention to emigrate. The (cumulative) distribution function of X_i gives the probability of having developed an intention at a given age, i.e. the probability of being in the intention stage.

The rate of transition from the assessment stage to the intention stage is derived from the empirical survival function and its Kaplan-Meier estimator. The rate may also be estimated using the Cox model. In this paper, the Cox model with skill level as stratification variable is used to estimate the cumulative hazard rate of transition into the intention stage, by age. Using skill level as a stratification variable results in two baseline hazard curves, one for the low/medium-skilled and one for the highly skilled. Note that the transition rates vary with age as a consequence of the effects of SN and PBC.

An individual with an interest in emigration but who drops out of an emigration decision process may develop an interest later again. That implies a transition from a later stage in the decision process to an earlier stage. In this paper, I disregard such transitions. Those who abandon the process are removed from the population at risk of emigration. A transition to an earlier stage may be incorporated easily in a multistate model of the decision process, of which the Markov decision process is a well known example (e.g. Guo and Hernández-Lerma 2009).³

Individuals who develop an intention to emigrate move to the next stage of the decision process: planning and preparation. During the planning and preparation stage, some individuals abandon the decision process and stay. Others complete the planning and preparation, and emigrate. Emigration and drop-out of the decision process (stay) are competing risks. The outcome of the decision depends on the actual behavioural control (ABC). Individuals with adequate ABC are likely to leave. Those without adequate financial, human, social and cultural resources and those who do not get the necessary permits are likely to stay. In the model, ABC is PBC plus a random factor. If the random factor is negative, the actual behavioural control is less than the perceived behavioural control. Some individuals may also have an ABC score that exceeds the PBC score. The random factor is drawn from a uniform distribution. The minimum and maximum values of the uniform distribution determine the largest differences between ABC and PBC. The ABC score is expressed as a figure between 0 and 1, with 0 a complete absence of any actual control over resources and support and 1 unlimited supply of resources. Individuals with high ABC will almost certainly emigrate and individuals with low ABC are highly unlikely to perform the intended behaviour. They end the decision process. The duration in the intention stage at time of exit (emigration or dropout) is a random variable. The possible values follow an exponential waiting time distribution with a single parameter; namely, the exit rate from the intention

³The Markov decision process is an analytical tool for sequential decision making under uncertainty. A Markov decision process generalizes a continuous-time Markov process in that a decision process is embedded in a Markov model and the process involves a sequence of actions (Alagoz et al. 2010). A model of the Markov decision process distinguishes states and actions. The probability that an individual continues to the next stage depends on the current state and the action. An action results in a reward. The value of the reward is unknown in advance. An individual knows the expected value, however. Markov decision processes are used to determine the times of transitions to the next stage that maximize lifetime rewards. In the model presented in this paper, an individual gets a reward if he/she emigrates. The reward is the net benefit of emigration.

stage. Exit rates vary between individuals because of differences in ABC scores. ABC scores are assumed to remain constant in time. An individual gets a score when he/she enters the intention stage and keeps that score until he/she leaves. For the cohort that enters the intention stage, exit rates will decrease in time because individuals with high emigration rates or dropout rates leave soon and the ‘survivors’ have lower rates. In the process model, emigration rates increase exponentially with ABC level. The emigration rate at a given value of ABC is

$$\mu_{ie}(ABC) = a_e \exp [b_e * ABC] \quad (10.3)$$

where a_e and b_e are nonnegative parameters to be estimated from data or to be fixed in simulation.

Dropout rates decrease exponentially with ABC:

$$\mu_{ic}(ABC) = a_c \exp [b_c * (1 - ABC)] \quad (10.4)$$

where c represents drop-out (censoring) and a_c and b_c are nonnegative parameters to be estimated or fixed. The rate of leaving the intention stage is the sum of the two transition rates:

$$\mu_{i+}(ABC) = \mu_{ie}(ABC) + \mu_{ic}(ABC) \quad (10.5)$$

The exit rate $\mu_{i+}(ABC)$ is a bathtub shaped hazard function. It is high at low values of ABC, because of dropout; it decreases when ABC increases but is still too low to affect emigration significantly; it increases when higher ABC levels push the emigration rate up; and it is high when ABC levels near their maximum value of one. Xie and Lai (1995) used a similar bathtub hazard rate function. Instead of using two Gompertz-like distributions, they used two Weibull distributions. The model was used later by Bebbington et al. (2006), among others. Bathtub distributions receive considerable interest in reliability engineering (see e.g. Almalki 2013).

Emigration and drop-out, i.e. stay, are competing risks. The probability that an individual who leaves the intention stage emigrates is $\frac{\mu_{ie}(ABC)}{\mu_{i+}(ABC)}$ and the probability that the individual stays is $\frac{\mu_{ic}(ABC)}{\mu_{i+}(ABC)}$.

The length of time an individual stays in the intention stage is a random variable. It depends on the rate of leaving the intention stage (exit rate), which depends on the ABC score. The probability that individual k exits the intention stage t years after entering the intention stage is the survival function $S_i(t, ABC_k) = \exp [-\mu_{i+}(ABC_k) t]$, where ABC_k is k 's level of actual behavioural control. The probability that individual k in the intention stage exits precisely after t years is the density function $f_{i+}(t, ABC_k) = \mu_{i+}(ABC_k) S_i(t, ABC_k)$. The probability that an individual, who intends to emigrate, emigrates at time t is $f_{ie}(t, ABC_k) = \frac{\mu_{ie}(ABC_k)}{\mu_{i+}(ABC_k)} f_{i+}(t, ABC_k)$ and the probability that he or she drops out of the decision

process at t is $f_{ic}(t, ABC_k) = \frac{\mu_{ic}(ABC_k)}{\mu_{i+}(ABC_k)} f_{i+}(t, ABC_k)$. The probability that individual k who intends to emigrate, emigrates within t years, is the cumulative incidence function $I_{ie}(t, ABC_k) = \int_0^t f_{ie}(\tau, ABC_k) d\tau = \int_0^t \mu_{ie}(ABC_k) S_i(\tau, ABC_k) d\tau$.

The probability that individual k , who intends to emigrate, drops out of the decision process within t years is the cumulative incidence $C_{ic}(t, ABC_k) = \int_0^t f_{ic}(\tau, ABC_k) d\tau = \int_0^t \mu_{ic}(ABC_k) S_i(\tau, ABC_k) d\tau$.

The transition rates $\mu_{ie}(ABC_k)$ and $\mu_{ic}(ABC_k)$ determine the timing of exit from the intention stage and the reason for exit (emigration or dropout). Individuals with the same ABC score exit at different times because of random factors (chance). The time at which individual k exits from the intention stage is obtained by a random draw from the exponential waiting time distribution with constant exit rate $\mu_{i+}(ABC_k)$. Let u denote a random draw from a standard uniform distribution $U(0,1)$. Individual k with ABC score ABC_k exits the intention stage at time ${}_kT_{i+}(ABC_k) = -\frac{\ln(u)}{\mu_{i+}(ABC_k)}$. Note that large exit rates lead to small exit times. The exit time depends on the random value u . When u is close to zero, the exit time is large; when u is close to one, the exit time is small. The reason for exit is determined by a random draw from a Bernoulli distribution with parameter $p_{ie}(ABC_k) = \frac{\mu_{ie}(ABC_k)}{\mu_{i+}(ABC_k)}$. Individual k emigrates if the value of the random number drawn is less than $p_{ie}(ABC_k)$, otherwise k stays.

The age at emigration is the sum of four random variables: the age at considering emigration, the time it takes for an individual who considers emigration to determine whether emigration is beneficial or not, the time it takes for an individual who considers emigration beneficial to develop an intention to emigrate, and the time an individual with an intention to emigrate needs to plan and prepare the departure. The age at emigration is:

$$X_e = X_a + T_{ai} + T_{ie} \tag{10.6}$$

The distribution of X_e depends on the distributions of X_a , T_{ai} and T_{ie} . It is not a theoretical probability distribution unless X_a , T_{ai} and T_{ie} follow theoretical distributions. For instance, if X_a is normally distributed and T_{ai} and T_{ie} are exponential waiting time distributions, then X_e follows a double exponential distribution (Coale and McNeil 1972).

10.4 Data

Although the model is not designed for a particular data set, it should reproduce at least stylised facts about emigration. The facts are: (1) observations on proportions of the world population that desire to emigrate, that intend to emigrate and that actually emigrate, (2) observations on levels of international migration, expressed as the emigration rate, (3) observations on lifetime international migrations in the

world, and (4) the typical age profile of migration, known in the literature as the Rogers-Castro migration age profile because the pattern was first documented extensively by Rogers and Castro (1981). The typical pattern is used widely in migration studies, in particular projections. The United Nations use a simplified version of the Rogers-Castro migration age profile in the World Population Projections 2010 (for a discussion, see Abel et al. 2014). Since the age at emigration is not an input variable in the process model of the TPB, but an output, the validity of the model is larger if it reproduces the typical migration age profile. The model should also allow for migrant selection. To that end, two skill levels are distinguished: low/medium and high.

A worldwide Gallup survey in 2005 among 750 thousand adults found that 14 % of the world's adults (15+) population (630 million) say they would like to emigrate if they could. Only 8 % of them are planning to do so within 12 months and less than half (39 %) of those planning to move say they have already started making preparations (Esipova et al. 2011). That is less than 1 % of the world population. Most individuals stay in what Esipova et al. call the *dream stage* and do not continue to the planning stage and preparation stage. The Gallup World Poll also found that emigration is selective. Adults with at least some secondary education tend to be more likely to want to go than those with less education. Employment status and job prospects also matter. Personal circumstances (finance, family situation) are important too. Most adults are discouraged because of policies that create roadblocks to leaving or entering a country. While age and education strongly relate to people's *desire* to migrate, they do not matter as much in whether potential migrants are *planning* to move in the next 12 months. However, education and employment status are important factors in the transition from planning to *preparation*. The most educated are twice as likely to start preparation than those in other education groups. Employed persons planning to migrate are much more likely to start preparation than those not employed. The Gallup study reveals that the majority (54 %) of people with professional skills planning to migrate also prepare to leave. This may be a consequence of employer-generated international migration.

The Gallup study is the first major study that provides empirical support for the process character of the migration decision. The study distinguishes stages that are close to the stages distinguished in the process model of the TPB. Therefore the Gallup survey is used to generate parameters of the process model.

A major finding of the Gallup study is the low level of emigration. Other studies found similar low values of emigration. Abel and Sander (2014) estimated that the volume of global international migration flows declined from 7.5 per thousand of the world population in the 5-year period 1990–1995 to 5.7 per thousand of the world population during 1995–2000. Since the year 2000, the global 5-year emigration rate remained stable around 6 per thousand. The estimates are based on data on lifetime migrants, which are persons living in a county different from their country of birth (foreign-born population). That means that annually a little over 1 per thousand of the population emigrates (between 1.1 and 1.5 per thousand). That figure is an average. The emigration rate is larger in some countries. For instance, in Europe in 2012, 5 per thousand of the population emigrated (Eurostat 2014). Most went

to another country of Europe. Using data from the Migration between Africa and Europe (MAFE) project, Lessault and Flahaux (2013) found that the emigration rate of Senegal is 0.7% (see also Beauchemin et al. 2014, p. 6). To be valid, the process model of the TPB should be able to produce an overall (global) emigration rate that is below 0.15% per year.

A comparison of the emigration produced by the process model and the rate obtained by Abel and Sander (2014) is not straightforward because the definitions of the emigration rate differ. In the process model, the average emigration rate is the ratio of the total number of emigrations and the total person-years of exposure to the risk of emigration for all cohort members combined. It is an occurrence-exposure rate. Exposure starts at age 15 and ends when an individual emigrates, drops out of the decision process or reaches age 50. Abel and Sander (2014) define the emigration rate as the number of emigrants during a period of 5 years divided by the total population in the country of origin, irrespective of the risk status. In the process model, their definition is approximated by the ratio of the number of emigrants between ages 15 and 50 and the person-years lived in the country of origin between 15 and 50, irrespective of the risk status (at risk, i.e. involved in the decision process, or not at risk, i.e. having dropped out). If the population is a stationary population, then the person-years lived between ages 15 and 50 is equal to the population aged 15–50 (Preston 1982). The comparison of the emigration rate produced by the simulation model and the emigration rate Abel and Sander estimate holds if the emigration rate of those 15–50 does not differ much from the emigration rate of the entire population.

The United Nations published data on population by country of residence and country of birth. The data show that 3% of the world population is foreign-born (IOM 2010). That percentage remains remarkably stable in time. In the model, the foreign-born population is approximated by the proportion of 50-year olds in the world that is born in a country other than their country of residence. The model starts with a cohort of 15-year olds. To pass the test of validity, the model should be able to produce the outcome that about 3% of the cohort of 15-year olds emigrates before age 50.

The typical age profile of migration is a skewed distribution with migration rates increasing rapidly at young ages, a peak at an age between 25 and 30, and declining more slowly than they increased. The shape has been described by a double exponential distribution (Rogers and Castro 1981; Raymer and Rogers 2008). The process model should produce an age profile of emigration that is close to the shape of a double exponential distribution. If a process model of the theory of planned behaviour results naturally in a shape that resembles a double exponential distribution, then the Rogers-Castro model age profile of migration can be given a behavioural interpretation. That would be an important bonus of the process model of emigration. It would replicate for migration what Coale and McNeil (1972) found for first marriage: a process model of first marriage gives a behavioural interpretation to the age profile of first marriage (see also Billari et al. 2007).

The distribution of ages at emigration is compared to the mathematical representation of the migration age profile, developed by Rogers and Castro (1981). In this

paper, the model schedule is limited to ages from 15 to 50. In this age category, the age profile of migration may be described by a double exponential distribution:

$$m(x) = c \exp [-\alpha (x - \mu) - \exp (-\lambda (x - \mu))]$$

where x denotes age, $m(x)$ is the proportion of emigrants that is aged x , and c , α , λ and μ are parameters to be estimated from data. The parameter c is a scaling factor, μ controls the location of the peak of the migration age profile, λ reflects the steepness of the ascending side and α represents the steepness of the descending side. If $\lambda > \alpha$, the location of the peak (mode) is larger than μ ; it is smaller if $\lambda < \alpha$.

The parameters c , α , λ and μ are estimated from the realizations of X_e , i.e. from the ages at emigration in the virtual population. The parameters then enter the above equation to determine the model age profile. To be valid, the process model of the TPB should produce an age profile of emigration that is close to the model age profile and exhibits the typical pattern observed in migration age profiles around the world.

10.5 Parameters of the Process Model of Planned Behaviour

The aim of the process model of the theory of planned behaviour is to describe an emigration decision process that produces a realistic macroscopic (population-level) pattern of emigration. In this section, the parameters of the process model are presented. The parameters are guesstimates. No statistical technique is used to estimate the parameters from data because the necessary data are missing. The validity of the model and the guesstimates are determined by how well the model reproduces the stylised facts presented in Sect. 10.4. In order to produce the stylized facts, the range of most parameter values is limited.

Consider a virtual cohort of 100,000 15-year olds. The cohort is followed until its members reach age 50. The focus is on the emigration decision. Emigration is the endpoint. Individuals differ in skill level: 80 % have a low/medium skill level and 20 % are highly skilled. Skill level has an important effect on the susceptibility to considering emigration and the outcome of the assessment of benefits and costs of emigration (deliberation). In order to reproduce the proportion of persons aged 15+ that desires to emigrate, observed in the Gallup study (14 %) and the selection effect of education, we must determine what a desire to emigrate means in the process model of the TPB. I assume that an individual has a desire to emigrate if emigration is considered beneficial, which means that the individual has completed the formation of behavioural beliefs and developed a positive attitude towards emigration.

Assume that 25 % of all individuals with low/medium skills develop an interest in emigration, i.e. consider emigration. Of these, 48 % believe that the benefits of emigration exceed the costs. Hence 12 % of the individuals with low/medium skills consider emigration beneficial; they develop a positive attitude towards emigration.

Among highly skilled individuals, 48 % develop an interest in emigration and, of those, 50 % believe that benefits exceed costs. Hence, 24 % of the highly skilled develop a positive attitude towards emigration. The difference reflects the observation in the Gallup study. The proportion of cohort members that develop an interest in emigration is 14.4 %. It results from 12 % for individuals with low/medium skills and 24 % for highly skilled individuals ($80 \cdot 0.12 + 20 \cdot 0.24$).

The age at which an individual develops an interest in emigration is the de facto onset of the emigration decision. The minimum age is fixed at 15. I assume that the range of ages at which individuals with low/medium skill level develop an interest in emigration is larger than that of the highly skilled. The ages at onset of the emigration decision follow a one-sided truncated normal distribution (see Sect. 10.3). For individuals with low/medium skill level, the mean of the original, untruncated normal distribution is 19 years and the standard deviation is 4 years. I assume that it takes on average 2 years to determine a belief in the benefits and costs of emigration and to develop a positive attitude (desire) or negative attitude towards emigration, irrespective of the skill level. The formation of that belief shifts the distribution 2 years to the right. The mean age of the truncated distribution that results is 21.6 years. The age at which highly skilled individuals consider emigration is 20 years, on average, with a standard deviation of 2 years. The mean age at which individuals develop an attitude towards emigration is 22 years. The mean of the truncated distribution is also 22 years. The difference in standard deviation of the normal distributions implies that highly skilled individuals develop an interest in emigration in a narrower age range than individuals with low/medium skill level.

The distributions of age at developing an attitude towards emigration (desire or no desire) are shown in Fig. 10.3. Individuals who do not consider emigration stay in the country of birth. We need to consider the age at which the decision not to consider emigration is made because, at the time individuals decide to stay, they are no longer at risk of considering emigration.

Many individuals who consider emigration beneficial do not intent to emigrate. They drop out of the decision process because of low levels of SN and PBC. SN and PBC determine the probability of continuation or dropout, and the timing of these transitions. A beta distribution describes the distribution of SN in the population. The distribution is defined on the interval $[0, 1]$. It has two positive shape parameters, α and β . The distribution is symmetric if $\alpha = \beta > 1$, has a positive skewness if $\alpha > \beta > 1$ and a negative skewness if $1 < \alpha < \beta$. In this paper, it is assumed that individuals with low/medium skill level are more likely to have a lower SN score than individuals with high skills, which means that the skilled individuals have a stronger social pressure to emigrate and are more willing to comply. Among the people with low/medium skill level, few individuals have very low SN scores, most have a low to moderate SN score and the prevalence of SN scores in the population declines with increasing SN score. That shape is described by a beta distribution with shape parameters $\alpha = 3$ and $\beta = 5$. Highly skilled have a higher SN score but not much higher. The shape is described by a beta distribution with $\alpha = 4$ and $\beta = 5$. The mean SN score is 0.37 for individuals with low/medium skill level and 0.44 for individuals with high skill. Figure 10.4 shows the distribution of SN scores, by skill level.




Fig. 10.3 Age at developing a desire to emigrate, by skill level

The PBC score differs by skill level, with highly skilled individuals having a higher score. Instead of using a unipolar 7-points scale, I use a bipolar scoring from -50 to $+50$ such that large negative values represent low scores of PBC and large positive values high scores of PBC. Using bipolar scoring has implications for the value-expectancy model (Ajzen 2006), but they are beyond the scope of this paper. Bipolar scores are converted to values between zero and one using a logistic distribution (logit model). Individuals with low/medium skill level are assumed to have a PBC score that is normally distributed around -10 , with a standard deviation of 8, which implies a considerable spread in the population with low/medium skill level. Van Dalen et al. (2005) provide evidence for the effect of education on PBC score. No literature could be found that addresses differences in spread. The PBC score of highly skilled is normally distributed with mean 10, and standard deviation 5. Figure 10.5 shows the cumulative distribution of the PBC scores for each of the skill levels and for the total population (tick line). With each PBC score is associated a probability, which transforms the original score to a value between 0 and 1. That value is used in further calculations. The figure may also be used to derive the original score from the score on the 0–1 scale. For each skill level,




Fig. 10.4 Distribution of subjective norm (SN) in population, by skill level

the original score associated with a score on the 0–1 scale (empirical distribution) can be retrieved using the quantile function. Figure 10.6 shows the distribution of the PBC score in the population, by skill level. The two shades of grey show the number of individuals with a given PBC score, by skill level. The dark shade is a result of combining two shades of grey. The solid line shows the total number or individuals in the population with a given PBC score. It represents the sum of the number with low/medium skill and the number with high skill. The distribution of PBC scores in the population is a mixture of two normal distributions.

The effects of SN and PBC are combined into a single score, using a Cobb-Douglas utility function. The parameters are: $\gamma=1$, $\alpha=0.6$ and $\beta=1-\alpha=0.4$. The composite score is denoted by V . The distribution of SN and PBC scores in the population and the parameters α , β and γ produce an average composite score of 0.33 for the individuals with low/medium skill level (median 0.32), 0.56 for the




Fig. 10.5 Cumulative distribution of PBC scores, by skill level

highly skilled (median 0.56), and 0.40 for the two skill levels combined (median 0.40). Twenty five percent of the virtual population has a score less than 0.27 and 25% has a score of 0.53 or larger. The highly skilled are much more likely to have a high composite score than those with low/medium skill level. The distribution of the composite score by skill level and for the total population is shown in Fig. 10.7.

The composite score V determines whether an individual who considers emigration beneficial develops an intention to emigrate or drops out of the decision process. An intention is developed if the composite score V exceeds a threshold value. We have no data on the threshold value of V , but we have information on the proportion of individuals with a desire to emigrate that plan to emigrate. The Gallup survey of 2005 revealed that 8% of those who desire to emigrate also plan to emigrate. That figure may be used as a proxy of the proportion developing an intention to emigrate. When that figure is used, the overall emigration rate and the proportion of




Fig. 10.6 Distribution of PBC scores in the population, by skill level

lifetime migrants (proportion of foreign-born population in the world) are greatly underestimated, however. The estimates improve substantially if 30–40 % of those who desire to emigrate develop an intention to emigrate. Therefore, 40 % is used. The Gallup study reports the proportion *planning* to emigrate. That is likely to be a fraction of those intending to emigrate. The threshold value is derived from that proxy of the proportion of individuals with a desire to emigrate that intends to emigrate. Note that the proportion is the share of the total population. The proportion is not distinguished by skill level. For the proportion developing an intention to be 40 %, the threshold value of V needs to be 0.45, given the distribution of V in the population.⁴ Forty percent of the population has a threshold value of V that is 0.45 or higher. The threshold is 0.67 if only 8 % develop an intention. In the simulation, it is

⁴The median value of V is 0.40.




Fig. 10.7 Distribution of composite score V in the population, by skill level

assumed that individuals who consider emigration beneficial and have a composite score V of at least 0.45 develop an intention to emigrate. Individuals with a lower score drop out. The threshold value is the same for everyone. Figure 10.7, which shows the distribution of V scores in the population, also shows the threshold value of V .

The time people need to determine SN and PBC scores is given by Eq. 10.2, with θ equal to 1. The waiting time is equal to θ if the scores are $V = 0.03$ or $V = 0.77$. If all individuals would have a score of 0.77, then they would need not more than 1 year to determine that SN and PBC scores are sufficiently high to produce an intention to emigrate. If all have a score of 0.03, they need an average of 1 year to determine that the SN and PBC scores are low and that the only reasonable option is to abandon the decision process and stay in the country of birth. Individuals who need 1 year or less to determine the SN and PBC scores

have either a very low or a very high V score. They have very limited support and lack self-efficacy or they combine strong self-efficacy with full support. The ultimate emigration rate is sensitive to the θ value. A higher value of θ suppresses the emigration rate without affecting the number of individuals emigrating in a lifetime (before age 50). The reason is the increase in the person-years at risk of emigration with an increase in the waiting time in the assessment stage (a). Some individuals move on to develop an intention to emigrate, but most drop out of the decision process. Given the distribution of V in the virtual population, the average time people need to determine SN and PBC scores can be obtained directly from the distribution without first running the simulation (using Eq. 10.2). The average duration is 2.32 years (2.45 years for individuals with low/medium skill level and 2.07 years for highly skilled). Forty percent develop and intention to emigrate and 60% drop out. Recall that this proportion is not an outcome of the model, but an input parameter. The average durations and the exit rate reported in the results section are slightly different because of random factors.

Individuals intending to emigrate, i.e. who are in the planning and preparation stage, emigrate or decide to stay. Actual behavioural control moderates the effect of intention on behaviour. Sheeran et al. (2003) assert that the performance of intentions as predictors of behaviour depends on the difference between the perceived (PBC) and the actual behavioural control (ABC). In the model, the ABC is the PBC plus a random factor. If we assume that most people overestimate their self-efficacy and the support they will get, and underestimate the barriers to emigration, then the random factor is negative for most individuals. Some individuals have an ABC score that exceeds their PBC score. The ABC score is a value between 0 and 1. In order to assure that the addition of a random factor does not violate that condition, a random factor is added to the logit of PBC, resulting in the logit of ABC. The random factor is a draw from the uniform distribution $U[-7.0,1.0]$. A decrease in the value of the logit of ABC reduces the value of ABC. The rate of leaving the planning and preparation stage, $\mu_{i+}(ABC)$ depends on the level of actual behavioural control (ABC). People complete the planning and preparation, and emigrate, or they abandon the process and stay. The proportion that emigrates depends on the emigration rate $\mu_{ie}(ABC)$ and the drop-out rate $\mu_{ic}(ABC)$, which depends on ABC too. The rates of emigration increase exponentially with ABC level (Eq. 10.3). The parameter a_e is 0.1 and b_e is 2. The following parameter values are used in the equation that determines the rate of dropout (Eq. 10.4): $a_c = 0.01$ and $b_c = 2$. These values are selected because they give a plausible bathtub shaped hazard function relating the hazard rate (emigration or drop-out) to ABC levels. The time at exit from the intention stage and the direction of exit (emigration or stay) depend on these parameter values. The time at exit from the intention stage is a random duration drawn from an exponential distribution with parameter $\mu_{i+}(ABC)$. The probability that an exit from the intention stage results in emigration is $p_i = \mu_{ie}(ABC) / \mu_{i+}(ABC)$. The probability that an exit results in a drop-out is $1 - p_i$. An individual emigrates if a random draw from a Bernoulli distribution with parameter p_i gives a 1. If the draw results in a 0, the individual drops out.

10.6 Results

The population consists of a virtual cohort of 100,000 individuals aged 15, 80,204 with a low/medium skill level and 19,796 highly skilled. Highly skilled means having a potential to develop high skills.

(a) Considering emigration

Among the 100,000 individuals, 14,396 consider emigration beneficial, which is 14.4 %. It is the percentage that develop an interest in emigration AND develop positive attitude towards emigration.

Figure 10.3 shows the ages at which individuals end the assessment and develop a desire to emigrate or drop out of the decision process. The age distribution is fixed by the assumption of normal distributions with given parameter values.

On average an individual considers emigration at age 21.5 if he/she has a low/medium skill level and at age 22.0 if he/she is highly skilled. One of eight cohort members, who ever consider emigration, consider emigration before age 18. It is higher among those with low/medium skill level: 1 in 6 versus 1 in 44 among highly skilled. Of the 18-year olds, who consider emigration, 94 % has a low/medium skill level and 6 % belongs to the category of highly skilled, i.e. follows a trajectory that results in high skills.

(b) Assessment of pros and cons of emigration

The proportion of individuals with an interest in emigration that considers emigration beneficial is fixed at 50 %. Hence, half continue to the assessment stage and half abandon the decision process. Among the 100,000 individuals, 14,396 desire to emigrate. It is the percentage reported in the Gallup study. The small difference between the sample value and the theoretical value (14.4; see Sect. 10.5) is an outcome of the random mechanism. Among individuals with low/medium skill level, 9595 consider emigration beneficial. Among the highly skilled, it is 4801.

The time individuals take to assess the pros and cons of emigration follows an exponential waiting time distribution with rate 0.5. On average an individual takes 2 years to determine whether emigration is beneficial.

(c) Intention to emigrate

The 14,396 individuals who consider emigration beneficial develop an intention to emigrate if their subjective norm (SN) and perceived behavioural control (PBC) are sufficiently high. If the scores are insufficient, they decide to stay and drop out of the emigration decision process. Individuals with low/medium skill have an SN score of 0.37, on average, and highly skilled 0.44. The average bipolar PBC score is -3.34 . It is -10.04 for those with a low/medium skill level and 10.07 for highly skilled. The PBC scores on the (0,1)-scale are 0.34 and 0.82, respectively. The population average is 0.50. The SN score and the PBC score are combined in the composite V score. The average V score is 0.40. It is 0.33 for individuals with low/medium skill level and 0.56 for highly skilled individuals. An individual

develops an intention to emigrate if the V score is at least equal to 0.45, which is the threshold value. The threshold value is the value of V at which 40 % of the people with a positive attitude towards emigration develop an intention to emigrate. The number of individuals with a V score of at least the threshold value is 5759, 1952 with a low/medium skill level and 3807 highly skilled. The majority of the 14,396 individuals with a desire to emigrate drop out of the decision process (8637), 7643 with low/medium skill level and 994 highly skilled.

Of the 14,396 individuals who consider emigration beneficial, 5759 individuals develop an intention to emigrate and 8637 drop out during the assessment stage. Hence, 5.8 % of the initial cohort develops an intention to emigrate. The majority of those who believe that they will benefit from emigration abandon the emigration decision because of inadequate SN and PBC scores. The duration in the Assessment stage at exit is 2.31 years. It is slightly larger for individuals with low/medium skill level (2.4 years) than for highly skilled (2.1 years). In other words, an individual with low/medium skill level and with a desire to emigrate takes a little longer, on average, to decide between emigration (intention) and stay than an individual with high skill level. The effect of skill level on the time it takes to reach a decision is due to its effect on SN and PBC scores. The difference has an effect on the age at which an individual develops an intention to emigrate or decides to stay. The mean age is 24.0 and is the same for individuals with low/medium skill level and highly skilled individuals. The average duration of stay in the attitude stage determines the rate of exit from that stage. The average exit rate is 0.432. It is equal to $1/2.31$. The rate of developing an intention to emigrate is much lower. It is 0.173. In its estimation, dropout is treated as censored observations. The distribution of ages at developing intentions to emigrate or to drop out is shown in Fig. 10.8.

(d) Emigration

Of the 5759 individuals who develop an intention to emigrate, 4338 do emigrate, 1336 with low/medium skill level and 3002 highly skilled. That is 75 % of those who intend to emigrate and 4.4 % of the initial cohort. The remaining individuals (1421) drop out during the intention stage, 616 with low/medium skill level and 805 highly skilled. Individuals who do emigrate have a high self-efficacy and have the resources to remove barriers and take advantage of opportunities. Their ABC score does not differ much from their PBC score. The model therefore accurately captures the theory, which states that a weak performance of intentions as predictors of behaviour is the discrepancy between perceived behavioural control and actual behavioural control. Table 10.1 shows the average PBC and ABC scores, by skill level, for emigrants and individuals who drop out during the intention stage. The ABC score of emigrants is close to the PBC score, which means that they were able to accurately predict the actual behavioural control. This applies to all emigrants irrespective of their skill level. Individuals intending to migrate but deciding to stay have a ABC score that is much lower than their PBC score. They have lower self-efficacy than they expected initially, they cannot mobilize the resources and support needed for emigration, or the barriers are much larger than expected.




Fig. 10.8 Age distribution of developing intention to emigrate or to drop out of the assessment stage, by skill level

Table 10.1 PBC and ABC scores, by skill level and migrant status

Skill level	Migrant status	PBC	ABC
Low/medium	Emigrant	0.531	0.222
	Stayer	0.307	0.086
High	Emigrant	0.837	0.448
	Stayer	0.799	0.137
Overall		0.500	0.313

What is the probability that an individual who intends to emigrate and who has an ABC score of 0.313 (average ABC score) emigrates? Emigration and dropout are competing risks. If the individual leaves the intentions stage before age 50, then the probability that the exit is because of emigration rather than dropout is computed using Eqs. 10.3 to 10.5. It is $100 * m_{ie}(ABC) / (m_{ie}(ABC) + m_{ic}(ABC)) =$

$100 * 0.187 / (0.187 + 0.040) = 82.4$ percent. The probability that the individual drops out of the decision process is 17.6 %. We may also obtain the probability of emigration for an individual with a desire to emigrate and a given SN score and bipolar PBC score. Suppose the individual has SN score of 0.5 and PBC score of -5 . The PBC score on the probability (0,1) scale is 0.491. The composite score V is 0.496, which is larger than the threshold value of 0.446. Hence the individual intends to emigrate. The ABC score is the PBC score with a random factor. The ABC score is 0.116. The emigration rate is given by Eq. 10.3 and the probability of emigration is $100 * m_{ie}(ABC) / (m_{ie}(ABC) + m_{ic}(ABC)) = 0.126 / (0.126 + 0.059) = 68.3$ percent. The probability of dropout is 31.7 %.

The proportion of the initial cohort that emigrates between ages 15 and 50 is 4.4 %. That figure is higher than the 3 % of the world population that is living in a country other than the county of birth.

The average emigration *rate* is 3.6 per thousand if the emigration rate is defined as the ratio of the number of emigrations during a period and the person-years exposed during that period to the risk of emigration. The average number of years a cohort member is at risk of emigration between ages 15 and 50 is 12.0 years. It is the number of years between age 15 and emigration or dropout of the emigration decision. Individuals who leave the country or drop out are no longer at risk of emigration. The emigration rate is therefore $0.044/12.0 = 0.0036$ or 3.6 per thousand. It is higher for the highly skilled population (11.7 per thousand) than for those with low/medium skill level (1.4 per thousand). This rate is an occurrence-exposure rate; it relates the number of occurrences to the duration of exposure. The occurrence-exposure emigration rate is larger than the emigration rate estimated by Abel and Sander (2014) (1.2 per thousand). Abel and Sander use a different type of rate and a different unit interval. They define the emigration rate as the ratio of the volume of migrations during a period of 5 years and the size of the world population. Their estimate of a 5-year emigration rate is 6 per thousand, which is 1.2 per thousand per year. The population they consider is the total population in the country of origin, irrespective of the risk status. To obtain the equivalent of the emigration rate defined by Abel and Sander, the probability of emigration should be divided by the average number of years an individual spends in the country of origin between ages 15 and 50 (before emigration or reaching age 50). The average number of years spent in the country between 15 and 50 is 34.0 years. The emigration rate is $0.044/34.0 = 0.00127$ or 1.27 per thousand. Using the definition of emigration rate adopted by Abel and Sander, the emigration rate produced by the process model is very close to Abel and Sander's estimate. The comparison is not without problem. Abel and Sander considered the total population of the world, including children and persons over 50. The comparison holds if the rate at which individuals aged 15–50 emigrate is close to the average emigration rate of the entire population, an assumption that is not implausible.




Fig. 10.9 Age distribution of emigration and drop out of intention stage, by skill level

(e) Age profile of emigration

The age profile of emigration and dropout of the intention stage, by skill level, is shown in Fig. 10.9. In the intention stage, more individuals emigrate than drop out. The difference between the two exits is higher for the highly skilled than for those with low/medium skill level.

Figure 10.10 shows the distribution of ages at emigration by skill level, produced by the simulation model and the model emigration schedules estimated from these ages at emigration. The dark line presents the outcome of the process model. The light line shows the model migration schedule estimated using the double exponential distribution, estimated from the simulated ages at emigration. The process model yields a description of the emigration age profile that is very similar to that described by the double exponential distribution. The emigration profile peaks between ages 25 and 30, which is realistic, and the shape resembles the




Fig. 10.10 Emigration age profile: simulated data and model migration schedule

shape of migration age profiles reported in empirical studies. Note that the age pattern is not determined by age-specific migration rates but is determined entirely by the parameters that govern the waiting time distributions of the transitions between the stages of the emigration decision process. The double exponential distribution is known to emerge as a convolution of a normal distribution and a set of exponential distributions (Coale-McNeil 1972). In the process model, the waiting time distributions are more complex than the exponential distribution, although the exponential distribution serves as a basis. The fact that the model produces the typical age profile of migration is important for the validity of the model.

The mean age at emigration is 27.7 years, 27.3 years for individuals with low/medium skill level and 28.4 years for highly skilled. The mean age of individuals who drop out during the intention stage is higher: 29.3 for those with low/medium skill level, 29.4 years for highly skilled and 29.3 for the two skill levels combined.

Figure 10.11 shows the ages at transition between stages of the process model of the TPB, by skill level. Many people develop an interest in emigration and a desire to emigrate. Many individuals with low/medium skill level drop out in the first stage of the decision process. Many highly skilled individuals with a desire to emigrate develop an intention to emigrate and start planning and preparation. Many of them drop out while planning and preparing, however, because the actual behavioural control (ABC) is lower than the perceived behavioural control (PBC).




Fig. 10.11 Ages at transition between stages of the process model of the TPB, by skill level

10.7 Conclusion

The theory of planned behaviour states that intention is the best predictor of behaviour. The theory is widely used to explain and predict behaviour. Many people who consider emigration beneficial never intend to emigrate because they believe that others do not want them to leave the country or they believe that they do not have sufficient resources and cannot generate the necessary social support. The simulation model presented in this paper describes the stages of the emigration decision process an individual goes through when considering emigration. The model operationalizes the process theory of planned behaviour. The process model of the TPB is consistent with the ‘horse race’ random utility model. The two models are competing risks models. The information processing and evidence accumulation mechanism is implicit in the process model. Some individuals need more time to

make a decision than other individuals. Individuals who have access to information because of the composition of their social network or because they learned how to access and process useful information efficiently and effectively (and have high levels of self-efficacy) have more extreme PBC or ABC scores than individuals without these networks or capabilities. In the process model, these individuals need less time to choose whether to continue the decision process or to abandon it. The outcome is also easy to predict. The action of individuals with scores close to average values (most individuals) is difficult to predict because small changes in scores may have significant effects.

The model, although simple, reproduces stylized facts remarkably well. It reproduces the typical age profile of emigration. It also reproduces the global international migration rate recently estimated by Abel and Sander (2014). The model slightly overpredicts the proportion of the world population not living in their country of birth, estimated by the United Nations. The reason is return migration, which is not considered in the model. It is remarkable that a simple simulation model of the emigration decision process is able to reproduce these facts of international migration. It makes the model a potentially powerful instrument to help explain and predict international migration. The models in use today for forecasting international migration derive future levels of migration from past levels, statistical associations between levels of migration and characteristics of the population, and expert judgments about changes in reasons for migration and effects of opportunities and restrictions (see e.g. de Beer 2008; Bijak 2011; Raymer et al. 2013; Azose and Raftery 2015). The predictive performance of these models is acceptable most of the time, when conditions do not change abruptly. In the presence of shocks, predictions are poor, probably because more people develop an interest in emigration and fewer people drop out of the emigration decision process. To determine under what conditions people leave their country, predictive models should incorporate individual decision processes.

The model has several limitations that may be removed in future research. It is limited to ages 15 to 50 and excludes child migration and retirement migration. The model considers a single endpoint (emigration) and does not produce a migration history. As a consequence, the emigration intensity cannot depend on migration experience. Return migration and onward migration are not considered either. The main subject of the model is the emigration decision process, not the migration history. The subjective norm (SN) is a characteristic of the individual and is not an outcome of interactions between the individual and significant others and institutions. For instance, the SN score does not depend on the migration experience in one's social network, which is important in most agent-based models of migration (see e.g. Klabunde 2014). SN, PBC and ABC scores are not updated after the occurrence of life events, such as marriage, divorce, job loss, and major health conditions have no direct effects on the emigration decision. Contextual (exogenous) factors, such as political conflict and environmental degradation, are not included in the model. Opportunities and barriers are not modelled explicitly and individuals do not respond to opportunities (e.g. job offers). Individuals do not anticipate events and conditions; they do not predict. They do not learn from experience

either. The model does not consider reasons for emigration and does not distinguish between employment migration, marriage migration, family reunion, and other forms of migration. Including these factors would change the way interests and attitudes develop and SN, PBC and ABC scores are generated, but they would not significantly change the process model. For these reasons, the model is referred to as a simulation model and not an agent-based model. The strength of the model presented in this paper is the operationalization of the theory of planned behaviour in a stochastic decision process model. A process model of the theory of planned behaviour can accommodate all factors and actors that have a significant effect on the emigration decision, but a comprehensive model that includes all these factors is beyond the scope of this paper.

Acknowledgment Icek Ajzen, Anna Klabunde, Matthias Leuchter, the students of the European Doctoral School of Demography 2014–2015, and two anonymous referees read the paper and provided extensive and very helpful comments. I am grateful for their help.

Annex A. Parameters of the Microsimulation Model

Parameters

```

Sample size 100000
Proportion highly skilled 0.20
Probability of developing an interest in emigration,
by skill level 0.25 0.48
Age at developing an interest in emigration: truncated
normal distribution
  Mean of untruncated normal distribution, by skill
  level 19 20
  Standard deviation of untruncated normal
  distribution, by skill level 4 2
Average years it takes to develop an attitude towards
emigration, by skill level 2.00 2.00
Age at decision not to develop a desire but to stay,
by skill level
  Mean of untruncated normal distribution, by skill
  level 27 28
  Standard deviation of untruncated normal
  distribution, by skill level 5 3
Truncated normal: lower bound 15.00 15.00
Truncated normal: upper bound 1000.00 1000.00
Subjective norm (SN): beta distribution
  Shape parameter 1 (alpha) 3.00 4.00
  Shape parameter 2 (beta) 5.00 5.00
Perceived behavioural control (PBC): normal
distribution
  Mean -10.00 10.00
  Standard deviation 8.00 5.00
Composite score V based on SN and PBC

```

```

gamma 1.000
alpha 0.600
beta 0.400
Waiting time in attitude stage: scaling factor theta 1
Proportion of persons in attitude stage that develop
intention to emigrate 0.4
Actual behavioural control (ABC):
Uniformly distributed random factor
Lower bound -7.00
Upper bound 1.00
Transition rate as function of ABC
Emigration rate:
a 0.10
b 2.00
Dropout rate:
a 0.01
b 2.00

```

References

- Aalen, O. O., Borgan, Ø., & Gjessing, H. K. (2008). *Survival and event history analysis. A process point of view*. New York: Springer.
- Abel, G. J., & Sander, N. (2014). Quantifying global international migration flows. *Science*, *342*(6178), 1520–1522.
- Abel, G. J., Samir, K. C., & Sander, N. (2014). *Examining the role of international migration in global population projections*. Paper presented at the annual meeting of the population association of America, Boston, and presented at the joint Eurostat/UNECE Work Session on Demographic Projections, Istat, October 2013. <http://paa2014.princeton.edu/papers/142574>.
- Adsera, A., & Pytlikoa, M. (2012). *The role of language in shaping international migration* (IZA discussion paper no. 6333). Bonn: Institute for the Study of Labor (IZA).
- Ajzen, I. (1985). From intentions to actions: A theory of planned behavior. In J. Kuhl & J. Beckmann (Eds.), *Action control: From cognition to behavior* (pp. 11–39). Heidelberg: Springer.
- Ajzen, I. (2002). Perceived behavioral control, self-efficacy, locus of control, and the theory of planned behavior. *Journal of Applied Social Psychology*, *32*(4), 1–20.
- Ajzen, I. (2006). *Constructing a theory of planned behavior questionnaire*. <http://people.umass.edu/ajzen/pdf/tpb.measurement.pdf>.
- Ajzen, I. (2011). The theory of planned behavior: Reactions and reflections. *Psychology & Health*, *26*(9), 1113–1127.
- Ajzen, I., & Klobas, J. (2013). Fertility intentions: An approach based on the theory of planned behavior. *Demographic Research*, *29*(8), 203–232.
- Alagoz, O., Hsu, H., Schaefer, A. J., & Roberts, M. S. (2010). Markov decision processes: A tool for sequential decision making under uncertainty (tutorial). *Medical Decision Making*, *30*(4), 474–483.
- Almalki, S. J. (2013). *A reduced new modified Weibull distribution*. http://www.researchgate.net/publication/249316479_A_reduced_new_modified_Weibull_distribution.
- Azose, J. J., & Raftery, A. (2015). *Probabilistic projection of international migration for all countries*. Paper presented at the annual meeting of the population association of America, San Diego, April 30–2 May 2015.
- Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, *84*(2), 191–215.

- Beauchemin, C., Sakho, P., Schoumaker, B., & Flahaux, M. L. (2014). *New patterns of migration between Senegal and Europe* (MAFE working paper, 21). <http://mafeproject.site.ined.fr/fichier/rte/29/WP%2021.pdf>.
- Bebbington, M., Lai, C. D., & Zitilik, R. (2006). Useful periods for lifetime distributions with bathtub shaped hazard rate functions. *IEEE Transactions on Reliability*, 55(2), 245–251.
- Belot, M., & Ederveen, S. (2012). Cultural barriers in migration between OECD countries. *Journal of Population Economics*, 25(3), 1077–1105.
- Ben-Akiva, M., McFadden, D., Train, K., Walker, J., Bhat, C., Bierlaire, M., Bolduc, D., Boersch-Supan, A., Brownstone, D., Bunch, D. S., Daly, A., de Palma, A., Gopinath, D., Karlstrom, A., & Munizaga, M. A. (2002). Hybrid choice models: Progress and challenges. *Marketing Letters*, 13(3), 163–175.
- Ben-Akiva, M., de Palma, A., McFadden, D., and others (2012). Process and context in choice models. *Marketing Letters*, 23(2), 439–456.
- Bijak, J. (2011). *Forecasting international migration in Europe*. New York: Springer.
- Billari, F. C., Prskawetz, A., Diaz, B. A., & Fent, T. (2007). The “Wedding-Ring”: An agent-based marriage model based on social interaction. *Demographic Research*, 17(3), 59–82.
- Bilsborrow, R. E. (2012). *Key issues of survey and sample design for surveys of international migration*. Presented at conference on comparative and multi-sited approaches to international migration INED, Paris, 12–14 Dec 2012.
- Boudewyns, V. (2013). *A meta-analytical test of perceived behavioral control interactions in the theory of planned behavior*. PhD dissertation, University of Maryland.
- Burkardt, J. (2014). *The truncated normal distribution*. http://people.sc.fsu.edu/~jburkardt/presentations/truncated_normal.pdf. Accessed 17 Feb 2015.
- Bussemeyer, J. R., & Rieskamp, J. (2014). Psychological research and theories on preferential choice. In S. Hess & A. Daly (Eds.), *Handbook of choice modeling* (pp. 49–72). Cheltenham: Edward Elgar Publication.
- Coale, A. J., & McNeil, D. R. (1972). The distribution by age of the frequency of first marriage in female cohort. *Journal of American Statistical Association*, 67(340), 743–749.
- Cohen, J. H., & Sirkeci, I. (2011). *Cultures of migration. The global nature of contemporary mobility*. Austin, Texas: The University of Texas Press.
- Colonius, A. (2001). Random utility models of choice and response time. In A. A. J. Marley (Ed.), *International encyclopedia of the social and behavioral sciences* (pp. 3291–3296). Amsterdam: Pergamon Press.
- Coulter, R. (2013). Wishful thinking and the abandonment of moving desires over the life course. *Environment and Planning A*, 45(8), 1944–1962.
- Darnton, A. (2008). *Reference report: an overview of behaviour change models and their uses*. GSR behaviour change knowledge review, Government Social Research (GSR), http://www.civilservice.gov.uk/wp-content/uploads/2011/09/Behaviour_change_reference_report_tcm6-9697.pdf. Accessed 13 Feb 2015.
- De Beer, J. (2008). Forecasting international migration: Time series projections vs argument-based forecasts. In J. Raymer & F. Willekens (Eds.), *International migration in Europe. Data, models and estimates* (pp. 285–306). Chichester: Wiley.
- De Groot, C., Mulder, C. H., & Manting, D. (2011). Intentions to move and actual moving behaviour in the Netherlands. *Housing Studies*, 26(3), 307–328.
- De Jong, G. F. (1994) *Choice processes in migration intentions and behavior*. Paper presented at the International Summer School “Choice processes in demography”, The Netherlands Graduate School of Research in Demography, Wassenaar, 6–20 July 1994.
- De Jong, G. F., & Fawcett, J. T. (1981). Motivations for migration: An assessment and a value-expectancy research model. In G. F. De Jong & R. W. Gardner (Eds.), *Migration decision making* (pp. 13–58). New York: Pergamon Press.
- Esipova, N., Ray J., & Pugliese, A. (2011) *Gallup world poll: The many faces of migration. Based on research in more than 150 countries* (IOM migration research series no. 43). International Organization for Migration (IOM), Geneva, in cooperation with GALLUP.

- Eurostat (2014). *Migration and migrant population statistics*. http://ec.europa.eu/eurostat/statistics-explained/index.php/Migration_and_migrant_population_statistics. Accessed 3 Apr 2015.
- Fife-Schaw, C., Sheeran, P., & Norman, P. (2007). Simulating behaviour change interventions based on the theory of planned behaviour: Impacts of intention and action. *British Journal of Social Psychology*, 46(1), 43–68.
- Fishbein, M., & Ajzen, I. (2010). *Predicting and changing behaviour*. New York: Psychology Press.
- Guo, X., & Hernández-Lerma, O. (2009). *Continuous-time Markov decision processes. Theory and applications*. Berlin: Springer.
- Hawkins, G. E., Marley, A. A. J., Heathcode, A., Flynn, T. N., Louviere, J. J., & Brown, S. D. (2014). Integrating cognitive process and descriptive models of attitudes and preferences. *Cognitive Science*, 38(4), 701–735.
- Heckhausen, H. (1991). *Motivation and action*. New York: Springer.
- Heckhausen, J., & Heckhausen, H. (2010). Motivation and action: Introduction and overview. In J. Heckhausen & H. Heckhausen (Eds.), *Motivation and action* (2nd ed., pp. 1–10). Cambridge: Cambridge University Press.
- Heckhausen, J., Wrosch, C., & Schulz, R. (2010). A motivational theory of life-span development. *Psychological Review*, 117(1), 32–60.
- Hess, S., & Daly, A. (Eds.). (2014). *Handbook of choice modeling*. Cheltenham: Edward Elgar Publication.
- IOM (International Organization for Migration). (2010). *World migration report 2010*. Geneva: IOM.
- Ispording, I. E., & Otten, S. (2014). Linguistic barriers in the destination language acquisition of immigrants. *Journal of Economic Behavior and Organization*, 105, 30–50.
- Janis, I. L., & Mann, L. (1977). *Decision making: a psychological analysis of conflict, choice and commitment*. Free Press.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Macmillan.
- Klabunde, A. (2014). *An agent-based decision model of migration, embedded in the life course – Model description in ODD+D format*. Unpublished paper. MPIDR, Rostock.
- Klabunde, A., & Willekens, F. (2016) Decision-making in agent-based models of migration. State of the art and challenges. *European Journal of Population*, 32(1), 73–97.
- Kley, S. (2011). Explaining the stages of migration within a life-course framework. *European Sociological Review*, 27(4), 469–486.
- Kniveton, D., Smith, C., & Wood, S. (2011). Agent-based model simulations of future changes in migration flows for Burkina Faso. *Global Environmental Change*, 21(1), S34–S40.
- Kniveton, D. R., Smith, C. D., & Black, R. (2012). Emerging migration flows in a changing climate in dryland Africa. *Nature Climate Change*, 2(6), 444–447.
- Kōu, A., & Bailey, A. (2014). ‘Movement is a constant feature in my life’: Contextualising migration processes of highly skilled Indians. *Geoforum*, 52, 113–122.
- Lessault, D., & Flahaux, M.-L. (2013). Regards statistiques sur l’histoire de l’émigration internationale au Sénégal. *Revue Européenne des Migrations Internationales*, 29(4), 59–88.
- Marley, A. A. J., & Colonius, H. (1992). The “Horse Race” random utility model for choice probabilities and reaction times, and its competing risks interpretation. *Journal of Mathematical Psychology*, 36(1), 1–20.
- McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 84(1), 113–126.
- Montgomery, M., & Casterline, J. (1996). *Social learning, social influence, and new models of fertility*. New York: The Population Council.
- Moses, J. W. (2015). EMIG 1.2: A global time series of annual emigration flows. *International Migration*, 53(5), 47–60.
- Orr, M. G., Thrush, R., & Plaut, D. C. (2013). The theory of reasoned action as parallel constraint satisfaction: Towards a dynamic computational model of health behavior. *Plos One*, 8(5), 1–11.
- Pindyck, D., & Rubinfeld, B. (2013). *Microeconomics* (Pearson series in economics 8th ed.). Upper Saddle River: Prentice Hall.

- Preston, S. (1982). Relations between individual life cycles and population characteristics. *American Sociological Review*, 47(2), 253–264.
- Prochaska, J. O., DiClemente, C. C., & Norcross, J. C. (1992). In search of how people change. Applications to addictive behaviors. *American Psychologist*, 47(9), 1102–1114.
- Pudney, S. (1989). *Modeling individual choice. The econometrics of corners, kinks and holes*. Oxford: Basil Blackwell.
- Raymer, J., & Rogers, A. (2008). Applying model migration schedules to represent age-specific migration flows. In J. Raymer & F. Willekens (Eds.), *International migration in Europe: Data, models and estimates* (pp. 175–192). Chichester: Wiley.
- Raymer, J., Wisniowski, A., Forster, J. J., Smith, P. W. F., & Bijak, J. (2013). Integrated modeling of European migration. *Journal of the American Statistical Association*, 108(503), 801–819.
- Rodriguez, C. A., Turner, B. M., & McClure, S. M. (2014). Intertemporal choice as discounted value accumulation. *Plos One*, 9(2), 1–9.
- Rogers, A., & Castro, L. (1981). *Model migration schedules* (Research report RR-81-30). Laxenburg: International Institute for Applied Systems Analysis (IIASA).
- Sheeran, P., Trafimov, D., & Armitage, C. (2003). Predicting behaviour from perceived behavioural control: Tests of the accuracy assumption of the theory of planned behaviour. *British Journal of Social Psychology*, 42(3), 393–410.
- Stark, O., Behrens, D. A., & Wand, Y. (2009). On the evolutionary edge of migration as an assortative mating device. *Journal of Evolutionary Economics*, 19(1), 95–109.
- Steele, F. (2005). *Event history analysis*. ESRC National Centre for Research Methods (UK), NCRM Methods Review Papers NCRM/004, September 2005. <http://eprints.ncrm.ac.uk/88/1/MethodsReviewPaperNCRM-004.pdf>. Accessed 1 Feb 2015.
- Townsend, J., & Oomen, C. (2015). *Before the boat. Understanding the migrant journey*. Brussels: Migration Policy Institute.
- Usher, M., Tsetsos, K., Yu, E. C., & Lagnado, D. A. (2013). Dynamics of decision-making: From evidence accumulation to preference and belief. *Frontiers in Psychology*, 4, 1–2.
- Van Dalen, H., & Henkens, K. (2012). Explaining low international labour mobility: The role of networks, personality and perceived labour market opportunities. *Population, Space and Place*, 18(1), 31–44.
- Van Dalen, H., & Henkens, K. (2013). Explaining emigration intentions and behaviour in the Netherlands, 2005–10. *Population Studies*, 67(2), 225–241.
- Van Dalen, H., Groenewold, G., & Schoorl, J. J. (2005). Out of Africa: What drives the pressure to emigrate? *Journal of Population Economics*, 18(4), 741–778.
- Willekens, F. (2013) *Agent-based modeling of international migration*. Research plan for independent research group, Max Planck Institute for Demographic Research, Rostock.
- Xie, M., & Lai, C.-D. (1995) Reliability analysis using an additive Weibull model with bathtub-shaped failure function. *Reliability Engineering and System Safety*, 52(1), 87–93
- Yser, M. C. (2012). Perceived behavioral control in reasoned action theory: A dual-aspect interpretation. *The Annals of the American Academy of Political and Social Sciences*, 640(1), 101–117.

Chapter 11

Deciding to Disclose: A Decision Theoretic Agent Model of Pregnancy and Alcohol Misuse

Jonathan Gray, Jakub Bijak, and Seth Bullock

11.1 Introduction

The case in favour of Agent Based Modelling (ABM) as a general analytical approach has been made numerous and elegantly (e.g. Epstein and Axtell 1994; Resnick 1994; Axelrod 1997; Gilbert 1999; Macy and Willer 2002; Epstein 2014; Silverman et al. 2011, 2013, amongst others). As such we will not belabour the point, and instead turn to addressing some of the concerns expressed about the approach. In this instance we focus on the perception of ABM as ad hoc in nature, reflecting the assumptions of the modeller rather than being empirically or theoretically grounded (Waldherr and Wijermans 2013). To ameliorate this concern, we draw on decision theory to produce simple rule based and learning decision making agents and show that they are able to play a form of signalling game¹ (Kreps and Cho 1987) with a basic form of intragroup social learning. Four decision models of varying complexity and behavioural plausibility are contrasted, by way of demonstrating the significance of the operationalisation of decision making in ABM.

¹In a signalling game, one player (the signaller), has some piece of information that is known only to them which affects the outcome of the game for both players. The signaller has a choice as to what they tell the other player about this hidden information, and the responding player as to what they believe the information to be.

J. Gray (✉) • J. Bijak

Department of Social Statistics and Demography, University of Southampton, Southampton, UK
e-mail: j.gray@soton.ac.uk; j.bijak@soton.ac.uk

S. Bullock

Department of Computer Science, University of Bristol, UK
e-mail: sgb@ecs.soton.ac.uk

This exercise is framed in the context of disclosure decisions, taking drinking patterns in pregnant women as a motivating example. Alcohol consumption in the antenatal period is a significant issue in itself, although there is not a clear consensus on the associated risk. In terms of official guidance in the UK, the National Institute for Health and Care Excellence (NICE) acknowledge that evidence of harm to the fetus is less than conclusive, but advise not drinking at all, or significant moderation (National Institute for Health and Care Excellence 2010a), with similar advice from the UK Department of Health (2008).

Turning more specifically to disclosure of alcohol use by women to healthcare professionals during their pregnancy, research is relatively sparse, although qualitative trends are reported by Phillips et al. (2007) and Alvik et al. (2006). The former explored factors impacting disclosure through a small case study, highlighting the need to build up rapport between woman and midwife over several appointments; the latter compared post partum reports of consumption with contemporaneous accounts, finding apparent under reporting during pregnancy which was amplified by increased drinking. The simulation model described in this chapter is able to replicate both qualitative trends, i.e. an increase in disclosure over appointments, and more honest behaviour by moderate as compared to heavier drinkers.

The resulting scenario is of substantial independent interest, and shows the potential utility of a simulation approach in domains where the process is obscured, here both because of the interest in concealment and obvious ethical concerns. With this said, the lack of a strong quantitative evidence base against which to validate the behaviour of the model augurs for caution in interpreting the results, and is a necessary reminder that in this instance the model is primarily a tool for formalisation of the thought process (Epstein 2008), rather than a machine for predicting.

A game theoretic approach to generating an abstract form of the problem gives a convenient and well known framework to reason about the processes involved in the scenario. While scenarios may map to a number of games, exploring one candidate game still allows for a principled comparison between interpretations, and enforces explicit assumptions. But equilibrium is the sine qua non of game theory, which is concerned with the stable outcome of an infinite contest of second guesses. We wish to see the system in motion rather than just at rest, even if it does eventually settle to some stable point. Instead, we choose to focus on the behavioural processes driving a system in motion, a system out of equilibrium, to understand how these processes interact with the movement. Introducing decision theory takes a step down the ladder of abstraction from the mental chess of game theory. Dealing instead in the mechanics of decision making, and the calculus of choice, allows us to explore not only paths that arrive at the destinations we might consider in game theory, but also avenues not accessible where we constrain ourselves to a sometimes implausible degree of rationality.

This does not preclude a strategic dimension, since decision rules are to a great extent modular, and as demonstrated in this chapter can be exchanged without altering the decision problem. In addition, rules are agnostic as to the source of

information, suggesting room for multi-stage processes – for example, a game theoretic, “model of the opponent’s mind”, approach could act as an information source for a decision rule. As a corollary, the decision problem agents attempt to answer can change, allowing behaviour in novel problems to be informed by beliefs derived under other conditions. This is also indicative of the broader benefits to ABM as an approach. Embedding these abstract rules in a simulated environment allows for mechanics which cannot be readily explored using purely analytic or predictive approaches, for example, the social learning dynamic of the disclosure game model.

While there is no universal theory of human behaviour to sit at the centre of ABM as a method, a key motivation for decision rules is their claim to provide an account of decision making that is behaviourally and cognitively plausible. Their mooted capability in this regard is to some extent supported by work from neuroeconomics, which aims to empirically test theories of decision making (Rustichini 2009). Many key aspects common to decision rules, for example the idea that a common currency is used by the brain to compare outcomes (Padoa-Schioppa and Assad 2006, 2008), are supported by neurological findings. In addition, a single decision rule represents a parsimonious alternative to numerous case specific production rules.

Given these features, the application of decision and game theory to ABM is an attractive approach to computational social science, where the locus of interest is process and decision making. Taking a balance between models focused on replication of low level neurological mechanics, and those with a higher level emphasis where individual behaviours are abstracted away, yields a computationally tractable approach. Despite the relative simplicity, it nonetheless captures some of the nuance and sophistication of human decisions.

The remainder of this chapter proceeds to outline the proposed approach to model development (Sect. 11.2), and experiments (Sect. 11.3), with selected results (Sect. 11.4), followed by a discussion contrasting the decision models (Sect. 11.5), and conclusions (Sect. 11.6).

11.2 Disclosure Game Model

In this section we sketch² the process of moving from a real world scenario to a minimal game which sufficiently captures reality, expressing the result as a decision problem representation, and translating this to a simulation model. We then outline four possible decision rules and, as an example of additional flexibility of process models and simulation in contrast to purely predictive or analytical approaches, extend the model to allow a simple form of social learning.

²A complete example of this for the alcohol misuse in pregnancy model is given in Sect. A.1, with a schedule of simulation provided in Sect. A.2.

11.2.1 Modelling Approach

To model a scenario, we take the approach of first creating a formal game to represent it, capturing the key features as far as possible in the structure of that game. This game is in essence a conjecture about the real data generating process, which can be played out in simulation.

The appropriate game representative of the scenario of interest, which captures the desired strategic dynamics, may not be immediately obvious. We suggest that an iterative process is beneficial, beginning from the simplest possible game, and progressively augmenting it.

Transitioning from the resulting game, to a set of decision problems is a relatively simple task. We treat the n player game as n one-player games (Insua et al. 2009), where the moves of other players are drawn from a probability distribution – nature, in game theoretic parlance. As with the game, the decision problem representation admits a degree of variation, and may need to be adjusted to reflect the decision rules that will be used.

These decision problems may then form the basis of an agent model, where agents use learning and decision rules to play out the game. Simulation can then support features which are not readily representable within an analytic framework, for example, populations of heterogeneous players, individual and social learning, or network effects. In addition, the ability to observe the system in a state of flux rather than at equilibrium is desirable, since even where a social system reaches a stable state, the process by which we arrive at it is significant.

11.2.2 Scenario

Typically in the UK, women have 12 appointments with a midwife during the antenatal period, and in the majority of cases will encounter several different midwives (Redshaw and Henderson 2014) in the course of their care. In the UK, and unlike most healthcare contexts, maternity notes are held by the patient, so midwives do not have extensive information prior to an appointment unless they have encountered the woman previously. Maternity notes are not generally linked to extra-departmental records, meaning that a history of alcohol related admissions to another service may remain unknown unless revealed by the woman.

According to NICE guidance (National Institute for Health and Care Excellence 2010a, 2010b) the issue of substance misuse should be raised at the initial booking appointment, followed by subsequent action if a concern is raised at the discretion of the midwife. This may take the form of specific guidance to reduce intake, or if deemed necessary a referral to a specialist midwife and relevant interdisciplinary team. On alcohol consumption, policy regarding how to determine the level of consumption is at the time of writing generally at the level of the local health authority, hospital trust, or according to the best judgement of the individual

midwife, with no guidance provided by NICE. This commonly takes the form of average units per week, but may include Tolerance, Annoyance, Cut down, Eye-opener (T-ACE)³ (Sokol et al. 1989) and similar measures.

Beyond the “booking” appointment, the onus is on women to raise concerns about their drinking behaviour, or the midwife to probe further if they feel it is warranted. In either case, once a concern has been raised the midwife must respond clinically, and inevitably personally, to the information.

In an ideal world, all interactions with healthcare providers would be immediately and fully disclosive, with no repercussions for the patient. However, alcohol misuse by women is known to attract stigma (Gomberg 1988), and is a recognised barrier to appropriate treatment in the maternity context (National Institute for Health and Care Excellence 2010b; Radcliffe 2011).

11.2.3 *Disclosure Game*

In order to translate the scenario sketched above into a more abstract, tractable form, we cast it as a signalling game, and assume that women’s disclosures (or not), are signals. We also impose a discretisation on the continuum of alcohol use, and use three types of behaviour – light,⁴ moderate, or heavy. Correspondingly, they are limited in what signals they may send when claiming to be one of these three types.

Midwives are treated in a similar fashion, where their type corresponds to how negatively they regard a drinking pattern – non-judgemental, moderately judgemental, and harshly judgemental. The expression of this judgement is not a matter of choice on their part, and is assumed to have no impact on their clinical response, which is to either refer the woman for specialist treatment, or do nothing.

At the end of a game, each player receives a payoff dependent on the actions and types of both players. Because both women and midwives have an interest in the outcome of the pregnancy, and would prefer a healthy baby, the payoff has a common interest component. Hence, both players receive a payoff based on the outcome of pregnancy, but women bear a social cost dependent on the signal they sent and the midwife’s reaction to it. Similarly, midwives pay a cost if they refer to a specialist, mirroring the organisational cost of non-routine care. Table 11.1 shows the three payoff matrices which together describe the game.

As an example, consider the challenge faced by an agent of the heavy drinking type. In order to get the best health outcome, they must be referred and would ideally achieve this without paying any social cost at all. The best move depends on the type and beliefs of the midwife. For example, a particularly unlucky scenario might be

³The T-ACE is a four question screening test for alcohol misuse intended specifically for use with pregnant women.

⁴Or abstinent, the extent of alcohol consumption being such that it would generally be felt to pose essentially no risk.

Table 11.1 Payoff matrices

Midwife	Woman		
	Heavy	Moderate	Light
Harsh	-2	-1	0
Moderate	-1	0	0
Non	0	0	0

a Social cost, X_s , for women, given their signal, and the midwife's type

Midwife	Woman		
	Heavy	Moderate	Light
Refer	10	10	10
Don't refer	-2	-1	10

b Health outcome, X_h , for women and midwives, given the midwife's action, and woman's type

Midwife	Woman		
	Heavy	Moderate	Light
Refer	-9	-9	-9
Don't refer	0	0	0

c Referral cost, X_c , for midwives, given their action and the woman's type

for the midwife to not only be of a harshly judgemental disposition, but to believe that no women really need to be referred (i.e. that all women are light drinkers). Even a relatively weak belief in this possibility can make the honest signal look like an unwarranted risk.

To formally define the game, let $N = \{m, w\}$ be the set of players each with a private type $\theta_i \in \Theta$, and a set of types $\Theta = \{l, m, h\}$, with pure strategies $A_m = \{r, n\}$ and $A_w = \{l, m, h\}$. Here, $\{l, m, h\}$ correspond to light, moderate, and heavy alcohol consumption for women, and non-judgemental, moderately judgemental, and harshly judgemental for midwives. Midwives' pure strategies $\{r, n\}$ are to refer, or do nothing, and those for women are to signal that they have one of the possible drinking patterns. Additionally, we define two utility functions:

$$u_w(s_w, s_m, \theta_w, \theta_m) = X_{s,s_w,\theta_m} + X_{h,\theta_w,s_m} \tag{11.1}$$

$$u_m(s_w, s_m, \theta_w) = X_{h,\theta_w,s_m} + X_{c,\theta_w,s_m}, \tag{11.2}$$

with X_c, X_h , and X_s being the payoff matrices as in Table 11.1, s_w and s_m denoting a specific signal by a woman and referral response by a midwife. Lastly let $p_w(l, m, h)$, $p_m(l, m, h)$ be distributions over types of women, and midwives respectively.

As noted, rather than solve the game, we allow populations of agents to play it, and hence stipulate further that women are drawn in order from a queue of n_w women (where $n_w = 1000$ in all simulations), and play against a midwife chosen at random from a population of n_m ($n_m = 100$). They play for a maximum of r_w rounds ($r_w = 12$ following the routine number of ante-natal appointments in the UK (National Institute for Health and Care Excellence 2010a)) or until they are referred, and a new player is drawn from the same distribution that produced the original players to replace them. If they are not referred, they rejoin the back of the queue after their appointment. In either case, they are informed of their payoff after

each round and update their beliefs accordingly using one of the rules described in Sect. 11.2.5.

Midwives play for r_m rounds ($r_m = 1000$ in all experiments), and conduct appointments in parallel, i.e. if there are five midwives, then five women are drawn from the queue and assigned at random to the midwives. Unlike women, midwives are only informed of their payoff if they choose to make a referral. Both groups of agents have perfect recall, and midwives are assumed to retrospectively update their observations if they make a referral after a number of appointments.

11.2.4 Social Learning

In reality, learning is not exclusively from personal experience, and social learning plays an important role. This social dynamic fits naturally into an agent framework, but is difficult to address without using an approach concerned with process, so we take advantage of this to show a naïve take on it here.

In the disclosure game model, this takes the form of having each midwife recount their play history to their colleagues with some probability q . Individuals then incorporate shared information into their beliefs using weighted updates, e.g. for a midwife a shared observation of a low type signal contributes to their beliefs by w , and $0 \leq w \leq 1$ (i.e. $n_j = n_j + w$). Women share only when they have finished play, and provide their complete history of games, because they have accurate information about the outcomes. By the same rationale, midwives share only their history with the most recent woman they referred. Sharing occurs simultaneously for all players at the end of each round, and all memories are either shared immediately or discarded.⁵ Accounts are shared with some probability to all fellow players. For example, a heavy drinker finishes play having claimed to be a light drinker, without ever being referred, and their account is selected to be shared with some probability q_w with all other women.

Because of their differing problem representations, the simple payoff reasoners and their more complex counterparts incorporate this exogenous information differently. The simple payoff based rule relies on a belief structure relating actions directly to rewards which is essentially model free. Because payoffs differ by the agent's private type, the information shared may not correspond to the experience of the listening agent in the same scenario. As a result, payoff reasoners have a belief bias towards the most common player type, and can believe in outcomes that are, for them, impossible.

A payoff based agent, who is a light drinker, hears the account of the heavy drinker. They take the account as literally happening to them, and update their beliefs to include the possibility that there is a negative outcome attached to claiming to be a light drinker.

⁵More precisely, memories of games remain, but it is assumed that only the most current information is relevant enough to be shared.

By contrast, representing the problem in terms of the probabilities of the individual lotteries imposes a model that abstracts the new information from payoffs, and allows the agent to discard implausible outcomes. This stronger assumption as to the static and known qualities of payoffs does however reduce the flexibility of the decision rule.

Returning to our example, a light drinker using this decision rule would follow the account through from their position in the game tree, correctly inferring that the outcome in their case would be positive.

11.2.5 Agent Models

While in principle a wide variety of agent models are possible, given that decision rules operate on essentially the same information, and produce the same output (a decision), we limit ourselves here to four. The simplest is a lexicographic rule (1), in the spirit of a Fast and Frugal Heuristic (Gigerenzer 2004) which uses only information about payoffs given actions; this is followed by a Bayesian risk minimisation rule (2) using the same information; a second Bayesian risk rule (3) which uses information about the underlying lottery; and a two-stage Cumulative Prospect Theory (CPT) (Hau et al. 2008) agent (4) which is identical to 3, but uses the CPT decision rule (Tversky and Kahneman 1992). Hence, each successive decision model adds a layer of sophistication to the problem representation while retaining the same input-output characteristics.

Agents have perfect recall and midwives recognise women if they repeatedly encounter them, making use of new information for retrospective updates. However, all four agent models make decisions ‘as-if’ they were always facing a new “opponent”.

A simplifying assumption is made that all midwives have just qualified after receiving identical training. As a result, they have homogeneous beliefs about women and assume to some extent that they are honest. Women have heterogeneous beliefs, which correspond to experiencing k randomly chosen paths through the game, and following each path at least once.

11.2.5.1 Lexicographic Heuristic

The lexicographic heuristic (Algorithm 1) follows the form of that used in Hau et al. (2008), and assumes a simplified problem representation, where an action is a choice between combined lotteries. Functionally, the heuristic maintains a count of the number of times that each action was followed by a payoff, and chooses the action which most commonly has the best payoff, i.e. one reason decision making. Where there is no clear best action, but one or more is evidently worse, a choice must be made as to whether to discard the poorer action; in this case we have elected to retain it. This approach requires minimal computation, and does not assume that u_i is static, or known.

Women resolve this by approximating the utility function, as a function $f(s_w, \sigma)$ on their choice of signal and an unknown distribution σ , which maps to u_w – i.e. s_w is a choice between simple lotteries. The algorithm maintains a count, n , of the number of occurrences of each outcome given the choice from s_w .

Midwives solve a slightly different problem with more information, where s_w is known, and s_m is the lottery choice – $f(s_w, s_m, \sigma)$. This is resolved by maintaining a separate count for each signal (i.e. n_{s_w, s_m}), and otherwise following the same Algorithm 1.

Algorithm 1 Lexicographic heuristic

```

n ← 1, action ← none
while action = none do
  Calculate the nth most common outcome following each action.
  Sort actions by the value of the nth most common outcome.
  if clear winner then
    action ← best
  end if
  n ← n + 1
end while
return action

```

11.2.5.2 Bayesian Payoff

The Bayesian payoff agent uses the same subset of information as the lexicographic method, but updates beliefs on the link between actions and payoffs using the Bayes rule, and attempts to choose the action which minimises risk.

Given the discrete nature of actions and payoffs, coupled with a desire for tractability of the simulation, the Dirichlet distribution is employed as a prior to represent these beliefs. The distribution is particularly convenient, in that to infer the probability of a signal implying a payoff is simply:

$$p(x = j|D, \alpha) = \frac{\alpha_j + n_j}{\sum_j (\alpha_j + n_j)}, \quad (11.3)$$

where n_j is simply the count of occurrences of signal-payoff pair j , and α_j is the pseudo-count of prior observations⁶ for a pair j . Hence, the belief that a signal will lead to a payoff is the number of times that pairing has been observed (including the pseudo-count), over the total number of observations thus far. This makes computation of beliefs fast and simple, since all that must be maintained is a count

⁶Pseudo-counts are related to, but distinct from prior beliefs. Here, the pseudo-count is a parameter to the prior belief distribution and is nothing more than a hypothetical count of prior observations.

of observations. As before, midwives follow a similar pattern but maintain n_{s_w} independent counts of pairings between referral choice and payoff, updating their beliefs about the relationship between the choice to refer and payoff given the signal they have received.

Agents then choose the strategy s_i to minimise risk R_i , which is simply defined as:

$$R_w(s_w) = \sum_{x \in X} -xp(x|s_w) \quad (11.4)$$

$$R_m(s_w, s_m) = \sum_{x \in X} -xp(x|s_w \wedge s_m), \quad (11.5)$$

where X is the set of payoffs the agent has observed to follow s_i .

11.2.5.3 Bayesian Risk Minimisation

The second Bayesian agent augments the reasoning of the simple payoff model, making the stronger assumption that the utility function is static and known. Women maintain two sets of beliefs, corresponding respectively to p_m and the probability of referral given signal choice. This leads to the risk function, minimised with respect to s_w :

$$R_w(s_w, \theta_w) = \sum_{i \in A_m} \sum_{j \in \Theta} -u_w(s_w, i, \theta_w, j)p(j)p(i|s_w), \quad (11.6)$$

so that the risk of a signal is the sum of the products of all payoffs with the probabilities of their entailed midwife types and responses.

“Midwives” reasoning centres on determining the meaning of signals, since given the knowledge of what some signal s_w conveys about the true type of the sender, the payoff for an action is known. As such, their inference process is the same as for the simple Bayesian agent but over signal-type pairs, and they attempt to minimise the following risk function, minimised with respect to s_m :

$$R_m(s_w, s_m) = \sum_{i \in \Theta} -u_m(s_w, s_m, i)p(i|s_w). \quad (11.7)$$

11.2.5.4 Descriptive Decision Theory

The most complex decision rule used is CPT, which attempts to reproduce a number of systematic deviations from rationality observed in humans. Rather than risk, ‘prospects’ (i.e. the sequence of payoff-probability pairings in ascending order of payoff associated with an action) are used as decision criteria. While CPT has primarily been applied in the context of decisions from description, it has been

modified to deal with decisions from experience by incorporating a first stage where probabilities are estimates from observations (Fox and Tversky 1998). In this instance the Bayesian inference process fills the first stage role.

CPT uses transformed probabilities underweighting small probabilities and overweighting large ones. This is intended to reflect the observed behaviour of humans, where sufficiently high likelihoods are treated as certain, and contrastingly low probabilities as impossible. The correct weighting function is subject to some debate, but here we have used that of Tversky and Kahneman (1992), which treats probabilities differently for gains (Eq. 11.8) and losses (Eq. 11.9):

$$w^+(p) = \frac{p^\gamma}{(p^\gamma + (1 - p)^\gamma)^{\frac{1}{\gamma}}} \tag{11.8}$$

$$w^-(p) = \frac{p^\delta}{(p^\delta + (1 - p)^\delta)^{\frac{1}{\delta}}}, \tag{11.9}$$

where p is the unweighted probability, and γ and δ are the weights for gain and loss probabilities respectively. Along similar lines, the values of losses and gains are transformed to reflect a tendency to regard a loss as more significant than a gain –

$$v(u_i) = \begin{cases} f(u_i), & \text{if } u_i > 0 \\ 0, & \text{if } u_i = 0 \\ \lambda g(u_i), & \text{if } u_i < 0 \end{cases} \tag{11.10}$$

where

$$f(u_i) = \begin{cases} u_i^\alpha, & \text{if } \alpha > 0 \\ \ln(u_i), & \text{if } \alpha = 0 \\ 1 - (1 + u_i)^\alpha, & \text{if } \alpha < 0 \end{cases} \tag{11.11}$$

$$g(u_i) = \begin{cases} -(-u_i)^\beta, & \text{if } \beta > 0 \\ -\ln(-u_i), & \text{if } \beta = 0 \\ (1 - u_i)^\beta - 1, & \text{if } \beta < 0 \end{cases} \tag{11.12}$$

and α , and β are respectively the power of a gain, and a loss, and λ is a multiplier giving the aversion to loss.

Finally, the transformed probabilities are used to construct decision weights, π^+ , π^- for gains and losses, where,

$$\pi_n^+ = w^+(p_n) \tag{11.13}$$

$$\pi_{-m}^- = w^-(p_{-m}) \tag{11.14}$$

$$\pi_i^+ = w^+(p_i + \dots + p_n) - w^+(p_{i+1} + \dots + p_n), 0 \leq i \leq n - 1 \quad (11.15)$$

$$\pi_i^- = w^-(p_{-m} + \dots + p_i) - w^-(p_{-m} + \dots + p_{i-1}), 1 - m \leq i \leq 0. \quad (11.16)$$

The CPT value of a single outcome prospect $f = (u_i; p_i)$, is $v(u_i)\pi^+(p_i)$ if $u_i \geq 0$, and $v(u_i)\pi^-(p_i)$ otherwise. For any given action the CPT value V is the sum of the value of the prospects of that action, as in the Bayesian risk model, and the agent chooses the option which maximises this quantity.

11.3 Method

This section provides details of experiments conducted to examine the ability of the model to reproduce qualitative trends reported in the midwifery literature by Alvik et al. (2006) and Phillips et al. (2007); as well as a global sensitivity analysis and construction of statistical emulators to explore and contrast the response surfaces of the four decision rules.

11.3.1 Qualitative Trends

Throughout this chapter, parameters for the CPT model were the same as those used in Tversky and Kahneman (1992) (Table 11.2). While there has been significant work on determining appropriate parametrisation for the model (e.g. Neilson and Stowe 2002; Nilsson et al. 2011; Glöckner and Pachur 2012; and particularly Byrnes et al. (1999) and Booij et al. (2009) addressing risk aversion and gender), a full exploration of the impact of these parameters, or heterogeneous values within populations is beyond the scope of this work. For simplicity, it was assumed that all three drinking types are equally prevalent within the population, although results derived from the Avon Longitudinal Study of Parents and Children suggest that the reality is far more positive⁷ (Humphriss et al. 2013). The scenario was biased towards disclosure as the better option by presuming a distribution of midwives strongly skewed towards non-judgemental types, with beliefs initially favouring honesty. Payoffs were as in Table 11.1, which ensure that it is always strictly preferable to refer drinkers and, together with the initial belief that signals will be honest, not to refer those claiming otherwise.

Two key measures were used: the fraction of the subpopulation who had ever signalled honestly and the proportion of the population who were referred. Both measures were taken after every round of play, and were taken relative to the agent's

⁷95.5% of women in the sample reported consumption at, or below, NICE recommended safe levels.

Table 11.2 Model parameters

Name	Description	Value
n_w	Number of women	1000
n_m	Number of midwives	100
r_m	Number of appointments per midwife	1000
r_w	Maximum number of appointments per woman	12
Runs	Simulation runs	1000
$p_w(h)$	Proportion of heavy drinkers	1/3
$p_w(m)$	Proportion of moderate drinkers	1/3
$p_w(l)$	Proportion of light drinkers	1/3
$p_m(h)$	Proportion of harsh midwives	5/100
$p_m(m)$	Proportion of moderate midwives	10/100
$p_m(l)$	Proportion of non-judgemental midwives	85/100
q_w	Probability of women sharing	0
w_w	Weight of shared information for women	0
q_m	Probability of midwives sharing	0
w_m	Weight of shared information for midwives	0
$s_i[a_i] : s_i[a_{-i}]$	Pseudo-count favouring honesty	10:1
γ	Probability weighting for gains	0.61
δ	Probability weighting for losses	0.69
α	Power for gains	0.88
β	Power for losses	0.88
λ	Loss aversion	2.25

position in their sequence of appointments giving the probability of signalling honestly, or being referred having had a given number of appointments.

In addition to assessing the adequacy of the rules in capturing qualitative trends, we also examined the impact of simple social learning within the population of women (Sect. 11.2.4) on the robustness of these trends. The original experiment was repeated at $q_w, w_w \in \{0.25, 0.5, 0.75, 1 | q_w = w_w\}$, with 100 runs under each condition.

11.3.2 Global Sensitivity Analysis

In general, we have followed the example of Bijak et al. (2013) for global sensitivity analysis of stochastic agent based models, although see Thiele et al. (2014) for a review of alternative techniques. For this purpose the Gaussian Emulation Machine for Sensitivity Analysis (GEM-SA) software (Kennedy 2004) was used, which implements the Bayesian Analysis of Computer Code Outputs (BACCO) method developed by Oakley and O'Hagan (2002, 2004), Oakley et al. (2006). This is a form of variance-based sensitivity analysis, which assumes that the model output

Table 11.3 Parameter ranges

Name	Description	Min	Max
$p_w(m)$	Proportion of moderate drinkers	0	1
$p_w(l)$	Proportion of light drinkers	0	1
$p_m(m)$	Proportion of moderate midwives	0	1
$p_m(l)$	Proportion of non-judgemental midwives	0	1
q_w	Probability of women sharing	0	1
w_w	Weight of shared information for women	0	1
q_m	Probability of midwives sharing	0	1
w_m	Weight of shared information for midwives	0	1
x_h	Health payoff for healthy delivery	1	100
x_r	Cost for referral	$-(x_h - 1)$	
$s_i[a_i] : s_i[a_{-i}]$	Pseudo-count favouring honesty	1:1	100:1

is an unknown, smooth function of the inputs. The unknown function can then be approximated as a Gaussian Process, which is fitted to the training data using Bayes' Theorem and then serves as an emulator for the simulator. The emulator is then able to provide an indication about the extent to which uncertainty in a parameter propagates to uncertainty in the output, and how sharply the output responds to change in each parameter.

Parameters for training were generated in R (R Core Team 2014) using an appropriately transformed Latin Hypercube Sample (Carnell 2012) over the space of parameters given in Table 11.3, giving eleven free parameters which were treated as uniformly distributed in the range given. Given the limitation of 400 design points for the GEM-SA software, we produced exactly that many parameter combinations and collected results for 100 runs of each, with emulator quality assessed by leave-one-out cross validation. A fixed set of 100 random seeds was used,⁸ such that each parameter set was run once with each seed, for every decision rule.

To capture the response characteristics for the model, we measured four outcome variables: (1) the interquartile range (IQR) of the average signal sent by each type of agent in a run, (2) the average signal of moderate drinking agents in a run, and (3, 4) the IQR of 1 & 2 between simulation runs. Together these four metrics give an indication of how far women are separable by their signalling behaviour (1), the behaviour of the at risk drinking groups⁹ (2), and finally the variability of the system in response to changes to the parameters (3 & 4).

Measurements were taken at the end of 1000 rounds of play, and emulators were built against 400 sample points from the full set of simulation results (1 & 2), and

⁸Fixed random seeds were used to allow simulation results to be reproducible, since the combination of a parameter set and a random seed yields a deterministic process.

⁹Under most conditions, the behaviour of heavy drinkers tracks closely with their moderate counterparts.

the IQR at each point (3 & 4) to assess both the overall trend, and the extent to which the parameters contribute to variance between runs.

Sixteen emulators were built, covering each of the four outputs on all decision models and used to conduct a probabilistic sensitivity analysis to assess the impact of parameters and interactions.

In addition to the sensitivity analysis, we also employed the resulting emulators to rapidly¹⁰ explore the parameter space. While emulated results are subject to inaccuracy, they do provide an indication which regions of the parameter space are plausible, and yield interesting results. Results for the outcomes of the interactions of $s_i[a_i] : s_i[a_{-i}]$ with x_h , and q_w with w_w are given in Sect. 11.4.3.

11.4 Results

11.4.1 Qualitative Trends

As shown in Fig. 11.1, all four decision rules were able to reproduce both qualitative trends towards more disclosure as women experience more appointments (Phillips et al. 2007), and a greater tendency towards underreporting of consumption by heavier drinkers (Alvik et al. 2006). Trends for all four rules are broadly similar, exhibiting a gradual increase across appointments which subsequently levels off. This levelling can in part be explained by the referral results (see Fig. 11.9 in Sect. A.4), which show that the majority of drinkers are referred, even with substantial concealment. Referrals continue to occur, in the absence of honest signals, because drinkers are able to achieve a referral by masquerading as higher or lower types, dependent on how their initial beliefs are biased. Despite this the results suggest that a minority of risky drinkers will evade detection altogether, with no notable distinction between heavy and moderate types. Under these parameters, light drinkers always signal honestly and are never referred since there is no perceived advantage in doing so, and the evidence of deceptive signalling is insufficient to outweigh the biased priors of the midwives.

11.4.2 Social Learning

Introducing social learning amongst women leads the behaviour of the decision rules to diverge markedly, which we explore possible reasons for in Sect. 11.5. Figure 11.2 shows the proportion of women who have signalled honestly at least once by their final appointment, under four sharing conditions.

¹⁰Once constructed, the emulator has an analytical solution conditional on the roughness parameters, which obviates the need to use MCMC.




Fig. 11.1 Average fraction of population ever signalled honestly by each appointment, after 1000 rounds, mean with 95 % confidence limit over 1000 runs. Note that the large number of runs leads to very tight confidence intervals

Aside from the lexicographic decision rule, the general tendency is towards less honest signalling by heavy drinkers, which is accompanied by a slight increase in referrals for the Bayesian, and CPT rules. For these decision models, this is because social learning exacerbates the existing tendency of heavy drinkers to impersonate moderate drinkers, who behave more honestly as heavy drinkers become less so. This arises because both classes of agent learn that the moderate signal is the lower risk option as it is both a reliable indicator of need, and does not attract strongly negative judgement. The reliability of the signal is self reinforcing, since the more the agents use it and get referred, the more confident midwives become that it indicates need.

Particularly notable, is the decline in honest signalling by light drinkers visible in both heuristic type rules at the 0.25 level of q_w & w_w , which is associated with an increase in false positives. This arises because of the lack of homophily in social learning, as light drinkers become informed about negative outcomes associated

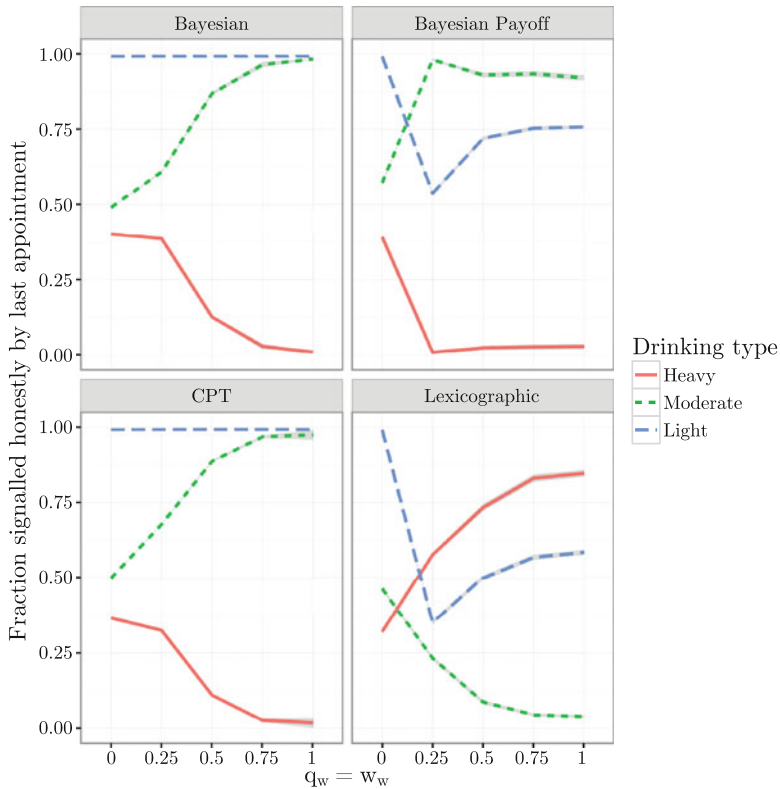


Fig. 11.2 Impact of social learning on trends in the average fraction of population ever signalled honestly by their final appointment, after 1000 rounds, mean with 95 % confidence interval over 100 runs

with concealment, despite having nothing to conceal. The relatively high referral rates of drinkers heighten the effect further, because shared information becomes dominated by their experiences.

The relationship is not, however, entirely straightforward, in that increasing social learning leads to greater variance between runs. A linear model was used to predict the between-runs interquartile range of the average signal sent by moderate drinkers. The predictors used were decision rule and level of social learning, together with the interaction between the two. The regression results were significant ($F_{7,12} = 25, p < 2.9 \times 10^{-6}$) with $R^2 = 94\%$, and intercept 0.07. The only significant coefficients were for the interaction terms, which were 0.44 ($p < 0.05$) for the Bayesian payoff rule and 0.69 ($p < 0.005$) for the lexicographic rule. This suggests that social learning for the heuristic style decision rules introduces considerable uncertainty to the model, which is explored further in the sensitivity analysis below.

11.4.3 Sensitivity Analysis

In this section we present a brief overview of the sensitivity analysis, followed by selected results highlighting the global effect of changes to perceived payoffs and degree of bias towards honesty, as well as social learning within women. The full results for the sensitivity analysis covering all sixteen emulators are available in Sect. A.5.

For the median signal choice of moderate drinkers, the results of the sensitivity analysis suggest that the proportion of light drinkers has a significant effect for all decision rules, accounting for 10, 38, 24, and 5 % of the variance in output for the Lexicographic, Bayesian Payoff, Bayesian, and CPT rules respectively. For the Lexicographic rule, the overwhelming majority of variance in signalling behaviour is reflective of the prevalence of stigmatisation by midwives (44 % $p_m(m)$, 7 % $p_l(m)$, and a further 15 % for their interaction). The proportions of midwives are also key drivers in group separation and the between run IQR of both measures for this rule.

Perhaps surprisingly, variance attributable to social learning between midwives is relatively low, with neither the weight nor probability accounting for more than 5 % of variance in any measure. While there are small contributions to variance in interaction with other parameters (e.g. 4 % to between groups IQR for the interaction with the proportion of light drinkers under the Bayesian rule), this may suggest that the model is lacking in this area, which we touch on in Sect. 11.6.

Figure 11.3 gives a qualitative picture of both emulator quality, and the divergent response surfaces of the decision rules in response to variations in social learning parameters. Emulator fit is clearly imperfect, but overall behaviour is qualitatively similar, with both emulated and simulated plots demonstrating separation in outcome space for the decision rules.

Following from the suggestive results for social learning introducing uncertainty (Sect. 11.4.2), Fig. 11.4 shows emulated points covering the parameter space in high resolution. These plots reflect the increase in uncertainty of outcome shown for the heuristic type rules, which is especially severe for the Bayesian payoff rule. They also suggest that the Bayesian decision rule is less stable under conditions where the weight of shared information is substantially higher than the probability of sharing. This indicates that placing a high weight on information from limited sources leads to greater variability, i.e. what information is shared matters.

For the CPT and Bayesian decision models, the interaction of bias towards honesty and distinction between payoffs has a significant and non-linear effect on instability and separability of groups. Figure 11.5 shows the effects, and also highlights the tendency towards poor separability of groups for both the heuristic type decision rules. The response surface of the Bayesian payoff rule is slightly more nuanced than the simple Lexicographic rule. Figure 11.5 shows better separation, close to partial pooling¹¹ at high payoff distinction, with relatively modest honesty

¹¹Pooling occurs when signallers of different types ‘pool’ their signals, and one adopts the signals of another.




Fig. 11.3 Median moderate drinker signal vs median between drinking type IQR for all decision rules, with signals coded as 0 = light, 1 = moderate, and 2 = heavy

bias, which is reflected by the variance contributions of 11 and 8 % respectively. For the more complex rules, the general tendency is towards less pooling for higher values of both, but with pockets where full pooling¹² occurs. The plots also suggest that the sensitivity of the CPT rule is marginally greater, which is supported by the significant contribution to variance of close to 15 % for all measures of x_h .

11.5 Discussion

From a pragmatic perspective, the differing response characteristics of the classes decision rules are substantial and significant, particularly when social learning is considered. There is a high level of uncertainty in the overall dynamics with the model free rules. This does not arise with the more complex rules, because they

¹²Indicating that all signaller types are using a the same signal.




Fig. 11.4 Emulated moderate drinker signal IQR in response to varying q_w and w_w

reframe information from others in the context of their own experiences, as what would happen to them in that situation. By contrast, the simpler rules treat the experiences as having literally happened to them, and since there is no mechanism of homophily, no way to listen only to accounts of agents similar to themselves, they can come to believe unreasonable things. Naturally, incorporating homophily, by, for example weighting shared information by the type of the sharer, would represent a trivial modification to the heuristic models. While to some extent this highlights the flexibility of the decision rule approach, it would of course sacrifice the parsimony of the model. This is an important consideration, given that part of the argument in favour of a decision theoretic approach lies in the minimal nature of the behavioural rules.

One of the notable features of the results is that the behaviour of rules within a class is very similar. To some extent this reflects poorly on the most complex rule, CPT, which diverges only minimally in behaviour from the Bayesian model. This might be to a degree anticipated since we have not elicited payoffs for obvious practical and ethical reasons, and they may be unrealistic, which limits the utility




Fig. 11.5 Emulated between groups IQR in response to varying $s_i[a_i] : s_i[a_{-i}]$, and x_h

of the CPT approach. Additionally, work by Glöckner and Pachur (2012) has shown that there is considerable variation in individual parameters for the decision model, whereas we have let them remain homogeneous here. In the same vein, utility functions should arguably vary between individual agents, which could potentially be addressed by replacing the fixed payoffs used here with a distribution. With this said, the significant increase in complexity, which entails both additional parameters and increases to simulation time may necessitate a middle ground, particularly where elicitation of payoffs is impractical. This, together with the variability associated with the heuristic type decision rules speaks to a trade off between capturing reality and replicating it.

Continuing the discussion of the issues raised by the representation of payoffs, the temporal aspect is significant, in that there is a timing difference in payoffs, since while the potential social pain of disclosure is immediate, the health outcome comes only later. In light of this, that there is a known impact of time on perceived utility Thaler (1981) suggests that incorporating some form of temporal discounting (e.g. exponential (Samuelson 1937), or hyperbolic (Ainslie 1991)), or a decision model which explicitly treats intertemporal choice, such as the CPT-like model of Loewenstein and Prelec (1992), is warranted.

As noted in Sect. 11.3.2, the impact of social learning in midwives is surprisingly minimal, where it might be expected to play a more significant role in reality. A possible explanation for this lies in the implementation, which may place an excessive constraint on how much information midwives can share. The restriction to sharing only after a referral, together with the disparity in population sizes and random allocation of appointments, leads to midwives rarely having more than a single interaction with woman to pass on to their colleagues. Furthermore, because midwives are only informed of the true type if a referral occurs, they have an inherent myopia since until they have evidence of deception they will not refer, with said evidence difficult to obtain without a referral.

In reality it might be anticipated that midwives would not withhold judgement, and would pass on concerns about specific women to their colleagues, or that particularly dramatic stories would persist and be passed. This might be addressed by incorporating noisy type information (Feltovich et al. 2002), capturing the unintentional information transmitted during appointments, together with a relaxation of the assumptions about when information may be shared and a more sophisticated model of information flow in general. This also highlights an advantage of the BACCO approach (which we describe in Sect. 11.3.2), in diagnosing issues with model design by giving insight into parameters which are contributing inappropriately to variance in output. Coupled with the ability of emulators to rapidly explore parameter space, this clearly suggests that statistical emulation is a powerful tool to support simulation based approaches. As noted in Sect. 11.4.3 the emulators here are indicative, but not definitive. Amongst the reasons discrepancy arises here are heteroskedasticity associated with social learning, the stochastic nature of the simulation, and a lack of precision given the large parameter range. The former issues could be addressed by a more comprehensive approach to setting the nugget, which explicitly incorporates point variance. The latter could be improved through iterative fitting procedures, where the simulation is sampled more heavily in plausible regions of parameter space, a procedure not possible here given the dearth of data to evaluate plausibility. That the discrepancy exists is not prohibitive in this instance, since we are not using the emulator for prediction, only to achieve a broad strokes picture of the behaviour of the simulator.

11.6 Conclusion

The conclusions that can be drawn about the behaviours of real life women, and their midwives, are necessarily limited by the paucity of data available to validate the model. While qualitative trends offer some indication, they are limited in scope, and do not permit strong claims about the drivers of disclosure. As such, further work will focus on applying the model to domains where validation data is more available, which will support a more comprehensive evaluation of the model discrepancy. With this said, the trends reported by Alvik et al. (2006), and Phillips et al. (2007) are borne out by the model, and predictions from the two more complex rules suggest

that encouraging information sharing between women may encourage disclosure, but at the expense of reducing accuracy. By contrast, if one takes the view that a Lexicographic model is a better approximation of real behaviour, then outcomes can best be influenced by controlling how far midwives punish their women socially. We would however suggest that there are better reasons than the outputs of a simulation for doing so.

More broadly, the results demonstrate the logistical feasibility, and its utility as a ‘tool for thinking’, of an agent model grounded in decision theory. The results also make clear that deciding the operationalisation of the decision making is of key significance.

Acknowledgements This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) grant EP/H021698/1 Care Life Cycle, funded within the Complexity Science in the Real World theme. The authors also gratefully acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

Appendix

A.1 Disclosure Game Model Development

This appendix provides a more in depth exploration of the model development process, beginning by deriving a game to serve as the basis for the model and decision problems.

A game, in the game theoretic sense, can be any interaction where the result for one person is dependent on the actions of another. In this scenario, the result for the woman would seem dependent on whether the midwife chooses to refer her for specialist support (although naturally the reality can only be thought of in terms of risk mitigation), and conversely, the right choice for the midwife is somewhat contingent on what the woman is willing to tell them.

A very simple way to represent this would be a game with two players, who both have two possible moves – ask for help, or not; and refer, or not (Fig. 11.6). Since both parties are invested in the outcome of the pregnancy, we might allow them to share the same payoff if everything ends well.

The first complication, is that there should be differentiation between referring, and doing nothing because specialist treatment incurs a cost. We can modify the payoffs to reflect this, by reducing the midwife’s payoffs when they refer. If the cost of referring is less than the value of a good outcome, then the effect of this is to make the only rational choice when not asked for help is to do nothing.

This simple game is however not very informative, and clearly neglects much of the nuance of the scenario. The wider difficulty here is that the real outcome depends on an attribute of one of the players, rather than their moves. In this case, we would expect the right choices to depend on the alcohol consumption of the

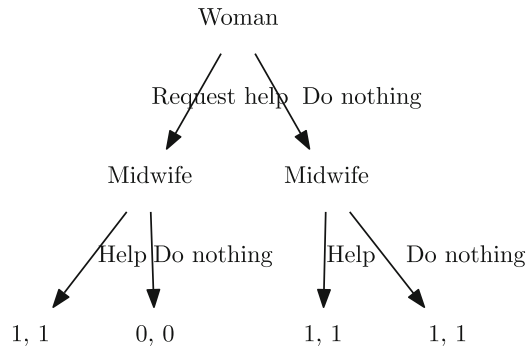


Fig. 11.6 A very simple two player game. The only time things in this very restricted world obviously end poorly is if the woman asks for help but does not get any. This implies that a rational player would always refer if asked for help, and is indifferent otherwise – in other words, there are three possible Nash equilibriums (a Nash equilibrium is a solution to a game between two or more players, where no player can gain from changing their move)

woman, rather than entirely on what she has claimed about it. To reflect this, we would need different variations on the same game to reflect this attribute.

To resolve this, we can do exactly that and cast it as a signalling game (Fig. 11.7), with three types of player corresponding to categories of drinking behaviour (light, moderate, and heavy). Each of these types of player, will play a different game. This also introduces a third player, who we will call nature. Nature takes the first move, and decides the type of the woman according to some probability distribution; in this case we will allow the probability of types to be uniform. This changes the dynamics of play substantially, since the midwife can no longer be certain of which game they are playing, and hence which move yields the best outcome. We must also amend the moves, and payoffs slightly. The woman now claims to be one of the types, and may send a signal to say, for example, that she a heavy drinker. We will also modify the common payoffs to allow light drinkers to get the best outcome no matter what, and moderate and heavy types to get the best outcome only if referred. We can also differentiate between the consequences of not getting help for these types by letting heavy drinkers have a very negative outcome, and moderate drinkers a slight one.

At this point, the game becomes challenging to analyse from a Nash equilibrium perspective (there are several hundred). But, having raised the issue of stigma, we would also like to incorporate this in the game. A possible approach to this is similar to the drinking behaviour of the women, and lets midwives have a type as well, corresponding to how judgemental they are when receiving signals: non-judgemental, moderately judgemental, and harshly judgemental. The expression of this judgement is not a matter of choice on their part, and is assumed to have no impact on their clinical response. Nature now has an additional move, to choose the type of the midwife, and we add costs for sending moderate and heavy signals. A heavy signal to a harshly judgemental midwife adds a heavy cost, and a moderate cost from a moderate midwife. The resulting game might reasonably be said to be intractable.

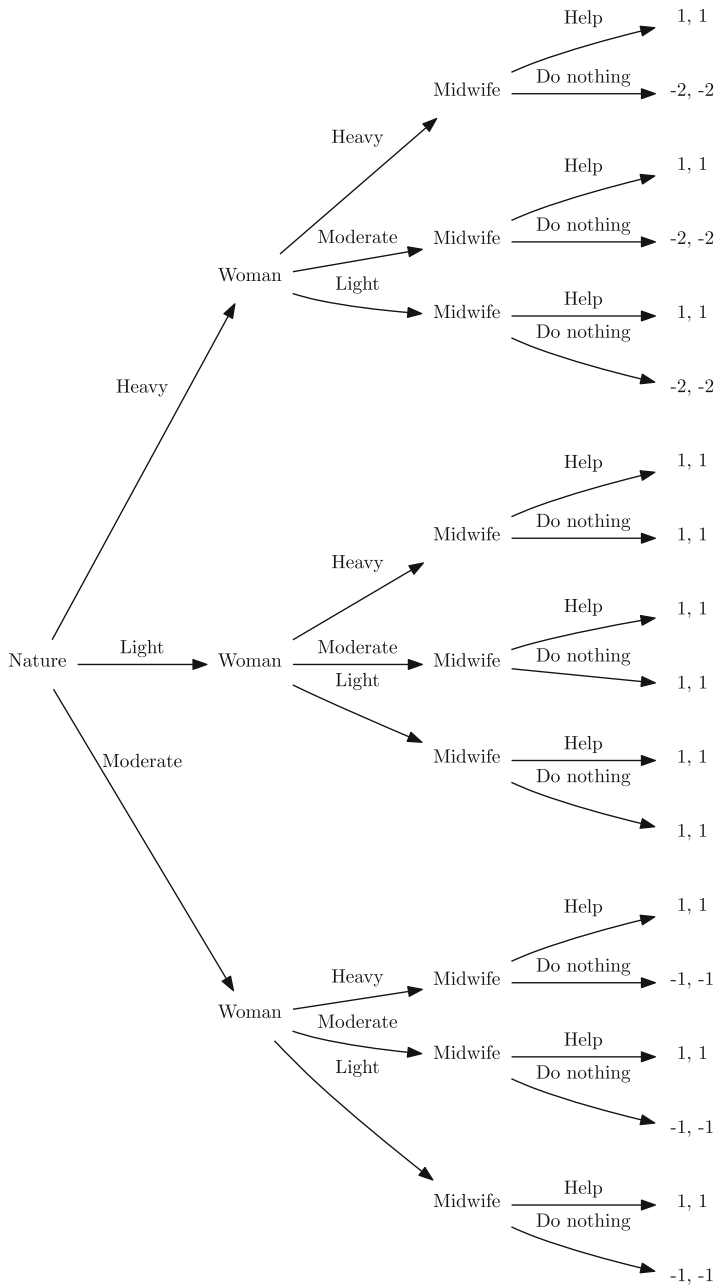


Fig. 11.7 A less simple two player signalling game

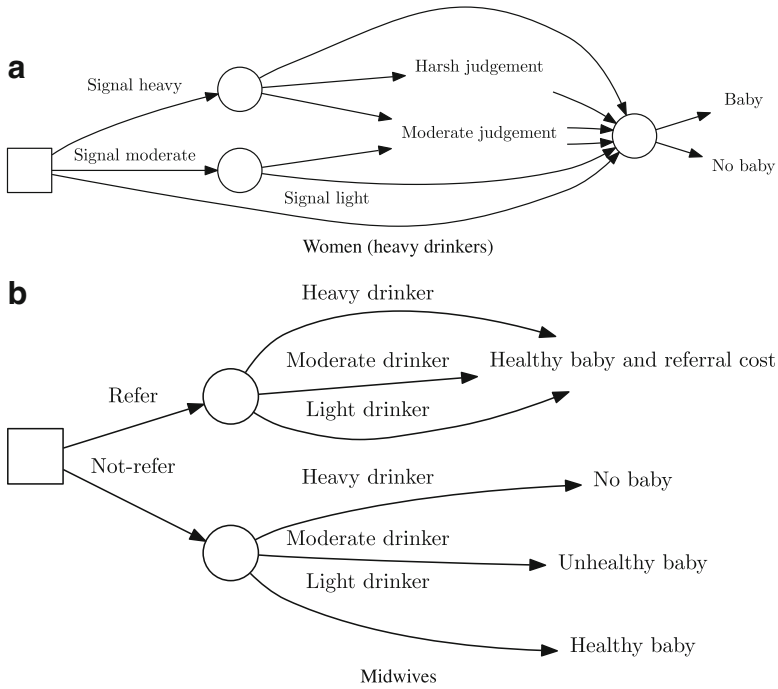


Fig. 11.8 Influence diagrams, showing the game broken into two decision problems. Squares indicate a decision node, while circles are (from the perspective of the agent) chance nodes (a) Women (heavy drinkers). (b) Midwives

At this juncture, we do not gain much further from the game representation, and instead separate it into multiple decision problems.

Breaking the game down into separate decision problems can be achieved by treating the moves of the other players as a chance node, and omitting moves by nature that are known to them. For women, there are two such nodes, corresponding to the move by nature determining the type of midwife they play against, and the midwife’s action. Midwives have a simpler problem with only a single chance node, because the woman’s move is known to them. Figure 11.8 shows the structure of the resulting decision problems. Note that there are in fact three distinct decision problems for the three types of woman, since the move by nature determining their type is known to them.

The precise structure of the decision problem is to some extent dependent on the decision rule in use, for example the Lexicographic heuristic rule is concerned only with a direct relationship between action and consequence. However, the literal translation from game to decision problem for women yields two chance nodes. As a result, solving this using the heuristic approach requires that the nodes be combined. By the same token, an arbitrarily complex problem could be resolved by rules without this limitation. This is significant, in that the decision problem is an

individual agent's model of the situation, which might not be expected to correspond perfectly with the true sequence of events.

From this position, simulating play and augmenting the basic conjecture is easily achievable, since together the game and the decision rules specify the basis for a simulation model. In the disclosure game case, we make additional stipulations on how many games agents play, order of play, the circumstances under which agents observe true types, and the structure of agent populations amongst others.

A.2 Simulation Schedule

In this section we give the step by step process for a single run of the disclosure game simulation.

1. Generate 1000 women, and place them in a queue
2. Generate 100 midwives.
3. For each round of the game
 - a. Take 100 women from the queue
 - b. Pair each one with a random midwife
 - c. For each pair
 - i. The woman sends a signal
 - ii. The midwife refers or not based on the signal
 - iii. The woman is informed of her payoff, the midwife's type, and whether she is referred
 - iv. The woman updates her beliefs
 - v. The midwife stores the game in their memory
 - vi. If the woman is referred
 - A. The midwife is informed of the woman's true type
 - B. The midwife retrospectively updates their beliefs using the true type, and memories of any games with this woman
 - C. The midwife is now eligible to share their memories of the games played with this woman
 - d. Women who have not been referred or had their baby, join the back of the queue
 - e. New women are generated to replace those referred or delivered
 - f. The new women are added to the back of the queue
 - g. For each referred or birthed woman
 - i. With probability p , her memory of games is shared with the active women
 - ii. She is removed from simulation
 - h. The active women update their beliefs
 - i. For each midwife with information to share

- i. With probability p , their memory of games with the referred woman is shared
- ii. The memory is no longer eligible to be shared
- j. The midwives update their beliefs

A.3 Agent Examples

This section provides a worked example for the learning and decision process of each agent model, focusing on the behaviour of the signalling agent.

A.3.1 Lexicographic Heuristic

As an example, take a light drinker who has played three rounds with a succession of particularly judgemental midwives, signalling honestly in two and claiming to be a moderate drinker in one. The most common outcome of the honest signal was a payoff of 10, which is clearly preferable to the 9 gained by claiming to be moderate. On that basis, they choose to signal honestly.

A.3.2 Bayesian Payoff

We take our light drinker from the lexicographic case and assume that they began with an uninformative prior. The 6 possible signal-payoffs pairings are then $[(l, 10), (m, 10), (h, 10), (m, 9), (h, 9), (h, 8)]$, with $\alpha_i = 1$ for all i . After playing the three rounds, $n_{l,10} = 2$, and $n_{m,9} = 1$.

The agent then evaluates R_w for each signal, e.g. for the light signal:

$$\begin{aligned}
 X &= \{10\} \\
 R_w(l) &= \sum_{x \in X} -xp(x|l) = -10p(10|l) \\
 R_w(l) &= -10\left(\frac{\alpha_{l,10} + n_{l,10}}{\sum_j (\alpha_j + n_j)}\right) = -10\left(\frac{1 + 2}{1 + 2}\right) \\
 R_w(l) &= -10\left(\frac{3}{3}\right) = -10
 \end{aligned}$$

and by the same method, $R_w(m) = -9\frac{1}{3}$, and $R_w(h) = -9$, concluding that signalling honestly is the best move.

A.3.3 Bayesian Risk Minimisation

Returning to our example agent, under this model the type of the midwife becomes salient, hence $n_h = 3$, and $n_{l,n} = 2$, $n_{m,n} = 1$. Their prior beliefs remain uninformative, i.e. $\alpha_j = 1, j \in \{l, m, h\}$, $\alpha_{i,j} = 1, i \in \{r, n\}, j \in \{l, m, h\}$. As before, the agent evaluates R_w for the three signals, and the process for the light signal is given below:

$$R_w(l, l) = \sum_{i \in A_m} \sum_{j \in \Theta} -u_w(l, i, l, j)p(j)p(i|l)$$

$$\begin{aligned} R_w(l, l) = & -u_w(l, r, l, l)p(l)p(r|l) - u_w(l, n, l, l)p(l)p(n|l) \\ & - u_w(l, r, l, m)p(m)p(r|l) - u_w(l, n, l, m)p(m)p(n|l) \\ & - u_w(l, r, l, h)p(h)p(r|l) - u_w(l, n, l, h)p(h)p(n|l) \end{aligned}$$

$$u_w(l, i, l, j) = 10$$

$$\begin{aligned} R_w(l, l) = & -10p(l)p(r|l) - 10p(l)p(n|l) - 10p(m)p(r|l) - 10p(m)p(n|l) \\ & - 10p(h)p(r|l) - 10p(h)p(n|l) \end{aligned}$$

$$p(l) = \frac{1 + 0}{1 + 1 + 1 + 3} = \frac{1}{6}$$

$$p(m) = \frac{1 + 0}{1 + 1 + 1 + 3} = \frac{1}{6}$$

$$p(h) = \frac{1 + 3}{1 + 1 + 1 + 3} = \frac{2}{3}$$

$$p(r|l) = \frac{1 + 0}{1 + 1 + 2} = \frac{1}{4}$$

$$p(n|l) = \frac{1 + 2}{1 + 1 + 2} = \frac{3}{4}$$

$$\begin{aligned} R_w(l, l) = & -10 \cdot \frac{1}{6} \cdot \frac{1}{4} - 10 \cdot \frac{1}{6} \cdot \frac{3}{4} - 10 \cdot \frac{1}{6} \cdot \frac{1}{4} - 10 \cdot \frac{1}{6} \cdot \frac{3}{4} - 10 \cdot \frac{2}{3} \cdot \frac{1}{4} \\ & - 10 \cdot \frac{2}{3} \cdot \frac{3}{4} \\ = & -10 \end{aligned}$$

and similarly for moderate ($R_w(m, l) = -9\frac{1}{3}$), and heavy ($R_w(h, l) = -8\frac{1}{2}$) signals, once again concluding that honesty is the better option.

Table 11.4 CPT parameters

Name	Description	Value
γ	Probability weighting for gains	0.61
δ	Probability weighting for losses	0.69
α	Power for gains	0.88
β	Power for losses	0.88
λ	Loss aversion	2.25

A.3.4 Descriptive Decision Theory

Once again, we return to the light drinker example. The inferential aspects are identical with the more complex Bayesian risk minimisation algorithm, hence $p(j)p(i|l)$, and $u_w(l, i, l, j)$ remain the same, but the agent additionally calculates $v(u_w(l, i, l, j))w^+(p(j))w^+(p(i|l))$. For the CPT parameters, the values are those originally given by Tversky and Kahneman (1992) and used in the actual simulations which are given in Table 11.4.

$$\alpha = 0.88$$

$$\gamma = 0.61$$

$$p(l) = \frac{1}{6}$$

$$p(m) = \frac{1}{6}$$

$$p(h) = \frac{2}{3}$$

$$p(r|l) = \frac{1}{4}$$

$$p(n|l) = \frac{3}{4}$$

$$u_w(l, i, l, j) = 10$$

$$f = (10; \frac{1}{24}, 10; \frac{1}{8}, 10; \frac{1}{24}, 10; \frac{1}{8}, 10; \frac{1}{6}, 10; \frac{1}{2})$$

$$f^+ = f, f^- = ()$$

$$n = 5$$

$$v(u_w) = f(u_w) = u_w^\alpha$$

$$v(u_w) = 10^{0.88} = 7.59$$

$$\begin{aligned}\pi_0^+ &= w^+\left(\frac{1}{24} + \frac{1}{8} + \frac{1}{24} + \frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) - w^+\left(\frac{1}{8} + \frac{1}{24} + \frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) \\ &= w^+(1) - w^+\left(\frac{23}{24}\right) \\ &= 0.19\end{aligned}$$

$$\begin{aligned}\pi_1^+ &= w^+\left(\frac{1}{8} + \frac{1}{24} + \frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) - w^+\left(\frac{1}{24} + \frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) \\ &= w^+\left(\frac{23}{24}\right) - w^+\left(\frac{5}{6}\right) \\ &= 0.17\end{aligned}$$

$$\begin{aligned}\pi_2^+ &= w^+\left(\frac{1}{24} + \frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) - w^+\left(\frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) = w^+\left(\frac{5}{6}\right) - w^+\left(\frac{19}{24}\right) \\ &= 0.04\end{aligned}$$

$$\begin{aligned}\pi_3^+ &= w^+\left(\frac{1}{8} + \frac{1}{6} + \frac{1}{2}\right) - w^+\left(\frac{1}{6} + \frac{1}{2}\right) = w^+\left(\frac{19}{24}\right) - w^+\left(\frac{2}{3}\right) \\ &= 0.09\end{aligned}$$

$$\begin{aligned}\pi_4^+ &= w^+\left(\frac{1}{6} + \frac{1}{2}\right) - w^+\left(\frac{1}{2}\right) = w^+\left(\frac{2}{3}\right) - w^+\left(\frac{1}{2}\right) \\ &= 0.09\end{aligned}$$

$$\begin{aligned}\pi_5^+ &= w^+\left(\frac{1}{2}\right) \\ &= 0.42\end{aligned}$$

$$V(f) = V(f^+) + V(f^-) = V(f^+) + 0$$

$$V(f^+) = \sum_i^n \pi_i^+(f^+)v_i^+(f^+) = 7.59$$

And as before, following the same process for moderate and heavy signals, which yields respectively 7.14, and 6.22, the agent chooses the higher valued action and sends an honest signal.

A.4 Supplementary Figures




Fig. 11.9 Average fraction of population referred by each appointment, after 1000 rounds, mean with 95 % confidence limit over 1000 runs. Note that the large number of runs leads to very tight confidence intervals

A.5 Sensitivity Analysis

This section provides complete variance based sensitivity analysis results for the disclosure game model. Each subsection gives results for one simulation output under all four decision rules, with tables providing the percentage of overall variance attributable to the individual parameters, emulator quality statistics, and the five most significant interaction contributions to variance in the output.

A.5.1 Median Moderate Drinker Signalling

Table 11.5 Median moderate drinker signalling parameter sensitivity

Parameter	Description	Lexicographic	Bayesian payoff	Bayesian	CPT
$p_w(m)$	Proportion of moderate drinkers	0.367	1.145	0.801	0.614
$p_w(l)$	Proportion of light drinkers	10.080	37.750	23.968	5.137
$p_m(m)$	Proportion of moderate midwives	6.715	13.017	0.894	1.485
$p_m(l)$	Proportion of non-judgemental midwives	43.942	1.655	1.602	2.618
q_w	Probability of women sharing	0.198	5.527	4.460	1.159
w_w	Weight of shared information for women	0.355	13.025	2.716	0.888
q_m	Probability of midwives sharing	0.145	0.667	0.368	0.157
w_m	Weight of shared information for midwives	0.118	0.376	0.176	0.200
x_h	Health payoff for healthy delivery	0.457	9.618	1.912	15.355
$s_i[a_i] : s_i[a_{-i}]$	Pseudo-count favouring honesty	0.140	7.537	10.427	7.795
Total	All parameters and two way interactions	86.777	96.527	85.529	74.123

Table 11.6 Median moderate drinker signalling emulator statistics

Rule	σ^2	Nugget σ^2	Mean output	Total output variance	Code uncertainty	RMSSE
Lexicographic	0.834	0.131	0.817	0.012	0.252	1.746
Bayesian payoff	1.667	0.475	0.662	0.003	0.181	3.12
Bayesian	3.352	0.534	1.160	0.001	0.068	2.423
CPT	1.503	0.331	1.241	0.002	0.101	1.842

Table 11.7 Top 5 interaction terms for CPT decision rule

Parameter	Variance
$x_h * s_i[a_i] : s_i[a_{-i}]$	20.814
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	5.698
$p_w(l) * x_h$	2.895
$s_i[a_i] : s_i[a_{-i}] * w_w$	2.799
$s_i[a_i] : s_i[a_{-i}] * q_m$	1.686

Table 11.8 Top 5 interaction terms for Bayesian decision rule

Parameter	Variance
$p_w(l)*s_i[a_i] : s_i[a_{-i}]$	17.270
$p_w(l)*q_m$	6.054
q_m*w_w	3.814
$s_i[a_i] : s_i[a_{-i}]*q_m$	3.538
$s_i[a_i] : s_i[a_{-i}]*w_w$	3.084

Table 11.9 Top 5 interaction terms for lexicographic decision rule

Parameter	Variance
$p_m(l)*p_m(m)$	15.331
$p_m(m)*p_w(l)$	3.682
$p_m(l)*p_w(l)$	3.581
$p_m(m)*q_m$	0.349
$p_m(l)*q_m$	0.279

Table 11.10 Top 5 interaction terms for Bayesian payoff decision rule

Parameter	Variance
$p_w(l)*w_w$	4.045
$x_h*s_i[a_i] : s_i[a_{-i}]$	1.856
$p_w(l)*q_m$	1.231
$s_i[a_i] : s_i[a_{-i}]*q_m$	0.997
q_m*w_w	0.929

A.5.2 Median Between Groups IQR

Table 11.11 Median between groups IQR parameter sensitivity

Parameter	Description	Lexicographic	Bayesian payoff	Bayesian	CPT
$p_w(m)$	Proportion of moderate drinkers	0.327	0.688	0.457	0.586
$p_w(l)$	Proportion of light drinkers	11.223	20.123	11.046	4.081
$p_m(m)$	Proportion of moderate midwives	36.630	1.160	0.364	1.945
$p_m(l)$	Proportion of non-judgemental midwives	6.228	4.487	0.0964	2.627
q_w	Probability of women sharing	0.498	0.235	2.537	1.812
w_w	Weight of shared information for women	1.018	2.307	1.889	0.740
q_m	Probability of midwives sharing	0.158	0.343	0.387	0.156
w_m	Weight of shared information for midwives	0.076	0.973	0.125	0.213
x_h	Health payoff for healthy delivery	0.317	10.960	3.305	16.493
$s_i[a_i] : s_i[a_{-i}]$	Pseudo-count favouring honesty	1.107	8.411	2.890	6.729
Total	All parameters and two way interactions	81.702	83.693	47.449	71.032

Table 11.12 Median between groups IQR emulator statistics

Rule	σ^2	Nugget σ^2	Mean output	Total output variance	Code uncertainty	RMSSE
Lexicographic	0.930	0.240	0.249	0.002	0.040	1.832
Bayesian payoff	1.242	0.417	0.232	0.001	0.0034	2.308
Bayesian	1.254	0.131	0.644	0.000	0.019	1.167
CPT	1.190	0.313	0.659	0.000	0.024	1.701

Table 11.13 Top 5 interaction terms for CPT decision rule

Parameter	Variance
$x_h * s_i[a_i] : s_i[a_{-i}]$	19.551
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	3.838
$s_i[a_i] : s_i[a_{-i}] * w_w$	2.450
$p_w(l) * x_h$	2.337
$s_i[a_i] : s_i[a_{-i}] * q_m$	2.046

Table 11.14 Top 5 interaction terms for Bayesian decision rule

Parameter	Variance
$s_i[a_i] : s_i[a_{-i}] * q_m$	4.284
$p_w(l) * q_m$	3.866
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	2.943
$x_h * s_i[a_i] : s_i[a_{-i}]$	2.680
$q_m * w_w$	2.282

Table 11.15 Top 5 interaction terms for lexicographic decision rule

Parameter	Variance
$p_m(l) * p_m(m)$	12.046
$p_m(m) * p_w(l)$	5.054
$p_m(l) * p_w(l)$	3.005
$p_m(l) * w_w$	0.819
$p_m(m) * w_w$	0.757

Table 11.16 Top 5 interaction terms for Bayesian payoff decision rule

Parameter	Variance
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	12.883
$p_w(l) * w_w$	5.667
$p_w(l) * x_h$	2.447
$x_h * s_i[a_i] : s_i[a_{-i}]$	2.360
$p_m(m) * p_w(l)$	1.919

A.5.3 Median Moderate Drinker Signalling IQR

Table 11.17 IQR of median moderate drinker signalling parameter sensitivity

Parameter	Description	Lexicographic	Bayesian payoff	Bayesian	CPT
$p_w(m)$	Proportion of moderate drinkers	0.428	9.828	3.816	2.068
$p_w(l)$	Proportion of light drinkers	8.369	13.791	4.045	2.400
$p_m(m)$	Proportion of moderate midwives	13.416	0.712	0.676	0.583
$p_m(l)$	Proportion of non-judgemental midwives	21.079	0.648	0.659	0.373
q_w	Probability of women sharing	2.307	3.481	0.891	0.600
w_w	Weight of shared information for women	6.021	6.009	0.562	0.937
q_m	Probability of midwives sharing	0.315	1.829	0.114	0.117
w_m	Weight of shared information for midwives	1.652	1.354	0.260	0.0.139
x_h	Health payoff for healthy delivery	0.253	0.612	4.889	15.146
$s_i[a_i] : s_i[a_{-i}]$	Pseudo-count favouring honesty	0.504	3.096	19.863	25.999
Total	All parameters and two way interactions	84.9968	77.413	57.125	83.322

Table 11.18 IQR of median between groups IQR emulator statistics

Rule	σ^2	Nugget σ^2	Mean output	Total output variance	Code uncertainty	RMSSE
Lexicographic	1.425	0.436	0.549	0.008	0.114	2.719
Bayesian payoff	1.223	0.496	0.747	0.012	0.207	2.034
Bayesian	1.065	0.000	0.230	0.002	0.088	1.015
CPT	0.874	0.213	0.233	0.001	0.066	1.806

Table 11.19 Top 5 interaction terms for CPT decision rule

Parameter	Variance
$x_h * s_i[a_i] : s_i[a_{-i}]$	17.377
$p_w(m) * s_i[a_i] : s_i[a_{-i}]$	3.356
$s_i[a_i] : s_i[a_{-i}] * w_w$	3.036
$s_i[a_i] : s_i[a_{-i}] * q_m$	2.067
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	1.721

Table 11.20 Top 5 interaction terms for Bayesian decision rule

Parameter	Variance
$p_w(m) * s_i[a_i] : s_i[a_{-i}]$	4.188
$x_h * s_i[a_i] : s_i[a_{-i}]$	3.120
$s_i[a_i] : s_i[a_{-i}] * q_m$	2.423
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	2.279
$p_w(l) * q_m$	1.489

Table 11.21 Top 5 interaction terms for lexicographic decision rule

Parameter	Variance
$p_m(m)*p_w(l)$	12.068
$p_m(l)*p_w(l)$	6.794
$p_m(l)*p_m(m)$	5.567
$p_m(m)*q_m$	0.886
$p_m(l)*q_m$	0.692

Table 11.22 Top 5 interaction terms for Bayesian payoff decision rule

Parameter	Variance
$p_w(l)*p_w(m)$	8.357
$p_w(l)*w_w$	7.431
$p_w(l)*s_i[a_i] : s_i[a_{-i}]$	4.882
$p_w(l)*q_m$	2.346
$p_w(l)*w_m$	2.025

A.5.4 IQR of Between Groups IQR

Table 11.23 IQR of median between groups IQR parameter sensitivity

Parameter	Description	Lexicographic	Bayesian payoff	Bayesian	CPT
$p_w(m)$	Proportion of moderate drinkers	0.691	5.926	1.053	1.265
$p_w(l)$	Proportion of light drinkers	3.664	17.047	4.877	3.656
$p_m(m)$	Proportion of moderate midwives	41.369	1.124	0.814	0.591
$p_m(l)$	Proportion of non-judgemental midwives	7.109	0.739	0.496	0.378
q_w	Probability of women sharing	1.963	2.038	0.733	0.589
w_w	Weight of shared information for women	7.932	11.193	2.289	1.960
q_m	Probability of midwives sharing	0.413	1.972	0.267	0.069
w_m	Weight of shared information for midwives	0.120	2.902	0.150	0.123
x_h	Health payoff for healthy delivery	0.228	3.190	6.308	14.777
$s_i[a_i] : s_i[a_{-i}]$	Pseudo-count favouring honesty	0.673	10.411	22.901	26.340
Total	All parameters and two way interactions	85.740	88.611	68.640	84.210

Table 11.24 IQR of median between groups IQR emulator statistics

Rule	σ^2	Nugget σ^2	Mean output	Total output variance	Code uncertainty	RMSSE
Lexicographic	0.826	0.409	0.259	0.002	0.034	2.364
Bayesian payoff	3.202	0.520	0.328	0.002	0.032	2.452
Bayesian	1.177	0.041	0.133	0.000	0.018	1.152
CPT	0.874	0.118	0.126	0.000	0.017	1.570

Table 11.25 Top 5 interaction terms for CPT decision rule

Parameter	Variance
$x_h * s_i[a_i] : s_i[a_{-i}]$	18.626
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	3.312
$s_i[a_i] : s_i[a_{-i}] * w_w$	2.823
$s_i[a_i] : s_i[a_{-i}] * q_m$	2.694
$x_h * q_m$	1.022

Table 11.26 Top 5 interaction terms for Bayesian decision rule

Parameter	Variance
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	7.947
$x_h * s_i[a_i] : s_i[a_{-i}]$	4.048
$s_i[a_i] : s_i[a_{-i}] * q_m$	3.134
$p_w(l) * q_m$	2.307
$p_m(m) * s_i[a_i] : s_i[a_{-i}]$	2.232

Table 11.27 Top 5 interaction terms for lexicographic decision rule

Parameter	Variance
$p_m(l) * p_m(m)$	8.659
$p_m(m) * p_w(l)$	3.726
$p_m(l) * p_w(l)$	3.237
$p_m(l) * w_w$	1.564
$p_m(m) * w_w$	1.558

Table 11.28 Top 5 interaction terms for Bayesian payoff decision rule

Parameter	Variance
$q_m * w_w$	3.830
$p_w(l) * p_w(m)$	3.808
$p_w(l) * s_i[a_i] : s_i[a_{-i}]$	3.401
$p_w(l) * q_m$	2.385
$x_h * s_i[a_i] : s_i[a_{-i}]$	2.294

References

Ainslie, G. (1991). Derivation of “rational” economic behavior from hyperbolic discount curves. *The American Economic Review*, 81(2), 334–340.

Alvik, A., Haldorsen, T., Groholt, B., & Lindemann, R. (2006). Alcohol consumption before and during pregnancy comparing concurrent and retrospective reports. *Alcoholism: Clinical & Experimental Research*, 30(3), 510–515.

Axelrod, R. (1997). Advancing the art of simulation in the social sciences. *Complexity*, 12(3), 1–13.

Bijak, J., Hilton, J., Silverman, E., & Cao V D (2013) Reforging the wedding ring. *Demographic Research*, 29, 729–766.

Booij, A. S., Praag, B. M. S., & Kuilen, G. (2009). A parametric analysis of prospect theory’s functionals for the general population. *Theory and Decision*, 68(1–2), 115–148.

Byrnes, J. P., Miller, D. C., & Schafer, W. D. (1999). Gender differences in risk taking: A meta-analysis. *Psychological Bulletin*, 125(3), 367–383.

Carnell, R. (2012). LHS: Latin Hypercube Samples. R package version 0.10. <https://CRAN.R-project.org/package=lhs>. Accessed 1 July 2014.

- Department of health. (2008). *Pregnancy and alcohol*. London: Department of Health.
- Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), 12.
- Epstein, J. M. (2014). *Agent_Zero: Toward neurocognitive foundations for generative social science*. Princeton: Princeton University Press.
- Epstein, J. M., & Axtell, R. L. (1994). *Growing artificial societies: Social science from the bottom up. Complex adaptive systems*. Washington, DC: Brookings Institution Press.
- Feltovich, N., Harbaugh, R., & To, T. (2002). Too cool for school? Signalling and countersignalling. *RAND Journal of Economics*, 33(4), 630–649.
- Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science*, 44(7), 879–895.
- Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In: D. J. Koehler, N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 62–88). Oxford: Blackwell Publishing.
- Gilbert, N. (1999). Simulation: A new way of doing social science. *American Behavioral Scientist*, 42(10), 1485–1487.
- Glöckner, A., & Pachur, T. (2012). Cognitive models of risky choice: Parameter stability and predictive accuracy of prospect theory. *Cognition*, 123(1), 21–32.
- Gomberg, E. S. (1988). Alcoholic women in treatment: The question of stigma and age. *Alcohol and Alcoholism*, 23(6), 507–14.
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: the role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21(5), 493–518.
- Humphriss, R., Hall, A., May, M., Zuccolo, L., & Macleod, J. (2013). Prenatal alcohol exposure and childhood balance ability: Findings from a UK birth cohort study. *BMJ Open*, 3(6), 1–8.
- Insua, D. R., Rios, J., & Banks, D. (2009). Adversarial risk analysis. *Journal of the American Statistical Association*, 104(486), 841–854.
- Kennedy, M. C. (2004). CTCD: GEM Software. <http://ctcd.group.shef.ac.uk/gem.html>. Accessed: 10 July 2014.
- Kreps, D. M., & Cho, I. (1987). Signalling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2), 179–221.
- Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2), 573–597.
- Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28(1), 143–166
- National Institute for Health and Care Excellence (2010a). Antenatal care. CG62. London: NICE.
- National Institute for Health and Care Excellence (2010b). Pregnancy and complex social factors. CG110. Manchester: NICE.
- Neilson, W., & Stowe, J. (2002). A further examination of cumulative prospect theory parameterizations. *Journal of Risk and Uncertainty*, 24(1), 31–47.
- Nilsson, H., Rieskamp, J., & Wagenmakers, E. J. (2011). Hierarchical bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, 55(1), 84–93.
- Oakley, A., Strange, V., Bonell, C., Allen, E., & Stephenson, J. (2006). Process evaluation in randomised controlled trials of complex interventions. *British Medical Journal*, 332(7538), 413–416.
- Oakley, J., & O’Hagan, A. (2002). Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*, 89(4), 769–784.
- Oakley, J. E., & O’Hagan, A. (2004). Probabilistic sensitivity analysis of complex models: A bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3), 751–769.
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090), 223–226.
- Padoa-Schioppa, C., & Assad, J. A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nature Neuroscience*, 11(1), 95–102.

- Phillips, D., Thomas, K., Cox, H., Ricciardelli, L. A., Ogle, J., Love, V., Steele, A., et al. (2007). Factors that influence women's disclosures of substance use during pregnancy: A qualitative study of ten midwives and ten pregnant women. *The Journal of Drug Issues*, 37(2), 357–376.
- R Core Team (2014). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing.
- Radcliffe, P. (2011). Substance-misusing women: Stigma in the maternity setting. *British Journal of Midwifery*, 19(8), 497–506.
- Redshaw, M., & Henderson, J. (2014). *Safely delivered: A national survey of women's experiences of maternity care 2014*. Oxford: The National Perinatal Epidemiology Unit.
- Resnick, M. (1994). *Turtles, termites and traffic jams: Explorations in massively parallel microworlds*. Cambridge: MIT Press.
- Rustichini, A. (2009). Neuroeconomics: what have we found, and what should we search for. *Current Opinion in Neurobiology* 19(6), 672–677.
- Samuelson, P. A. (1937). Note on measurement of utility. *The Review of Economic Studies*, 4(2), 155–161.
- Silverman, E., Bijak, J., & Noble, J. (2011). Feeding the beast: Can computational demographic models free us from the tyranny of data? In: T. Lenaerts, M. Giacobini, H. Bersini, P. Bourguine, M. Dorigo, & R. Doursat (Eds.), *Advances in artificial life, ECAL 2011* (Vol. 276, pp. 747–754). Cambridge: MIT Press.
- Silverman, E., Bijak, J., Hilton, J., Cao, V. D., & Noble, J. (2013). When demography met social simulation: A tale of two modelling approaches. *Journal of Artificial Societies and Social Simulation*, 16(4), 9.
- Sokol, R. J., Martier, S. S., & Ager, J. W. (1989). The t-ace questions: Practical prenatal detection of risk-drinking. *American Journal of Obstetrics and Gynecology*, 160(4), 863–870.
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3), 201–207.
- Thiele, J. C., Kurth, W., & Grimm, V. (2014). Facilitating parameter estimation and sensitivity analysis of agent-based models: A cookbook using NetLogo and R. *Journal of Artificial Societies and Simulation*, 17(3), 11.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.
- Waldherr, A., & Wijermans, N. (2013). Communicating social simulation models to sceptical minds. *Journal of Artificial Societies and Social Simulation*, 16(4), 13.

Part IV
Family Formation and Fertility

Chapter 12

An Agent-Based Model of Sex Ratio at Birth Distortions

Ridhi Kashyap and Francisco Villavicencio

12.1 Introduction

Since the 1980s and 1990s, a decline in total fertility levels across a number of countries in Asia and the Caucasus has been accompanied by an unprecedented rise in the proportion of male births compared to female births, usually expressed in terms of the sex ratio at birth (SRB) (Bongaarts 2013; Guilmoto 2009). According to available demographic estimates since the 1950s SRBs in most countries lie between 104–106 male births for every 100 female births and over the past two centuries SRB levels have remained unchanged across different settings where fertility has declined.¹ Although this unprecedented trend is not universal across the diverse demographic contexts in Asia, the sizeable populations where it has been

¹The United Nations (UN) Population Division publishes a comparative global SRB time series from 1950 onward and produces SRB forecasts until 2100 (United Nations 2013). SRB forecasts, unlike more recent probabilistic approaches adopted in the UN's fertility forecasts, are deterministic. These estimates as well as other demographic studies documenting SRB trends make use of different types of data sources with varying levels of reliability: birth registration data, which are the most reliable when available; birth-history estimates from large surveys; and census data on

R. Kashyap (✉)

Department of Sociology and Nuffield College, University of Oxford, Oxford, UK

Max Planck Institute for Demographic Research, Rostock, Germany

e-mail: ridhi.kashyap@nuffield.ox.ac.uk

F. Villavicencio

Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark

Max-Planck Odense Center on the Biodemography of Aging, Odense, Denmark

e-mail: villavicencio@imada.sdu.dk

noted—China, India, South Korea, Vietnam, Georgia, Azerbaijan, Albania, among others—make it one of the “most notable anomalies” of contemporary demography (Guilmoto 2009, p. 519). While earlier literature presumed that the presence of widespread son preferences would delay the fertility transition and keep fertility rates higher than they would be in their absence (Amin and Mariam 1987; Nath and Land 1994), the rise of SRBs despite steep fertility declines, as first clearly evidenced in the case of South Korea, suggested that sex-selective abortion (SSA) provided a strategy for reconciling both.²

This has led demographers to speculate that SRBs may rise in other contexts where son preference persists but safe, effective and cheap sex determination technology has yet to become available (Bongaarts 2013). Figure 12.1 shows United Nations estimates for total fertility rates (TFR) and SRB trends from 1970 to 2010 for South Korea, China and India. Also reported are data of an example (Turkey) where a similar rise in SRB has not accompanied the fertility decline.

The causes, patterns and demographic implications of distorted SRBs remain the subject of much research. Extensive demographic research has paid attention to highlighting SRB levels and trends in country-specific as well as comparative perspective across Asia (Attané and Guilmoto 2007; Duthé et al. 2012; Guilmoto 2009, 2012). In contrast, comparatively less attention has been paid to understanding the levels and trends of the micro-level factors underpinning SRB trajectories. The dynamics of son preference, its interactions with total family size aspirations and the availability of sex-selective abortion are important to model especially as Bongaarts notes “policymakers are hampered by an absence of methods for projecting trends in sex ratios at birth” (Bongaarts 2013, p. 185).

Recognizing this lacuna in the literature, the focus of emerging scholarship on SRBs has been more theoretical. Guilmoto (2009) theorizes the micro-level fertility calculus that couples engage in that leads to SRB distortions at the macro-level. In the same paper, he likens the shape of SRB trajectories to an “archetypal transition cycle” involving three stages: first a rise, followed by a levelling-off and an eventual decline towards normal levels that he terms the “sex ratio transition” (Guilmoto 2009, p. 519). Bongaarts (2013) develops a similar macro-level framework that relates different stages in SRB trajectories—a gradual, then steep rise, followed by levelling off and eventual decline—to different stages of the fertility transition.

recent births or age structure (Guilmoto 2009). Data quality problems for estimating SRB trends in Asia are detailed extensively elsewhere (Attané and Guilmoto 2007).

²Sex-selective abortion requires pre-natal sex-determination technologies as well as methods for abortion. Different technologies for pre-natal sex-determination technologies presently exist—amniocentesis, an invasive procedure that is conducted between 15 and 20 weeks of pregnancy, ultrasonography, which can determine the sex of the foetus as early as 11 weeks, and blood tests involving the analysis of the foetal DNA floating in the mother’s bloodstream, which are minimally invasive procedures and can be done at home as early as 7 weeks into the pregnancy (Bongaarts 2013).




Fig. 12.1 Total fertility rates (TFR) and sex ratio at birth (SRB) trends 1970–2010 (Source: United Nations (2013))

In Bongaarts’s schema, SRB levels hover at normal or near-normal levels at the early-mid stages of the fertility transition when fertility levels are high, and become distorted, at first gradually then steadily in the mid-stage of the fertility transition as fertility begins to fall. He hypothesizes an eventual turnaround emerging when low fertility levels are established at the late stage. How micro-level factors plausibly change and interact to generate macro-level SRB trajectories at each of the stages—the rise from normal levels, the levelling off phase and an eventual decline and return to normal levels—in the sex ratio transition is not explicitly explored by the authors.

Guilmoto attributes this to the limitations of existing data that preclude a “more detailed decomposition” of the three micro-level factors he identifies—son

preference, technological access and fertility decline—in explaining observed SRB levels and trends (Guilmoto 2009, p. 535). Existing data and approaches make it difficult to quantify, for example, what levels and rates of change in son preference over time as well as probabilities of sex-selective abortion underlie observed SRB and fertility trajectories. Furthermore, limited research has sought to disentangle the impact of SSA practice on reducing fertility as distinct from a population that exclusively practices differential stopping behaviour (DSB).³

This chapter dynamically models individual-level fertility preferences and reproductive behaviours bottom-up to examine emergent population-level SRB and total fertility patterns through the use of an agent-based model (ABM). Agent-based computational modelling methods have been proposed as a valuable set of techniques to model emergent demographic phenomena through the simulation of heterogeneous individuals who follow behavioural rules and adapt their behaviours in response to stimuli (Billari and Prskawetz 2003; Bonabeau 2002). The ability to simulate dynamic and endogenous processes such as changing rates of son preference or the impact of the fertility decline on increasing readiness to abort, model individual adaptation to exogenous stimuli such as the availability of sex selection technology, explore counterfactual scenarios and overcome data limitations make ABM a promising technique by which to examine micro-level reproductive behaviours that generate SRB distortions.

A detailed description of the model, study of model behaviour and experiments and model calibration for South Korea is presented elsewhere (Kashyap and Villavicencio [forthcoming](#)). This chapter extends the model to Indian SRB trajectories and contrasts the SRB trajectories of both India and South Korea to highlight how different micro-level patterns can be used to describe their distinctive trajectories. South Korea witnessed the emergence of skewed SRBs alongside a rapid fertility decline in the 1980s. By the 1990s SRBs started to level off and very soon afterwards by the mid-1990s SRBs showed a remarkable turnaround towards normalization (see Fig. 12.1). By the mid-2000s SRBs had already returned to near normal levels. As an exemplar of a country that has been through all three stages of the sex ratio transition, calibrating the ABM to the South Korean case can shed light on the levels and rates of change in son preference, rate of diffusion of technology and probabilities of sex-selective abortion that plausibly underpinned different stages of the SRB trajectory. The Indian case in contrast witnessed a later rise in SRBs in the 1990s, followed by a levelling off that began in the mid-to late-2000s. The future course of SRB trajectories in India remains open to question. Will the turnaround to normal or near-normal levels in India be as rapid as that witnessed in South Korea? By calibrating the ABM so that we are able to recreate the Indian

³Differential stopping behaviour (DSB) refers to fertility behaviour in which couples continue childbearing until they reach a desired number of sons by regulating their contraceptive use and childbearing behaviour based on the sex composition of existing children. Several papers report the presence of DSB—manifested in higher levels of contraceptive use after bearing sons, higher levels of parity progression in daughter-only families, or high sex ratios at last birth—across different son-preferring populations (Amin and Mariam 1987; Arnold et al. 1998; Bongaarts 2013; Clark 2000; Retherford and Roy 2003).

SRB trajectory bottom-up, we can also explore the future implications of changing fertility preferences and reproductive behaviours.

Our model calibration for both countries shows how even as son preference levels were likely declining in both contexts SRB distortions might have plausibly emerged due to the combined effects of technology diffusion and increasing probabilities to sex-selectively abort at lower birth parities resulting from declining total fertility levels. Indeed, an interesting insight from the South Korean calibration is that even relatively low levels of son preference, for example $\sim 30\%$ of the population desiring one son, can result in significant SRB distortions if access to technology diffuses steadily and fertility falls rapidly to encourage individuals to abort at the lowest parity—parity 0 or the transition to first birth. Model calibration highlights how the extent of SRB distortion is strongly linked to the parity at which individuals choose to sex-selectively abort. SRB distortion is greater when individuals are ready to abort at low parities such as parity 0 or 1 and less inclined to progress to higher births instead, a situation that becomes increasingly likely as fertility falls.

12.2 Modelling Reproductive Behaviours

12.2.1 Sex Selection

The model operationalizes the “ready, willing and able” (RWA) framework to conceptualize the micro-level causal processes leading to the practice of sex selection. We borrow these insights from Guilmo (2009) who adapts Ansley Coale’s RWA framework (Coale 1973), originally used by Coale to account for the European fertility decline in the nineteenth century, to the practice of sex selection. To explain the fertility decline in Europe, Coale posited that three conditions had to be met: (1) Were individuals *willing* to limit fertility? This condition required that limiting fertility had to be acceptable and not proscribed by existing normative considerations (e.g. religious or ethical) that individuals adhered to. (2) Did parents have access to technology (e.g. contraception) to be able to limit family size? This condition described the *ability* of parents to limit fertility. (3) Were parents *ready* to limit their family size? The readiness condition brought in considerations of utility, that is, limiting fertility had to be economically advantageous to the actor within their decision-making calculus.

In their more general reformulation of Coale’s classic framework, Lesthaeghe and Vanderhoeft (2001) highlight its usefulness as a conceptual apparatus in social demography to explain adaptation to new modes of behaviour involving processes of innovation and diffusion. In the same paper, the authors also apply the framework to the study of fertility transitions in the developing world, underscoring its wider applicability beyond the setting for which it had originally been developed. The strength of the RWA conceptual framework, they argue, is in its ability to integrate economic and noneconomic paradigms that include normative legitimacy, economic utility, as well as functional considerations of technological access in explaining

adaptations to new behaviours. Moreover, by explicitly identifying distinctive causal processes underlying the uptake of new behaviour, the framework highlights how transitions to new forms of behaviour can take different forms depending on differing rates and patterns of change in each of the R, W and A components.

Guilmoto (2009) adapts the RWA conceptual framework to the proximate causal processes underlying the practice of sex selection. By adopting the RWA approach from the “actor’s point of view”, sex selection can be seen as “a rational strategy in response to changing constraints and opportunities within existing gender regimes” (Guilmoto 2009, p. 526). The practice of sex selection from the actor or individual’s perspective can be seen as the outcome of three conditions being met: (1) The actor must be *willing* to consider sex selection because of the persistence of cultural norms that reinforce the value of male offspring. Sex preferences must be entrenched within a social and cultural context that allows agents to consider acting on them. Across the diverse contexts of South and East Asia, where strong SRB distortions have been observed, Das Gupta et al. (2003) argue that son preference is underpinned by commonalities in patrilineal and patrilocal kinship systems that reinforce the material and ideological value of male offspring. (2) Actors must be *able* to access sex-selective abortion due to the availability of relatively affordable and accurate pre-natal sex determination technology to detect the sex of the foetus, as well as have access to abortion that enables them to abort it when unwanted. (3) Even in the presence of son preference (*willingness*) and access to technology (*ability*), actors will not perform sex selection unless they are *ready* to do so. The idea of readiness usually gains importance in low fertility situations with the diffusion of norms towards smaller families wherein deliberate sex selection becomes a preferable goal instead of additional births as the means to realize son preference. Guilmoto describes the costlier trade-off between higher parity births and the realization of son preference with declining fertility in terms of a “fertility squeeze” that determines the readiness to practice sex selection. In what follows we elaborate on how each of these three components of the RWA framework are incorporated in the ABM.

12.2.2 *Son Preference (Willingness)*

Son preference sp is assigned as follows in the model: agents have either no preference for male offspring ($sp = 0$) or a desire for one male offspring ($sp = 1$). Those who have a son preference ($sp = 1$) practice differential stopping behaviour, that is, they have higher fertility rates than those who do not or have met their son preference (see Sect. 12.2.3). The model design allows for an individual’s probability to be son preferring ($sp = 1$) to be determined by the period as well as cohort she belongs to.

We choose to assign son preference dichotomously because we believe this effectively captures the way the social norm influences reproductive behaviour and allows for easy interpretation for the levels of son preference in the population. Son

preference assigned this way can be interpreted as the desire for a surviving son. While some individuals, particularly among older cohorts, may desire more than one son, the desire for several sons likely reflects the indirect influence of higher mortality conditions where bearing at least two sons might be considered a strategy to ensure at least one survived into adulthood. By simulating mortality dynamics for males until age 50 (described in greater detail in the Sect. 12.3.2) the feedback effect of high mortality levels on fertility behaviour is accounted for in the model. For example, if a woman were to lose her son before she completes her reproductive life, her son preference might get recoded as 1 and she would again be subject to the higher age-specific fertility rate that individuals with unmet son preference are subject to.⁴

To calibrate the model for India, we pool data from the fertility preferences questions across three waves of the National Family Health Survey (NFHS) (1992–93, 1998–99, 2005–06) to approximate a measure of son preference.⁵ The data reveal that mean son preference, as measured by the average of the ideal number of desired sons, fell at a rate of 3.7 % across successive 5-year cohorts from those born 1940–44 to 1970–74. This decline in mean son preference across cohorts was largely due to a decline in the proportion of women reporting a desire for more than one son (specifically declines in an ideal number of two or three sons). When we dichotomize this variable to see how proportions desiring at least one son vary by cohort the change is less salient at just under 0.2 % decline across successive younger birth cohorts. We interpret this as an indication that while the decline of higher order son preference (an ideal preference for two or three sons) may be related to a wider decline in the preference for large families, fertility decline does not equally rapidly erode norms surrounding a desire for one son.

Logistic regression analysis of the pooled NFHS data with the dichotomized son preference variable as the outcome revealed that both time of the survey and the cohort that the female respondent belonged to were statistically significant predictors ($p < 0.001$) of son preference. The impact of societal transformation and macro-level structural modernization factors, such as urbanization and educational expansion on demographic behaviour, are captured well by using the concept of cohort (Ryder 1965). Moreover, cohort changes in fertility preferences and behaviour (e.g. desired family size, contraceptive use) are well-documented across the developing world (Pasupuleti and Pathak 2011; Cleland et al. 1994). While cohort effects are statistically significant for son preference decline in India, period effect size is larger. This is in line with previous research in South Korea that also found period effects to be larger than cohort effects in the decline of son preference (Chung and Das Gupta 2007). We use the coefficients from the logistic regression to

⁴Whether her son preference gets recoded or not will depend on prevailing period- and cohort-levels of son preference for that time-step in the simulation. This takes into account the fact that a woman who might have had a son preference at one time-step and have acted on it then might not have the preference in a later period.

⁵The National Family Health Survey (NFHS) is the name by which the Demographic and Health Survey (DHS) is called in India.




Fig. 12.2 Proportion with son preference, India 1980–2050. Observed values from fertility preference data from National Family Health Survey (1992–1993, 1998–1999, 2005–2006)

project both backward and forward for time-steps and for cohorts for which we do not have observed data. The predicted proportions of individuals with son preference in the population at time t and from cohort c can be expressed as:

$$sp_c(t) = \frac{e^{\beta_0 + \beta_c + \rho t}}{1 + e^{\beta_0 + \beta_c + \rho t}} \quad (12.1)$$

In Eq. (12.1), β_c refers to the coefficient of son preference change by cohorts, ρ refers to the coefficient of son preference change over time and β_0 to the intercept. The three parameters β_0 , β_c and ρ are estimated from a logistic regression fitted to the NFHS data with the two covariates. Figure 12.2 plots mean son preference over time extrapolated backward to 1980 (the first year of the simulation) and forward until 2050 from the coefficients of the logistic regression. Since the proportion of son preferring individuals at any given time t is different for different cohorts, mean son preference here indicates the simple mean of son preference individuals aged 15–50 in that particular year (all fertile cohorts).

For South Korea, we obtain a measure of son preference that is slightly different than the “ideal number of sons desired” measure we use for India. For Korea, we use a measure which asks women if they feel they “must have a son”. This question has been asked in Korean fertility surveys carried out by the Korean Institute for Health and Social Affairs. Unfortunately, as we were unable to obtain the micro-data for these surveys we cannot explore the cohort and period effects separately in the decline of son preference as we do for India. Instead we rely on period (time) trends in proportions stating they must have a son reported in Chung and Das Gupta (2007). We fit a logistic regression with one covariate (time) to backward and forward project these son preference trends (see Fig. 12.3).

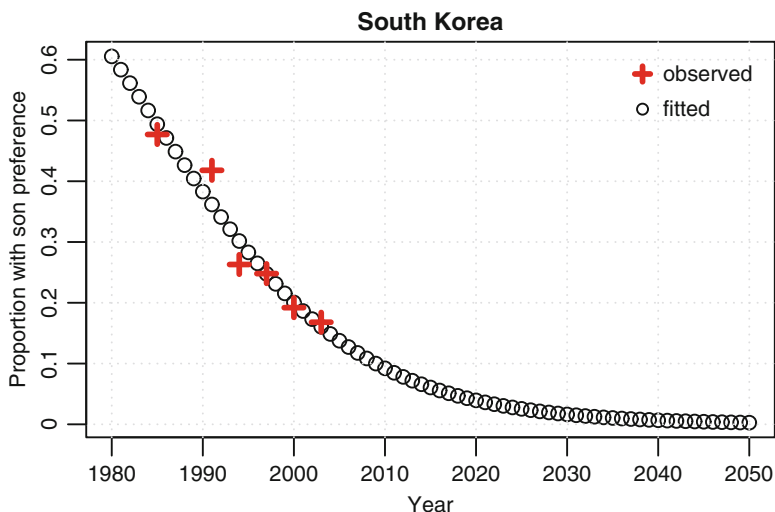


Fig. 12.3 Proportion with son preference, South Korea 1980–2050. Observed values from Chung and Das Gupta (2007) (Source: Kashyap and Villavicencio (forthcoming))

12.2.3 Differential Stopping Behaviour (DSB)

Differential stopping behaviour (DSB) is a common manifestation of son preference that is well documented in the literature (Amin and Mariam 1987; Arnold et al. 1998; Bongaarts 2013; Clark 2000; Retherford and Roy 2003). It is reflected in higher risk of parity progression and lower probability of using contraception for those who are sonless compared with those who already have sons. In the model, female agents with unmet son preference have a higher fertility risk, expressed as a deviation from the standard fertility schedule $h^*(x_i, t)$ by a proportional expansion factor $1 + \gamma$. When son-preferring female agents bear a son, their birth risk is adjusted down by a factor of $1 - \alpha$, indicating a reduced risk. For female agents with no son preference ($sp = 0$) the standard period age-specific rates $h^*(x_i, t)$ apply to them.

$$h_i(x_i, t) = \begin{cases} h^*(x_i, t) & \text{if } sp_i(t) = 0 \\ h^*(x_i, t) \times (1 + \gamma) & \text{if } so_i(t) < sp_i(t) \\ h^*(x_i, t) \times (1 - \alpha) & \text{if } so_i(t) \geq sp_i(t) \text{ and } sp_i(t) \neq 0 \end{cases} \quad (12.2)$$

As Eq. (12.2) shows, if the current number of sons $so_i(t)$ of female agent i is less than her son preference $sp_i(t)$, her period age-specific fertility rates $h_i(x_i, t)$

results from multiplying the standard schedule by a factor of $1 + \gamma$, where γ may be conceptualized as a son preference intensity parameter. For example, $\gamma = 0.3$ implies that a woman with unmet son preference experiences a fertility risk that is 30 % higher than the period age-specific schedule $h^*(x_i, t)$ that normally determines her risk for childbearing. A higher value of γ indicates a higher intensity of son preference through its impact on fertility behaviour. When an agent with son preference ($sp = 1$) has a son, her birth risk is adjusted downward by parameter α . For example, $\alpha = 0.15$ implies that a woman who meets her son preference experiences a fertility risk that is 15 % lower than the period age-specific schedule $h^*(x_i, t)$. In low fertility contexts, we would expect the birth risk adjustment factor α to take on a higher value compared with contexts where fertility levels are higher.

12.2.4 Access to Technology (Ability)

We use the logistic diffusion model, widely used to describe the diffusion of new technologies, to model the individual's *ability* or probability of gaining access to technology (Geroski 2000).

$$Ability(t) = \frac{e^{\nu(t-\phi)}}{1 + e^{\nu(t-\phi)}} \quad (12.3)$$

In Eq. (12.3) *Ability*(t) simulates an individual's probability of getting access to technology, which increases as a function of time t , where t corresponds to the time-step in the simulation (e.g. $t = 0, 1, \dots, 30$ for a 30-year simulation covering a period of 1980–2010), ν determines the slope or rate of increase, and ϕ the inflection point of the logistic diffusion curve. At the population level, *Ability*(t) can be interpreted as the proportion of individuals at a particular time-step gaining access to pre-natal sex-determination technology. At the initialization of the model no agent has access to technology ($tech_i(t) = \text{False}$ for all individuals at $t = 0$), but as *Ability*(t) gets recalculated at each time-step, a random number from a uniform distribution with a minimum value of 0 and maximum value of 1 is redrawn for each individual, which when less than *Ability*(t) sets the variable $tech_i(t) = \text{True}$ for that individual.

12.2.5 Fertility Decline (Readiness)

As fertility falls, norms surrounding smaller families become more entrenched. Individuals are likely to desire smaller families and if means are available to allow them to realize their son preference with a small family size, they are likely to do so. This is the motivating idea to generate an individual's propensity or *readiness*

to abort. Guilмото (2009) describes the readiness to abort as strongly related to the fertility squeeze felt by couples planning the size and composition of their families. From a modelling perspective, this fertility squeeze can be viewed as a form of social pressure that is closely related to prevailing total fertility levels and determines an individual’s readiness to abort. This readiness to abort is likely higher when fertility levels are lower and couples feel a greater “squeeze” or pressure to reconcile their son preference at lower parities than when average family size is higher and proceeding to higher parities is not out of step with prevailing total fertility norms.

If $sp_i(t) > so_i(t)$ and $tech_i(t) = \text{True}$, then

$$Readiness_i(t) = \begin{cases} \min \left\{ 1, \frac{\beta}{TFR(t-1)} \right\} & \text{if } parity_i(t) = 0 \\ \min \left\{ 1, \frac{parity_i(t) \times \sigma}{TFR(t-1)} \right\} & \text{if } parity_i(t) > 0 \end{cases} \quad (12.4)$$

Equation (12.4) shows how we model readiness to sex selectively abort. Readiness is the probability of an agent i to abort at time t and consequently lies between 0 and 1. An individual is only ready if she has unmet son preference (*willing*) and she has access to technology (*able*). When these two conditions are met, for an agent with $parity_i(t) > 0$, her *readiness* (probability) to abort depends on her current parity at the beginning of the period, $parity_i(t)$, and the most recent, prevailing, model-generated fertility levels, $TFR(t - 1)$.⁶ This ratio of an agent’s current parity to prevailing fertility levels is conceptualized as determining the extent of her fertility squeeze. When fertility levels are higher, for example at 3, a woman with unmet son preference and access to technology has a $0.33 \times \sigma$ probability to abort as she transitions from first to second parity compared with a woman who transitions from first to second parity when total fertility levels have fallen to 2.5 and the probability is $0.4 \times \sigma$.⁷ The parameter σ allows us to assess the impact of the fertility squeeze by scaling it up or down on SRB trajectories when calibrating the model. It allows us to account for the possibility that even if the fertility squeeze may be present in a population, there might be other counteracting forces, such as religious or cultural taboos against the practice of abortion or punitive measures against sex-selective abortion, that may not allow for the full extent of the fertility squeeze to be felt. Conversely, higher σ values indicate a greater intensity of the fertility squeeze.

Indeed, in some situations, particularly as fertility becomes very low, we may expect some individuals to abort at the lowest possible parity—parity 0 or before the transition to first birth. Although abortions before the first birth may become

⁶Since an agent’s current parity is determined at the beginning of the period t and TFR is calculated at the end of the period t , the most recently observed TFR with respect to an agent’s parity when she is at risk of childbearing in that period is $TFR(t - 1)$.

⁷At $parity_i(t) = 1$, the probability to abort will be $1/3 \times \sigma = 0.33 \times \sigma$ when $TFR(t - 1) = 3$, and $1/2.5 \times \sigma = 0.4 \times \sigma$ when $TFR(t - 1) = 2.5$.

more frequent as fertility falls, the literature suggests that these tend to be rarer events than higher parity abortions (Guilmoto 2009, p. 533). We therefore model parity 0 abortions as a function of prevailing fertility levels but subject to their own fertility squeeze parameter β than higher parity abortions, which are controlled by the parameter σ . The parameter β models to what extent individuals are ready to abort at parity 0. This allows us to generate scenarios where parity 0 abortions might be relatively low whilst higher parity abortions very high, which is in line with empirical evidence. Parity 0 abortion, while rare, is likely to be higher in contexts such as China or South Korea, in contrast with a context such as India. Even across the very low fertility settings, China, with a one-child policy where progression to the second parity for a large number of couples is restricted by law, a higher value of β to generate higher probabilities of parity 0 abortion would be likely, compared to South Korea where first births only displayed marginally skewed SRBs (Croll 2000, p. 43). The parameter β allows us to control this context-specific variation in parity 0 abortion and moderate the extent of the fertility squeeze at parity 0.

Figure 12.4 provides a diagram of the simulation process indicating which decisions are taken at each simulation step. Table 12.1 lists the relevant parameters that control the three RWA components—son preference, fertility decline and technology availability—in the model.

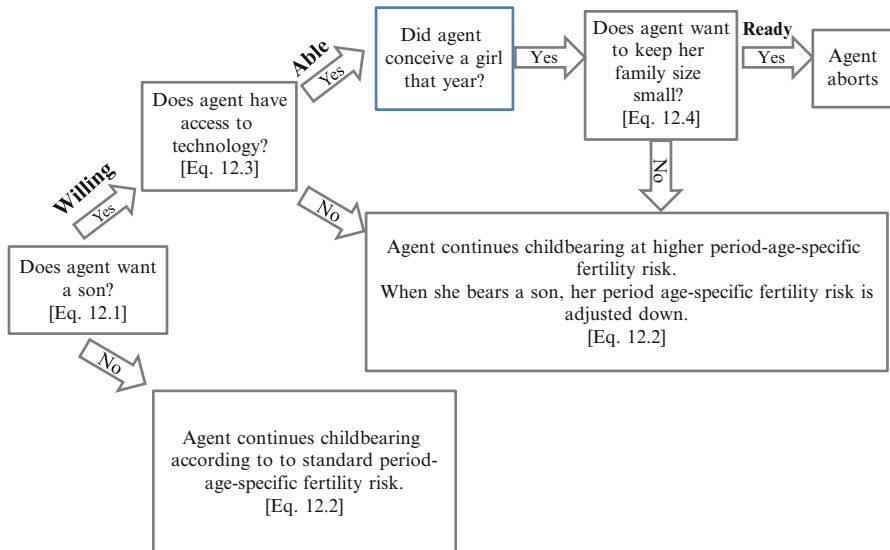


Fig. 12.4 Diagram of the simulation process at each time-step (Source: Kashyap and Villavicencio (forthcoming))

Table 12.1 Agent-based model parameters

Category	Parameters	Description
Son preference	γ	Son preference intensity
<i>Willingness</i>	α	Birth risk adjustment (when son preference realized)
Fertility decline	σ	Fertility squeeze scaling factor for parity 1 and higher
<i>Readiness</i>	β	Fertility squeeze at parity 0
Access to technology	ν	Rate of technology diffusion
<i>Ability</i>	ϕ	Inflection point for technology availability

12.3 Model Description

12.3.1 State Variables

The model comprises individual agents who each have an identity number (*id*), age (*x*), sex (*s*), cohort (*c*), son preference (*sp*), parity (*p*), sons (*so*), technology access (*tech*) and abortions (*ab*). The model is programmed in R (R Core Team 2015) and the code is optimized to take advantage of parallel computing using the `snowfall` R package (Knaus 2013). Table 12.2 lists agent’s state variables and their values. The model is initialized as a one-sex model with 100,000 initial female agents, but then as we model male as well as female births, from the first time-step onward the model becomes two-sex. We do not model partnership formation and focus only on reproductive behaviour. While partnership formation is no doubt important for reproductive behaviour, modelling partnerships and households would add complexity to the model that is not necessary for its immediate purpose.

An agent’s identity number, sex, and cohort are assigned to the agent at birth and remain the same throughout her life course. Cohorts are categorized by 5-years (1930–34, 1935–39, ...). The variable parity (*p*) corresponds to the current number of children of the female agent. Sons (*so*) refers to the number of sons she has. Technology access (*tech*) is a Boolean variable that takes a True or False value depending on whether an agent has access to pre-natal sex-determination technology. Abortions (*ab*) indicates how many sex-selective abortions a female agent has over the course of her life. An agent’s age, parity, son preference, sons, access to technology, and abortions values are time-varying and are updated at each time-step in the model.

When calibrating the model for South Korea and India, we are interested in approximating the shape and levels of observed SRB trajectories from 1980 onward when trajectories began to rise from normal to higher levels. As the abortion procedure depends on an individual’s parity it was important to approximate an initial parity distribution that closely approximates the observed population structures for both countries in 1980. To do this, we initialized the model 35 years earlier in 1945 to allow all women in reproductive ages (15+) to complete their fertility careers by 1980 and have their children belong to the starting population of 1980. Using UN data, we approximate the population structure of 1945 and

Table 12.2 Agent's state variables

Agent's state variables	Variable name	Values
Identity number	<i>id</i>	1, 2, 3, ...
Age	<i>x</i>	0–50
Sex	<i>s</i>	1: female 2: male
Cohort	<i>c</i>	5-year cohort (1930–34, ..., 2005–09)
Son preference	<i>sp</i>	0, 1
Parity	<i>p</i>	0, 1, 2, ...
Sons	<i>so</i>	0, 1, 2, ...
Access to technology	<i>tech</i>	1: True 0: False
Abortions	<i>ab</i>	0, 1, 2, ...

Source: Kashyap and Villavicencio ([forthcoming](#))

start a simulation process in which individuals die and reproduce according to their age-specific death and fertility rates for each year.⁸ The abortion procedure is not modelled prior to 1980 as due to limited technology availability there is little evidence to suggest SRBs were distorted before. The resulting population structure obtained in 1980 is very close to the population structure for both countries reported in UN data, and the minor differences that persist are likely attributable to migration dynamics that are not modelled.

12.3.2 Procedures

The model contains two procedures for agents: ageing and reproduction that are carried out at each time-step or tick in the model. Each tick corresponds to one year. At each tick an agent ages by one year. Since we focus on reproductive behaviour, all female agents die off in our model after reaching the maximum age of reproduction set at age 50. Male agents die at age 50 as well. The model is adapted to use death and fertility rates, from the United Nations World Population Prospects database, which are issued for 5-year age groups for 5-year periods (United Nations 2013). By simulating male agents until age 50, we can account for child and young adult mortality for males, which may have an impact on a woman's reproductive behaviour, as has already been mentioned.

The reproduction procedure models conception and birth. As defined in Eq. (12.2), the risk of childbirth $h_i(x_i, t)$ for each individual i is determined by age-specific fertility rates for that period t , and the current parity and son preference of the individual. The sex of the birth is determined at the point of conception

⁸The UN World Population Prospects data only offer data from 1950 onwards. Consequently, we approximated the population structure of 1945 with the one from 1950, and for the period 1945–1950 we used the same age-specific fertility rates and death rates as for the period 1950–1955.

by a probability of 0.5122 for male births and 0.4878 for female births, which corresponds to an SRB of 105. As we model sex-differential birth stopping behaviour and as our model moves in time and technology becomes available, we model the opportunity for female agents to have sex-selective abortions.

12.4 Results

We present results from simulations calibrated with mortality, fertility and son preference data for South Korea and India. For each country, we seek to approximate the son preference intensity, the extent of the fertility squeeze and associated abortion probabilities, and rates of technological diffusion that plausibly generated their specific SRB trajectories. We validate both calibration attempts against UN data on SRBs and fertility levels for the time period from 1980 to 2013 for each country provided in the UN World Population Prospects database (United Nations 2013). We also project our model forward after 2013 until 2050, continuing with the parameters and assumptions that help us approximate trends between 1980 and 2013. As described in the introduction, both India and South Korea have had distinctive SRB transitions. While South Korea's SRBs rose earlier and more rapidly, they also showed a remarkable turnaround in a very short period of time (see Fig. 12.1). In contrast, India has yet to undergo all three stages; although SRB trajectories first became skewed, and have leveled off since the mid-to late-2000s, they have yet to show the rapid turnaround they did in South Korea. Further, the rise in Indian SRBs was not as steep as the one in South Korea with peak SRBs remaining well below those observed in South Korea. We now explore what micro-level dynamics caused each of these distinctive trajectories.

12.4.1 *The South Korean Sex Ratio Transition*

Son preference data from South Korea indicate that the proportion of married women stating they “must have a son” declined from 48 % in 1985 to 26 % in 1994; over the same time however South Korean SRBs rose from about 108 to 114 and fertility declined from 2.1 to 1.6. It is interesting to note that son preference levels appear to be quite low ($\sim 30\%$ of the population stating “must have a son”) as SRBs reached their peak in 1990.

Figure 12.5 shows SRB and fertility trajectories from the model calibrated with UN mortality and fertility rates for South Korea, and son preference time series from Korean fertility surveys, across two sets of parameter values, one with $\beta = 0$ and the other set with $\beta = 0.2$. We validate⁹ our model by trying to match as much as possible the smoothed UN estimates shown in Fig. 12.5, which are indicated by

⁹Details on model calibration are available in Kashyap and Villavicencio (forthcoming).




Fig. 12.5 Simulated fertility (*solid line*) and sex ratio at birth (*dotted-dashed line*) trajectories, 5-year moving averages, South Korea 1980–2050

the lightest lines in grey in the figure.¹⁰ The simulated $\beta = 0$ curve in Fig. 12.5 shows 5-year moving averages of SRBs averaged across 25 simulation runs, with the following parameter values: $\gamma = 0.20$, $\alpha = 0.075$, $\nu = 0.5$, $\phi = 7$, $\sigma = 1.7$ and $\beta = 0$. These parameters allow for a rapid diffusion of technology with saturation of technology by the early 1990s and a very intense fertility squeeze at parity 1 and higher but with no parity 0 abortion ($\beta = 0$). A value of $\sigma = 1.7$ implies that given the levels of fertility decline between 2.4 to 1.6 between 1980–1990 when SRBs steadily rose, probabilities to abort for those with access to technology and unmet son preference who had a daughter as their first child rose from $\sim 71\%$ to 100% .¹¹ At parity 2 and higher and $\sigma = 1.7$, all individuals with an unmet son preference and access to technology transitioning to third birth would practice sex-selective abortion over the whole period. With abortion delayed to parity 1 and higher with

¹⁰UN SRB and fertility trends are issued for 5-year periods (e.g. 1970–75, 1980–85) where the 5-year period is assumed to run from July to July (e.g. July 1 1970–July 1 1975) and the estimates are assumed to refer to the mid-point of the period concerned (1 January 1973). Our model runs produce yearly estimates. To enable a better comparison between the UN estimates and our simulated results, we present smoothed UN estimates, in which we linearly interpolate the 5-year values for each year. For this, as advised in UN metadata we assume the value for each of the 5-year intervals corresponds to mid-point within the interval (e.g. 1970–75 rate corresponds to mid-point 1973).

¹¹If $TFR = 2.4$, $\min\{1, (parity \times \sigma / TFR)\} = \min\{1, (1 \times 1.7 / 2.4)\} = 0.708$. If $TFR = 1.6$, $\min\{1, (parity \times \sigma / TFR)\} = \min\{1, (1 \times 1.7 / 1.6)\} = \min\{1, 1.0625\} = 1$.

$\beta = 0$, even as son preference was declining, the steady diffusion of technology in the 1980s combined with the increase in the probabilities to abort at parity 1 over the period lead to increasing SRB distortions in the model between 1980 and 1990 as technological diffusion progresses to saturation by the early 1990s.

The most salient observation to be made from this simulated SRB curve with $\beta = 0$ is that while it matches a part of the rise in SRBs between the early to the mid-1980s due to technological diffusion and increasing readiness to abort at parity 1 due to declining fertility levels, it does not match the peak SRB levels observed in South Korea in 1990s. We fall about 2–3 points short of the peak, reaching 111–112 instead of the observed 114. We cannot match these peak SRB levels of 114–115 by adjusting technological diffusion parameters or increasing σ . By adjusting these parameters we worsen the fit over the 1980s with a curve that rises too early.

Instead, gradual adjustments in β that allow for small increasing probabilities of sex-selective abortion at parity 0 keeping all other parameters the same as before enable us to replicate the peak and turnaround in SRBs observed in South Korea. The simulated $\beta = 0.2$ SRB and TFR levels reported in Fig. 12.5 share all parameters in common with $\beta = 0$ simulated curve, except for their difference in allowing sex-selective abortion at parity 0 (these parameters are: $\gamma = 0.20$, $\alpha = 0.075$, $\nu = 0.5$, $\phi = 7$ and $\sigma = 1.7$). While our simulated trajectories do not match perfectly with the observed ones reported in UN data, they follow the general shape that was observed in South Korea between 1980 and 2000 relatively well. After the mid-2000s our simulated estimates fall slower than was observed in South Korea. The logistic fit for the son preference data for South Korea shown in the topmost panel of Fig. 12.6 shows that for 1991 our fit under-predicts son preference compared with the observed value, which may be partly why our turnaround starts slower.

A son preference intensity parameter of $\gamma = 0.20$ and birth risk adjustment of $\alpha = 0.075$ allow us to match Korean fertility trajectories over the period, with the prevailing levels of son preference very well. This implies that for those with unmet son preference fertility risk was higher by about 20%. The UN projections of the SRB for South Korea hold it constant at 107—slightly higher than what is considered biological normal levels of 104–106—starting 2010 until 2050. Allowing the underlying behaviours to change according to the trends reported in Fig. 12.6, our model suggests that SRBs will fall even lower towards levels of 105–106 after the mid-2020s.

12.4.2 *The Indian Sex Ratio Transition*

In contrast to the South Korean sex ratio transition, SRBs measured at the national-level in India peaked later and stayed lower at levels of 111. Moreover, unlike the South Korean case, SRBs appear to have leveled off in recent years but have not




Fig. 12.6 Son preference (*willingness*), access to technology (*ability*) and parity-specific abortion probabilities (*readiness*) over time, 1980–2050, for South Korea used in simulations. Modified from Kashyap and Villavicencio ([forthcoming](#))

shown the dramatic turnaround they did in South Korea. Son preference levels, using the measure we calibrate for the model from the NFHS, indicate that son preference was higher than levels in South Korea and declined slower in India.¹² What does

¹²Although here we treat them as roughly approximating proportions desiring one son at different time points in the model, we acknowledge these data come from different surveys and from questions regarding sex preference that were worded very differently in each of the surveys.

the model suggest were the reasons that SRBs did not reach levels comparable to those in South Korea? The fact that SRBs reached 111 when fertility levels were over 2.5 children per woman in India compared with the South Korean case where SRBs peaked when TFR was well below 2 suggest that the fertility squeeze was lower in India than in Korea. Individuals were much less likely to abort at lower parities, with evidence suggesting the greatest distortion was observed at parity 3 and higher (Arnold et al. 2002). Secondly, technological diffusion happened much quicker and more rapidly in Korea than India due to South Korea's comparative wealth and development. The literature widely acknowledges that technology was not widespread in the 1980s and diffused steadily in the 1990s in India (Arnold et al. 2002; Patel 2007). To approximate a general shape for the technological diffusion curve over the time from 1980 onward for India we rely on data points on proportion of women reporting ultrasound use during most recent pregnancy from the DLHS-I-III data reported in Akbulut-Yuksel and Rosenblum (2012).¹³ The best fit to these data points for a logistic diffusion curve from Eq. (12.3) used to model access to technology is of $\nu = 0.13$ and $\phi = 33$.

Figure 12.7 reports the dynamic trajectories and associated parameters of *willingness* (son preference), *ability* (access to technology) and *readiness* (parity-specific abortion probabilities) that we use to simulate SRB and fertility trajectories for India. The results from these simulations, 5-year moving averages across 25 simulations, are reported in Fig. 12.8. We use two sets of parameters for India: in simulations with the first set of parameters we use the technological diffusion parameters of $\nu = 0.13$ and $\phi = 33$ estimated from the technological diffusion data points from the DLHS-I and II, and test the model with different values of the readiness parameters σ and β to find the abortion probabilities that provide the best possible match to the UN estimates of the Indian SRB curve from 1980–2013, which are depicted by the lightest curve in grey in the Fig. 12.8. The best possible fit obtained with $\nu = 0.13$ and $\phi = 33$ is reported in the simulated curve with $\nu = 0.13$, $\phi = 33$ and $\sigma = 0.7$ in Fig. 12.8. To improve upon the fit provided by these parameters, we adjust technological diffusion by allowing for an earlier inflection point with $\phi = 27$, and a slightly slower rate of increase with $\nu = 0.12$, and a slightly lower fertility squeeze parameter of $\sigma = 0.65$. In both simulated SRB curves depicted in Fig. 12.8 we keep $\beta = 0.05$ to allow for negligible parity 0 abortion, and hold son preference intensity parameter $\gamma = 0.125$ and the birth risk adjustment parameter $\alpha = 0.075$.

¹³The authors of that paper report this proportion by regions for India and estimate this proportion for births during each year between 1999 and 2008. We average across all regions to approximate national-level technological diffusion. We recognize this erases the heterogeneity across different regions and in future amendments to the model we hope to incorporate issues surrounding regional heterogeneity in the model. We discuss these issues further at the end of the chapter. We should also note that we could not find similar data for South Korea so we estimated technological diffusion for South Korea by conducting sensitivity analyses for the model across a range of technology diffusion (ν and ϕ) parameters.

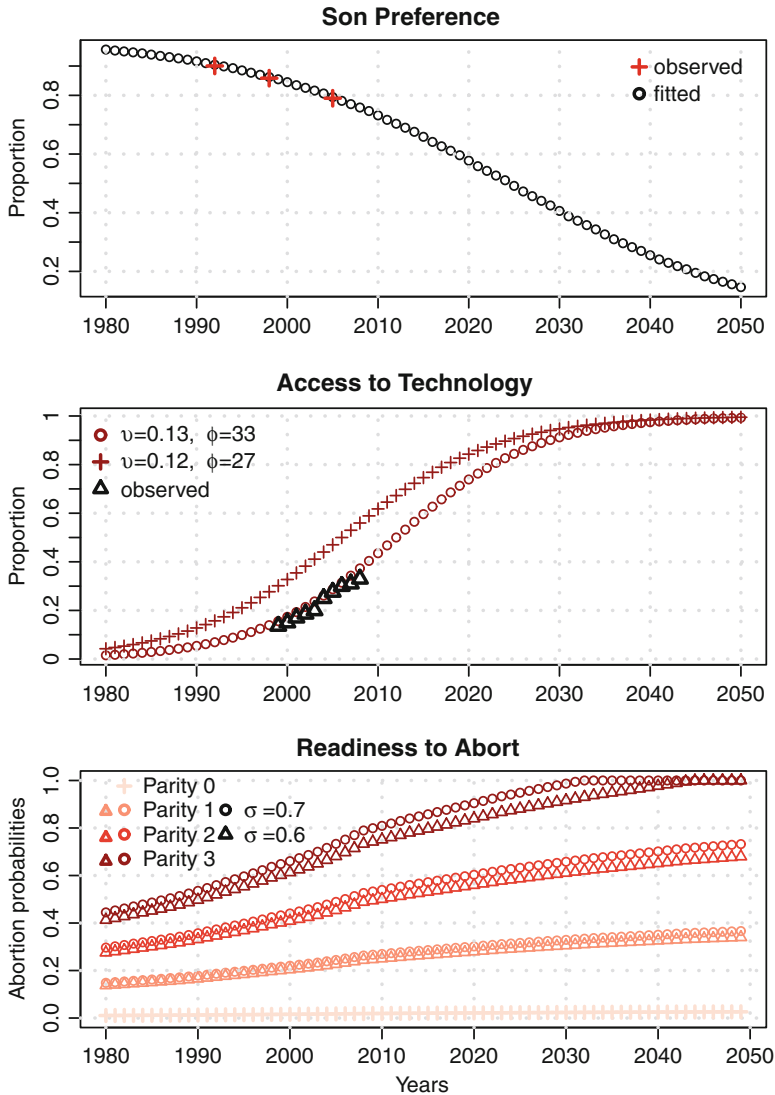


Fig. 12.7 Son preference (*willingness*), access to technology (*ability*) and parity-specific abortion probabilities (*readiness*) over time, 1980–2050, for India used in simulations

The simulated SRB curve with earlier technology diffusion $\nu = 0.12$, $\phi = 27$ and lower $\sigma = 0.65$ provides a better fit to Indian SRB trajectories between 1980 and 2013 than the curve generated by $\nu = 0.13$, $\phi = 33$ and $\sigma = 0.7$. The model fit demonstrates how the intensity of the fertility squeeze was much lower in India than in South Korea, with values for σ as well as β significantly lower than those used




Fig. 12.8 Simulated fertility (*solid line*) and sex ratio at birth (*dotted-dashed line*) trajectories, 5-year moving averages, India 1980–2050

for model fitting in South Korea. Our model suggests that few individuals likely aborted at low (parity 1 or parity 2) or very low parities (negligible at parity 0) in India (see probabilities to abort in Fig. 12.7). The rise in SRBs observed between the mid-1980s and 2010 was caused by the diffusion of technology starting in the 1990s and small increases in the probability to abort that accompanied the fertility decline, with the greatest probabilities to abort in the transition from second to third birth, and higher. The slower decline in son preference in India accompanied by a slower, flatter diffusion of technology generates a flatter SRB trajectory with a less sharp peak compared with South Korea. It is interesting to note that the SRB peak in South Korea roughly coincided with the saturation of technology. Peak SRBs are generated by the model in the mid-2000s when technological access is $\sim 40\%$ for the simulated $\nu = 0.12, \phi = 27$ and $\sigma = 0.65$ curve, and in 2010 for the simulated $\nu = 0.13, \phi = 33$ and $\sigma = 0.7$ SRB curve, when technological access is also $\sim 40\%$. Thereafter, even as technological access and abortion probabilities increase, SRBs begin to level off due to declining son preference.

The fertility trajectories generated by the best fit parameters ($\nu = 0.12, \phi = 27$ and $\sigma = 0.65$) are plotted in Fig. 12.8. These match UN TFR trajectories shown in the lightest grey curve particularly well indicating that a good balance between son preference intensity that amplifies fertility and the fertility squeeze generated

abortion probabilities that exert a downward influence on them. For the model calibrated for India, if an individual is son preferring their fertility rates are higher by 12.5 % ($\gamma = 0.125$) and upon meeting their son preference, their rates fall by 7.5 % ($\alpha = 0.075$). Following the best fit simulated SRB curve ($\nu = 0.12$, $\phi = 27$ and $\sigma = 0.65$) for the period after 2010 indicates that SRB trajectories are likely in the midst of a turnaround presently in India, and will gradually fall towards levels of 107 by 2050.

12.5 Methodological Challenges, Discussion and Future Work

The agent-based model presented in this chapter applies the insights of the *ready, willing and able* framework to the practice of sex selection to generate sex ratio at birth distortions at the macro-level from bottom-up. The model enables us to quantify the dynamic trajectories of son preference, technology diffusion and abortion probabilities as a function of the fertility decline that undergird the sex ratio transition. We are able to generate the shape and match levels of SRB trajectories at their different stages—the rise from normal levels, levelling off and their decline and return towards near-normal or normal levels—for the distinctive transitions of South Korea and India. The model illustrates a number of interesting insights. We show how son preference can be declining in a population but SRBs can nevertheless rise. We find that SRB levels are highly sensitive to sex-selective abortion practiced at low parities, and especially to abortions before the first birth at parity 0 as in the case of South Korea. Even relatively low levels of son preference of $\sim 30\%$, as we show in the case of South Korea, can result in significant SRB distortions when sex-selective abortion is practiced at low parities.

It is nevertheless important to be mindful of some of the challenges faced in developing the model. In seeking to develop the model, we generally relied on the approach to keep the model as simple as possible and minimize the number of parameters. The optimal number of parameters, however, was not clear at the outset and involved extensive sensitivity analysis across several parameter values to assess model behaviour and model fit. As we worked towards a more streamlined model, we carried out a number of robustness checks over different parameter values to ensure the model was behaving according to theoretical expectations. Approaches to model calibration and sensitivity analysis with the model are detailed elsewhere (Kashyap and Villavicencio [forthcoming](#)).

When designing and implementing an ABM, one significant recurring issue for model designers is the trade-off between model parsimony and complexity. By trying to adopt a parsimonious modelling approach, we may have simplified some processes in a manner that is either not easily approximated with existing data or perhaps a bit unrealistic. For example, we chose to model son preference dichotomously and its effect on fertility in time-constant γ (son preference intensity)

and α (birth risk adjustment after meeting son preference) parameters. From a modeling and a theoretical perspective there is a strong case to be made, as we have tried to do, about incorporating son preference dichotomously and allowing its effects on fertility to be constant across time. As the effect of son preference on fertility at any time point is decided dynamically as a balance between the proportion son-preferring at that time and their upward effect on total fertility levels, as well as those meeting their son preference who exert a downward effect, we sought to measure and calibrate model-generated TFR trajectories with observed TFR levels as a validity check in the model. These model-generated TFR values also informed the abortion probabilities that affected the SRB levels, which further endogenized model processes. Our reasoning was that if the survey data approximations of son preference were reasonable, then plausible, but not very small, values of time-constant γ and α parameters would allow us to generate simulated fertility trajectories that matched observed TFR and SRB trajectories from UN data.

A significant issue when designing and implementing the ABM, related to the parsimony and complexity trade-off, is the extent of heterogeneity that is necessary and useful for the model. For instance, we do not explicitly account for regional heterogeneity in the model in its current stage. Our primary reason for doing this was because we were keen to use standardized UN data to enable the model to be applied in different contexts, as well as validate against SRB trajectories over time that are readily available at the national level. On the one hand, we account for heterogeneity in the micro-level processes driving the macro-outcome in the model over time and regional heterogeneity may be seen as one implicit component of this overall heterogeneity. Nevertheless, an aggregate picture at the national level may be composed of extremely disparate trends at the regional level, whose impact cannot be assessed in the current version of our model. In India, for example, the regional nature of son preference has been widely noted in the literature, with son preference being concentrated in the north and northwest of the country (Das Gupta et al. 2003). Although national-level SRBs never reached levels comparable to those in South Korea and China, those in northwest India did (Guilmoto 2009). Thus, one possible reason why national-level technological diffusion parameters fitted to survey data do not fit the SRB curve observed quite as well may be due to regional differences in technological diffusion that are not explicitly accounted for in the model.

A word about model outcome uncertainty: although the ABM gives a good understanding of the plausible values of different micro-levels factors that generate SRB distortions, the model does not provide, at least in its current stage, any estimate of the uncertainty surrounding model outcomes. We can nevertheless use the ABM and its flexibility to experiment with values of different levels of son preference, fertility decline and technological diffusion and explore the impact of these on SRB levels and trends.

Despite these limitations, we believe the ABM presented here provides a contribution to the literature that is powerful in the insights it provides and flexible enough to be adapted to different contexts. While each geographical context has

its own set of underlying factors that will determine the course of future SRB trajectories, the ABM presented here provides a valuable computational tool that will enable researchers to test the implications of unique sets of micro-level forces—individual fertility preferences, behaviours, and technology use—to explain both current macro-level trends as well as project future ones.

We hope to develop the model further and explicitly engage with some of the limitations we have identified here in future work. We encourage other researchers working in the field to make use of and extend our model and develop ways to address its methodological issues and uncertainties.

References

- Akbulut-Yuksel, M., & Rosenblum, D. (2012). *The Indian ultrasound paradox*. SSRN Scholarly Paper ID 1989245. Rochester: Social Science Research Network.
- Amin, R., & Mariam, A. (1987). Son preference in Bangladesh: An emerging barrier to fertility regulation. *Journal of Biosocial Science*, 19(02), 221–228.
- Arnold, F., Choe, M. K., & Roy, T. K. (1998). Son preference, the family-building process and child mortality in India. *Population Studies*, 52(3), 301–315.
- Arnold, F., Kishor, S., & Roy, T. K. (2002). Sex-selective abortions in India. *Population and Development Review*, 28(4), 759–785.
- Attané, I., & Guilmoto, C. Z. (Eds.) (2007). *Watering the neighbour's garden: The growing demographic female deficit in Asia*. Paris: CICRED, CEPED.
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. D. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the united kingdom with Gaussian process emulators. *Demographic Research*, 29(27), 729–766.
- Billari, F. C., & Prskawetz, A. (2003). *Agent-based computational demography: Using simulation to improve our understanding of demographic behaviour*. New York: Springer.
- Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3), 7280–7287.
- Bongaarts, J. (2013). The implementation of preferences for male offspring. *Population and Development Review*, 39(2), 185–208.
- Chung, W., & Das Gupta, M. (2007). The decline of son preference in South Korea: The roles of development and public policy. *Population and Development Review*, 33(4), 757–783.
- Clark, S. (2000). Son preference and sex composition of children: Evidence from India. *Demography*, 37(1), 95–108.
- Cleland, J., Onuoha, N., & Timæus, I. (1994). Fertility change in sub-Saharan Africa: A review of the evidence. In T. Locooh & V. Hertich (Eds.), *The onset of fertility transition in sub-Saharan Africa* (pp. 1–20.) Liège: Ordina Editions.
- Coale, A. J. (1973). The demographic transition reconsidered. In *Proceedings of the International Population Conference 1973* (pp. 53–72.) Liège: Ordina Editions.
- Croll, E. (2000). *Endangered daughters: Discrimination and development in Asia*. London/New York: Routledge.
- Das Gupta, M., Zhenghua, J., Bohua, L., Zhenming, X., Chung, W., & Hwa-Ok, B. (2003). Why is son preference so persistent in East and South Asia? A cross-country study of China, India and the Republic of Korea. *The Journal of Development Studies*, 40(2), 153–187.
- Duthé, G., Meslé, F., Vallin, J., Badurashvili, I., & Kuyumjian, K. (2012). High sex ratios at birth in the Caucasus: Modern technology to satisfy old desires. *Population and Development Review*, 38(3), 487–501.
- Geroski, P. A. (2000). Models of technology diffusion. *Research Policy*, 29(45), 603–625.

- Guilmoto, C. Z. (2009). The sex ratio transition in Asia. *Population and Development Review*, 35(3), 519–549.
- Guilmoto, C. Z. (2012). Son preference, sex selection, and kinship in Vietnam. *Population and Development Review*, 38(1), 31–54.
- Kashyap, R., & Villavicencio, F. (forthcoming). The dynamics of son preference, technology diffusion, and fertility decline underlying sex ratio at birth distortions: A simulation approach. *Demography*.
- Kennedy, M. C., & O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3), 425–464.
- Knaus, J. (2013). Snowfall: Easier cluster computing (based on snow). R package version 1.84–6.
- Lesthaeghe, R., & Vanderhoeft, C. (2001). Ready, willing, and able: A conceptualization of transitions to new behavioral forms. In J. B. Casterline (Ed.), *Diffusion processes and fertility transition: Selected perspectives* (pp. 240–264). Washington, DC: The National Academies Press.
- Nath, D. C., & Land, K. C. (1994). Sex preference and third birth intervals in a traditional Indian society. *Journal of Biosocial Science*, 26(03), 377–388.
- Pasupuleti, S. S., & Pathak, P. (2011). Spatial and temporal changes in fertility behavior of Indian women cohorts. *Genus*, 66(3), 62–74.
- Patel, T. (2007). *Sex-selective abortion in India: Gender, society and new reproductive technologies*. New Delhi: SAGE.
- R Core Team (2015). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing.
- Retherford, R. D., & Roy, T. K. (2003). *Factors affecting sex-selective abortion in India and 17 major states*. Number 21 in National Family Health Survey Subject Reports. Mumbai: International Institute for Population Sciences and Honolulu: East-West Center.
- Ryder, N.B. (1965). The cohort as a concept in the study of social change. *American Sociological Review*, 30(6), 843.
- United Nations. (2013). *World Population Prospects: The 2012 Revision* (DVD Edition). New York: United Nations, Department of Economic and Social Affairs, Population Division.

Chapter 13

Exploring the Role of Communication in Shaping Fertility Transition Patterns in Space and Time

Sebastian Klüsener, Francesco Scalone, and Martin Dribe

13.1 Social Status, Communication Links, and the Fertility Transition in Space and Time

Research on historical demographic phenomena is currently benefiting from a number of developments that provide new perspectives on well-documented processes, such as the historical fertility decline during the demographic transition. In recent decades, there have been substantial advancements in the digitisation of historical census and vital registration data, particularly for north-western Europe and North America. Projects such as NAPP¹ or IPUMS² have helped to make historical census data accessible for researchers (see, e.g., Ruggles et al. 2011). In addition, because of the substantial advancements in historical Geographic Information Systems (GIS) and the automatised geocoding of address data, we are better able to control for spatial aspects of historical demographic processes. Parallel to the advancements in data availability, there have been significant improvements in computational

¹North Atlantic Population Project.

²Integrated Public Use Microdata Series.

S. Klüsener (✉)

Max Planck Institute for Demographic Research, Rostock, Germany

e-mail: kluesener@demogr.mpg.de

F. Scalone

Department of Statistical Sciences, University of Bologna, Bologna, Italy

e-mail: Francesco.Scalone@unibo.it

M. Dribe

Centre for Economic Demography, Lund University, Lund, Sweden

e-mail: Martin.Dribe@ekh.lu.se

power and modelling software. These developments allow today's researchers to analyse and simulate complex processes, such as the fertility transition, using large numbers of individual-level observations and a variety of modelling approaches. Among these approaches are (multilevel) event history, spatial econometric, and agent-based simulation techniques. All of these trends have contributed to a renewed interest among researchers in studying historical fertility decline patterns (e.g., Van Bavel 2004; van Poppel et al. 2012; González-Bailón and Murphy 2013; Dribe et al. 2014a, b; Goldstein and Klüsener 2014).

The main aim of this contribution is to demonstrate how the advancements in data availability and agent-based models can enable us to provide new insights into longstanding demographic debates about the factors that shape the decline in fertility during the demographic transition. Here we focus on the discussion of to what degree fertility decline patterns in space and time can be understood as being shaped by innovation or adjustment processes (Carlsson 1966). From an innovationist perspective, the decline in fertility in space and time is predominantly shaped by the diffusion of information and ideas that influence women in their decisions about whether to adopt deliberate, parity-specific fertility control strategies. The adaptationist perspective, by contrast, sees the adoption of fertility control strategies as a reaction to changing structural conditions, such as rising returns from focusing on the quality instead of the quantity of children, and changes in the costs associated with raising children. These two views are certainly not mutually exclusive, as in order to adapt to changing circumstances women need to learn about these changes. In most cases, women acquire knowledge not only through their own observations, but also through communication with others. At the same time, people who face strong adaptation pressures might be particularly likely to use communication strategies in their search for approaches to coping with these challenges. In addition, because people lack full access to information, there might be a temporal gap between the emergence of changes in structural conditions and the subsequent adaptation to these new conditions.

Based on Carlsson (1966), it could be expected that if information diffusion is the main force shaping the spatiotemporal decline patterns, these patterns should be structured by variation in communication pathways. If, on the other hand, adaptation pressure is the primary factor shaping the patterns, they should be characterised by spatiotemporal variation in changes in structural conditions. Proponents of the innovationist perspective have observed that once it has gained momentum, the fertility decline process tends to occur very rapidly within societies. Thus, it is very unlikely that the process can be directly linked to a single economic force to which people have adapted (Cleland and Wilson 1987). Meanwhile, proponents of the adaptationist perspective have stressed that, according to the outcomes of numerous econometric studies, economic factors are important predictors of fertility decline (e.g., Galloway et al. 1994).

It is relevant to point out that we believe the decline in fertility to be shaped by both innovation diffusion and adaptation to changing circumstances (see also Klüsener et al. 2013; Dribe and Scalone 2014). Thus, it is not our aim in this study to prove that fertility decline is driven by communication processes alone. Instead,

we intend to investigate which of the major spatiotemporal characteristics of the fertility decline in Sweden during the period 1880–1900, and across three different social classes (elite, farmers, workers and others), could have emerged as a result of communication processes that were structured by social and spatial variation in communication links.³ Related to this variation, our analysis will, for instance, provide support for the view that the elite had a higher density of communication links through space than the other social classes, and that the communication links to the cities were especially dense.

To achieve our objective, we run simulation models on a full individual-level sample of married women aged 20–49 who were living in Sweden in 1880. These models are based on a *ceteris paribus* approach in which we model fertility decline as a communication process within a static society that is not subject to adaptation pressure or other social change processes, apart from the simulated diffusion of the adoption of fertility control strategies. As proxies for communication links, we use information on migration links between regions (see Hägerstrand 1965; Rosero-Bixby and Casterline 1994). The results presented in this contribution are the initial outcomes of a larger research project in which we consider a range of scenarios that vary according to their starting conditions and assumptions on diffusion mechanisms.

Another motivation for our study is related to a recurring observation in spatial analyses of the fertility transition in various countries. Several studies have identified clusters of high fertility decline around early centres of the decline (e.g., big cities) that cannot be explained by the socio-economic characteristics of these areas (Schmertmann et al. 2010; Goldstein and Klüsener 2014; Costa 2014). These clusters might be caused by the diffusion of the adoption of fertility control behaviour as a result of social interaction (see also Hägerstrand 1965), or by unobserved variables representing changes in structural conditions that start to spread out from early centres of the decline as soon as the fertility transition gains momentum in this area. Because there is a lack of information on these processes, conventional regression methods offer few options for exploring the question of whether these clusters of the decline are indeed related to social interaction. Our simulation models, by contrast, allow us to investigate this question.

We believe that our study makes important contributions to the debate on the factors that shape the fertility decline patterns during the demographic transition. First, while we are certainly not the first to study the fertility transition with simulation techniques (González-Bailón and Murphy 2013; see also Casterline 2001), we are probably the first to model the fertility transition in space and time simultaneously for different social classes, while taking social and spatial variation in communication links into account. Second, our novel approach enables us to demonstrate that a substantial number of the major social class and spatiotemporal

³While our simulation models are driven by simulated communication processes only, in choosing our vanguard adopters for the starting conditions we do, however, consider in one scenario variation in adaptation pressure (see below).

characteristics of the fertility transition patterns that have been observed in Sweden (and in other countries around the globe) may have solely been shaped by communication processes that were structured by social and spatial variation in communication links. These findings are also relevant for research on other social change processes in which communication plays a role.

13.2 Theoretical Considerations

As our agent-based simulations focus on exploring the potential role of communication in shaping fertility transition patterns, we will restrict ourselves in the theoretical section to considerations on the interplay of social status and space in communication processes. For a more detailed account of how space and place might affect both spatial variation in communication diffusion and adaptation incentives, see, e.g., Klüsener et al. (2013). In our study we investigate a historical population among whom most of the social interactions were still local in character. Therefore, we assume that spatial distance substantially moderated both the frequency of social interactions and the quality of the information exchanged. Thus, especially in the early phase of the decline in Sweden, we would expect to find that the diffusion of ideas was particularly strong in areas where a group of pioneers had already adopted the behaviour, as we assume that these pioneers interacted more frequently with people in their own region of residence than with people in other regions (see also Hägerstrand 1965).

Our empirical analysis supports the general observation that big cities serve as important communication and transport hubs, both within a country and internationally. This is particularly the case for capital cities, which have links to other countries not only through trade networks, but also through diplomatic relations. Consequently, big cities might emerge as early centres of the fertility transition not only because urban populations generally are subject to greater adaptation incentives and pressures to reduce their fertility than rural populations, but also because information that might prompt people to adopt fertility control behaviours is likely to reach inhabitants of bigger cities much earlier than residents of rural areas.

In terms of social class differences, existing research has shown that elite groups tend to be early adopters in the fertility transition (Livi-Bacci 1986; Haines 1992; Bengtsson and Dribe 2014; Dribe and Scalone 2014). Farmers, on the other hand, tend to lag behind in the process (van Poppel 1985; Dribe et al. 2014a, b). Based on his research on social status differences in the fertility transition in Britain, Szreter (1996, p. 546 ff.) developed the concept of “communication communities”. He argued that the shifts in fertility behaviour in Britain were mediated by the membership of individuals in social groups who shared similar social norms and identities (e.g., occupational classes). For the working classes and farmers, these social groups were rather local in character; whereas for members of the middle class and the elite, these social communities often spanned across localities, and in some cases even

formed nation-wide networks. Based on these considerations, we would expect to find that information relevant to adopting a fertility control behaviour spread more quickly through space among elite groups than among farmers and workers. In addition, we would assume to find that social class boundaries limit the spread of information across social classes; an assumption for which we also find support in our analysis of the observed decline patterns (see below). These considerations on communication communities were highly relevant for the specification of our agent-based simulations. We should, however, acknowledge that Szreter (1996, p. 546) has expressed opposition to the idea that social classes are the right dimension for studying communication communities, arguing that these categories are too broad. But as it is not the aim of this contribution to grasp fertility decline in all of its details, we believe that the differentiation into broad social classes is justified when investigating the general spatiotemporal aspects of fertility decline by social status.

13.3 Data and Analytical Strategy

For our analysis we can draw upon complete individual-level datasets of the three Swedish censuses of 1880, 1890, and 1900. These datasets, which were prepared by the Swedish National Archives, are available through the web portal of the North Atlantic Population Project (Ruggles et al. 2011). For the purposes of mapping the observed and the simulated fertility decline patterns, we use a GIS file that provides the administrative boundaries of the 25 Swedish counties (*län*) during the period 1880–1900. This file has been derived from a GIS dataset of historical administrative boundaries in Sweden created by the Swedish National Archives.

Our simulations are based on the 1880 dataset, as 1880 is about the time when the fertility transition started in Sweden (Hofsten and Lundström 1976; Coale and Watkins 1986, Map 2.1). In addition, we use information from all three censuses to derive data on observed fertility changes between 1880 and 1900. We have data for 4.6 million individuals in the 1880 census, while the 1890 census covers 4.8 million and the 1900 census 5.2 million individuals. These individuals are grouped in households. Family members who were residing in the same household are linked by pointer variables. This allows us to connect a mother to her children and to her husband, if they are living in the same household.

As the majority of women in Sweden at that time did not provide occupational information, we decided to use the occupation of the husband to determine each woman's social status. Thus, in the analyses we consider only married women who were sharing a household with their husband at the time of the census. Social class is defined based on the husband's occupation, which is coded in HISCO (Historical International Standard Classification of Occupations; van Leeuwen et al. 2002) and is classified according to HISCLASS (Historical International Social Class Scheme; van Leeuwen and Maas 2011). We differentiate three social classes. The first class, *elite*, is comprised of individuals in elite and upper middle class occupations (HISCLASS 1–6). The second class, *farmers*, corresponds to HISCLASS 8. The

third class, *workers and others*, includes skilled, lower skilled, and unskilled workers; as well as all other groups (HISCLASS 7, 9–12). This third class is probably the most heterogeneous one. For the agent-based simulations, we consider just two other attributes in addition to social status: namely, the current region of residence of each woman and her region of birth. Information on fertility trends in the region of birth will be of relevance for our models, and we will not be able to model these trends for regions of birth outside of Sweden. We thus decided to exclude those women for whom the birth region is unknown or who were born outside of Sweden (around 2.1 % of all of the women who were chosen based on the described characteristics). The resulting total number of women (agents) in our simulation models is 488,438.

In obtaining information on observed fertility trends in the period 1880–1900, we were faced with the problem that the census data do not allow to calculate standard fertility rate measures. We therefore decided to use the child-woman ratio (CWR) as an indirect fertility indicator based on data on the number of linked children in the household. The CWR is commonly defined as the number of children aged 0–4 per woman aged 15–49 (Shyrock and Siegel 1980). For the 1900 census in Sweden, Scalone and Dribe (2012) compared fertility levels derived by applying the CWR with fertility levels obtained through other standard fertility measures (e.g., the total marital fertility rate), and using the own-children method as an alternative indirect method (see also Dribe and Scalone 2014). They demonstrated that the unadjusted CWR is a reasonably reliable indicator of socio-economic differentials in gross or total fertility. We decided to exclude married women under age 20, as we assume that these women had not been married for very long, and therefore likely had lower CWRs than other women.

While most of our diagnostics of observed and modelled characteristics are standard statistical measures, we also employ the Moran's I test of spatial autocorrelation to detect evidence of spatial clustering (i.e., spatial autocorrelation) of the fertility decline in specific regions of Sweden. The Moran's I is very similar to the Pearson's product moment correlation coefficient, except that we are not obtaining the correlation between two variables v and w for all regions i , but rather the correlation between the observed values of v in each region i and the mean value of v in regions j that are adjacent to region i (see also Bivand et al. 2013). Adjacency is defined as being one of the four nearest neighbouring regions.⁴ The Moran's I can take values from -1 (strong negative spatial autocorrelation) over 0 (no spatial autocorrelation) to 1 (strong positive spatial autocorrelation). If a pattern is characterised by a spatial clustering of regions with high or low levels, the Moran's I test would return elevated positive values.

⁴Measured by calculating the spherical distances between the regional geographical centroids.

13.4 Observed Spatiotemporal Fertility Decline by Social Class 1880–1900

Before we turn to the conceptual considerations for the agent-based simulation models, we will first present descriptive findings on the spatiotemporal aspects of the fertility decline in the three social classes in Sweden between 1880 and 1900. This information is instrumental for identifying the major characteristics of the observed fertility decline patterns that we attempt to reproduce with our simulations. In addition, the analysis of the observed patterns helps us to assess the validity of the assumptions that guide the specification of our simulation models.

We begin by looking at Table 13.1, which shows national-level means and regional variation in CWRs that were derived from the censuses of 1880, 1890, and 1900 for our three social classes and for all of the classes combined. Data are presented for both the obtained CWRs in the three census years (columns 1–3) and the changes between the censuses (columns 4–5). If we look at the trends for all women, it is visible that net fertility actually increased between 1880 and 1890. Given that fertility was actually decreasing in this period (Dribe 2009), we believe that this outcome is attributable to the decline in mortality among infants and children, which completely compensated for the fertility decline. Between 1890 and 1900, net fertility decreased by about 4%. However, fertility trends differed considerably between the social classes. In 1880, the CWR of the elite was still close to that of all women. But in the two decades that followed, the CWR declined by about 15% among the elite. Thus, in 1900 the elite group had the lowest CWR by far. While the CWR of working-class women increased between 1880 and 1890, it decreased thereafter. In 1900, working-class women had a CWR that was slightly lower than it was in 1880. Women in farming families experienced very modest declines in both decades.

To examine the regional variation in CWRs in the three social classes across the 25 Swedish counties, we present data for two measures of dispersion: the standard deviation and the Moran's I index described above. Based on our theoretical considerations, we would expect to find that the elite had the densest communication links through space; an assumption which is also supported by our analysis of migration links (see below). If ideational diffusion processes indeed played a role in the fertility decline, then these denser communication links among the elite should have contributed to a more homogenous spatial pattern in terms of both the CWR levels and the decline over time. The results for the standard deviation indicate that in all three censuses, the elites had the lowest standard deviation of CWR levels across regions, while the farmers had much higher levels. When we look at changes over time, we see that in the first decade the elite had a standard deviation that was above the average; a finding that is not in line with our expectations. This might be related to the fact that during this period the elite were the only social class who had already entered the transition, as regional variance tends to increase after the onset of the decline (Coale and Watkins 1986). In the second decade, when all of the social classes experienced a decline, the elite had indeed the lowest standard deviation across regions, even though they constituted the class who registered the sharpest decline in this period.

Table 13.1 Trends and regional variation in the child-woman ratio (CWR) by socio-economic status (1880–1900)

	CWR 1880	CWR 1890	CWR 1900	CWR change 1880–1890	CWR change 1890–1900
Mean (national level)					
Elite	0.98	0.93	0.83	−0.05	−0.10
Farmers	1.01	1.00	0.99	0.00	−0.01
Workers and others	1.02	1.05	1.01	0.03	−0.04
Total	1.01	1.02	0.98	0.01	−0.04
Standard deviation (25 regions)					
Elite	0.10	0.09	0.10	0.04	0.03
Farmers	0.16	0.17	0.16	0.04	0.04
Workers and others	0.10	0.09	0.11	0.03	0.05
Total	0.11	0.11	0.12	0.03	0.04
Moran's I index of spatial autocorrelation (25 regions)					
Elite	0.53	0.47	0.46	−0.05	−0.12
Farmers	0.41	0.43	0.42	0.05	−0.08
Workers and others	0.42	0.44	0.57	0.15	0.26
Total	0.42	0.41	0.49	0.22	0.16

Note: We only consider married women aged 20–49 born in Sweden with a spouse present in the household in the respective census (1880, 1890, 1900), and for whom the social status could be detected. In deriving the Moran's I we define the four nearest regions as neighbours. These were determined by calculating the spherical distance between the geographical centroids of the 25 Swedish regions

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

When we look at the spatial clustering of regions with high and low values measured by the Moran's I, we see that in 1880 the elite actually exhibited higher values than the farmers and the workers and others. This finding might again be related to the fact that the elite had already entered the fertility transition by 1880, as this process tends to create clustered areas with low fertility around early centres of the decline (Coale and Watkins 1986). The elite did, however, experience the strongest decreases in the spatial autocorrelation of the regional CWRs over time, while the workers and others registered the strongest increases. Particularly interesting are the outcomes for the changes over time, as the elite exhibited very low levels of spatial clustering. This finding is in line with our assumption that the dense communication links through space contributed to a more spatially homogenous decline pattern among this group.

To investigate in greater detail the spatial decline patterns by social class, we included Figs. 13.1 and 13.2. Figure 13.1 contains maps that present the spatial variation in the development of the CWRs between 1890 and 1900.⁵ Figure 13.2

⁵We decided to focus the map on the second decade, as in this decade all three social classes experienced at least some decline. We thus consider this second decade to be more informative in terms of the spatial decline patterns.

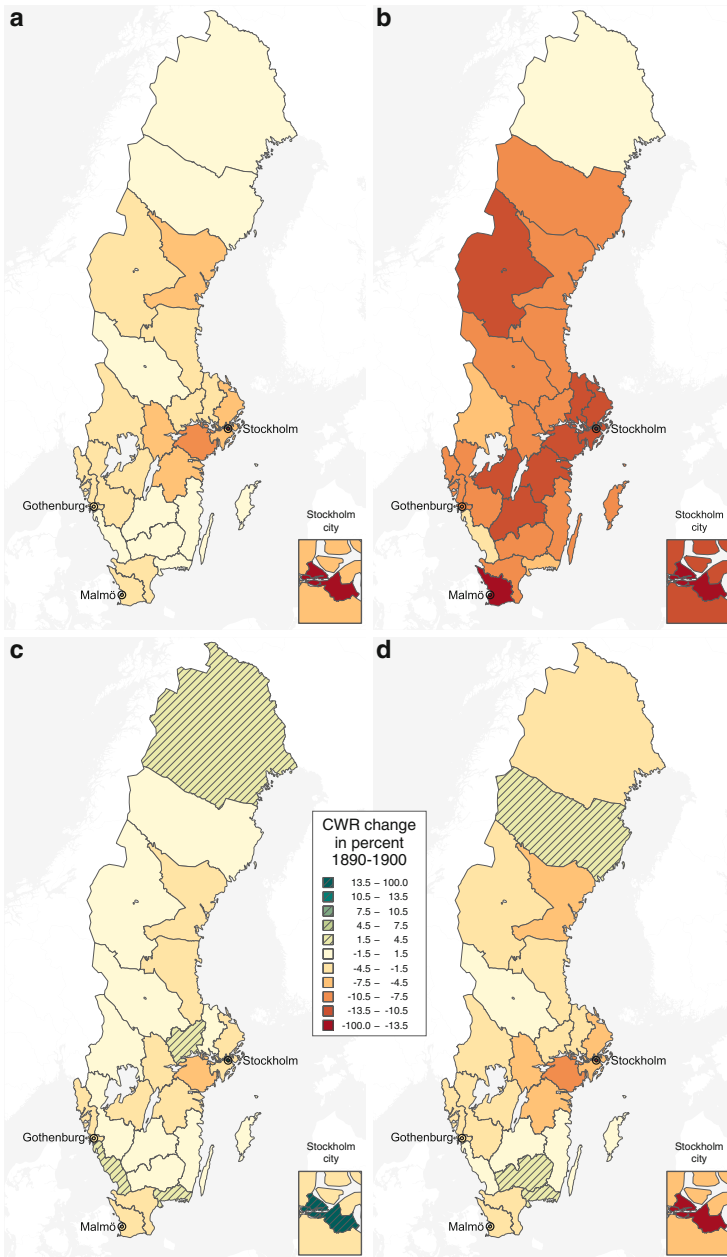


Fig. 13.1 Spatial variation of trends in the child-woman ratio (*CWR*) by socio-economic status (1890–1900, in percent) (a) Total (b) Elite (c) Farmers (d) Workers & others

Note: We only consider married women aged 20–49 born in Sweden with a spouse present in the household in the respective census (1880, 1890, 1900), for whom the social status could be detected

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations.
Base Map: The Swedish National Archives, MPIDR Population History GIS Collection

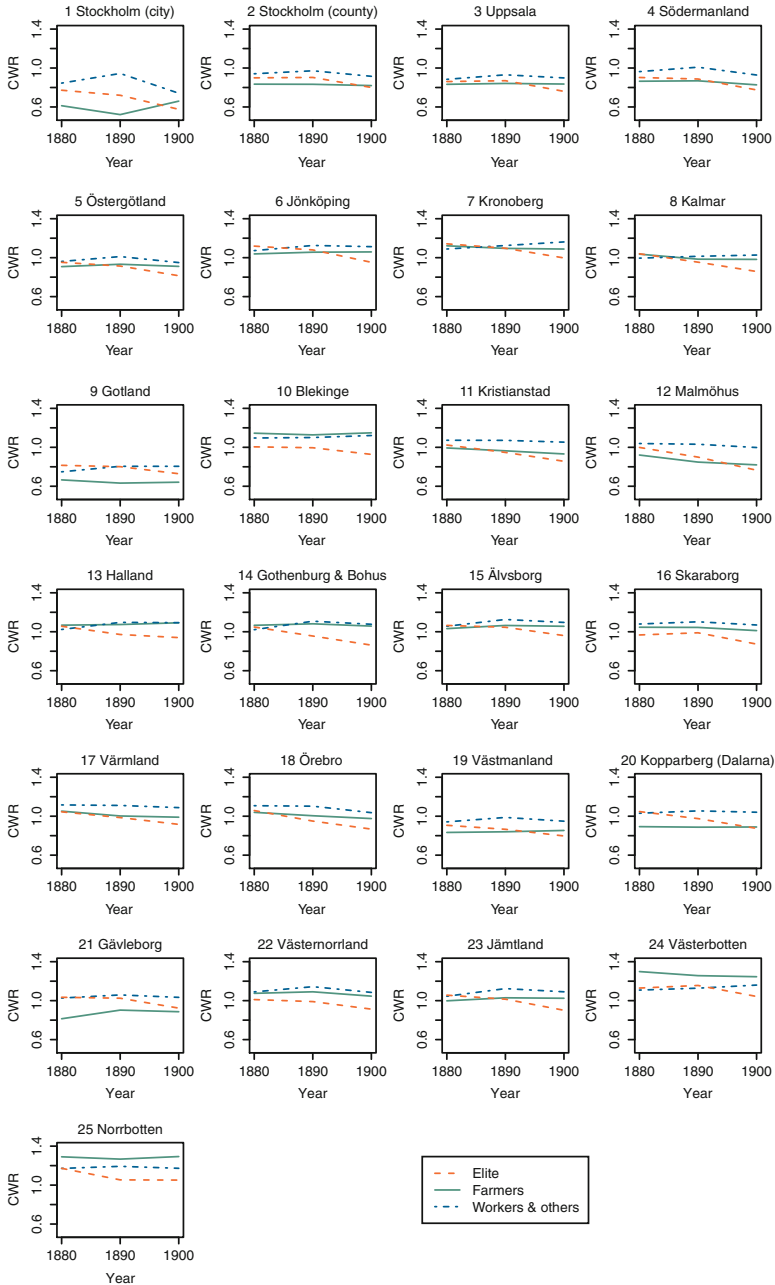


Fig. 13.2 Regional trends in the child-woman ratio (*CWR*) by socio-economic status (1880–1900)
 Note: We only consider married women aged 20–49 born in Sweden with a spouse present in the household in the respective census (1880, 1890, 1900), and for whom the social status could be detected

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

shows the development of the CWR levels in the 25 regions which we obtained from the censuses of 1880, 1890, and 1900.⁶ When we look at the map that displays the trends for all women (Fig. 13.1a), it is visible that the decline was particularly pronounced in the capital city of Stockholm and in the areas located to the southwest of Stockholm. In addition, an area in the northern part of Sweden (around Västernorrland) experienced a substantial decline. This observation does not fit with our assumption that the decline initially occurred mainly in large cities and the surrounding areas. However, the two regions in which Gothenburg and Malmö (the second and third biggest Swedish cities after Stockholm) are located also had relatively high levels of decline. Rather low levels of decline or even slight increases were recorded in central southern Sweden and in the peripheral north. Parts of central southern Sweden were characterised by elevated levels of religiosity (Lindström 2001), and have been referred to as the Swedish “Bible Belt”. As our models do not account for variation in religiosity, we might face problems in capturing how these aspects contributed to shape the decline patterns.

The spatial fertility change pattern of the elite (Fig. 13.1b) differed substantially from that of the total population. During this period, fertility declined in this social class to some extent in all of the regions except the very north of Sweden, and in many regions the CWRs decreased more than 10%. Overall, the pattern for the elite suggests that by the late nineteenth century information about the advantages of adopting fertility control strategies, and the techniques used to prevent conception, had already spread to virtually all parts of Sweden. The farmers (Fig. 13.1c) and the workers and others (Fig. 13.1d) experienced much smaller CWR decreases, and most of the regions with relatively large declines were in an area that comprised Stockholm area and counties southwest of the capital. The observation that fertility increased among the farmers in Stockholm city should not be given too much weight, as this group was very small (93 women in 1880). The regions in which Gothenburg and Malmö were situated registered overall declines, but in these regions the declines among the farmers and the workers and others were not particularly large. Overall, the findings displayed in the maps in Fig. 13.1 support the view that the spatial patterns varied substantially by social class.

These differences may become even clearer if we look at the trend patterns in the 25 regions in Fig. 13.2. In many regions, the CWRs among the farmers and the workers and others were actually increasing between 1880 and 1890, while the CWRs among the elite were already on a downward trajectory. This suggests that social class boundaries played an important role in shaping the fertility transition patterns, either because different classes were subject to varying levels of adjustment incentives and pressures, or because relatively little information with relevance to the decision to adopt fertility control strategies was diffusing across social class boundaries. It is important to note that Swedish society of the late nineteenth century was highly stratified: i.e., the nobility and high-level managers and professionals

⁶The location of the regions displayed in Fig. 13.2 can be obtained from an overview map in Appendix 1.

formed a distinctive elite, and opportunities to enter this elite group were limited (see, e.g., Dribe et al. 2015). Overall, the outcomes of the analysis of the observed patterns indicate that the adoption of deliberate fertility control strategies occurred among the elite earlier and more homogeneously across regions than it did in the other two social classes. Some of the more conservative regions in central southern Sweden and the peripheral north were lagging behind.

13.5 Conceptual Considerations for the Agent-Based Simulations

The main aim of our agent-based simulations is to explore which major social class and spatiotemporal characteristics of the fertility decline in Sweden may have resulted from communication processes structured by social and spatial variation in communication links. A limitation we face in analysing the observed patterns is that the available data do not cover the whole transition period, but only the decades during which the process gained momentum. Based on our observations of this period, and on the existing literature on the social and spatiotemporal characteristics of the decline (e.g., Livi-Bacci 1986; Szreter 1996; Schmertmann et al. 2010; González-Bailón and Murphy 2013; Goldstein and Klüsener 2014), we have identified the following patterns as characteristic of the fertility transition:

- Members of the elite are forerunners in the process, particularly at the beginning of the transition (Dribe and Scalone 2014; see also Livi-Bacci 1986; Szreter 1996), while the farmers are laggards (Dribe and Scalone 2014; see also van Poppel 1985).
- The elite experiences the decline more homogeneously across space than farmers and workers and others. A potential mechanism for this might be that the elite maintains long-distance communication links to a greater extent than other social classes (Szreter 1996; Klüsener et al. 2013). Particularly in more peripheral regions, the elite may also benefit from their smaller group size and their concentration in local urban centres.
- Big cities are early centres of the fertility transition. In addition, there may emerge diffusion clusters around early centres of the decline (see also Schmertmann et al. 2010; Goldstein and Klüsener 2014), while peripheral areas lag behind (Klüsener et al. 2013).

For our models, we assume that both social class boundaries and distances between regions act as moderators in shaping the fertility transition in space and time. Support for this assumption is provided by the findings presented in Fig. 13.2, which show that different classes in different regions experienced the onset of the decline at varying points in time, with the elite being a forerunner in the process. We certainly cannot rule out the possibility that the patterns we see in Fig. 13.2 could have been caused by social and spatial variation in adjustment

incentives and pressures. But if we assume that these observed patterns were indeed caused by diffusion processes, then these diffusion processes would need to have spread faster within social classes in particular regions than across social class and regional boundaries to have produced such patterns. The assumption that social class boundaries and distances influenced the decline also fits with Szreter's (1996) theoretical considerations on communication communities.

The simulations presented in this contribution have been programmed in R 3.2.1 solely based on the functions available in the pre-installed packages. The data and the code are available for download on the following website: www.openabm.org/model/5028. The full individual-level sample of women of childbearing age in Sweden obtained from the census of 1880 that we use for our models has been described in the data section. For our models, we consider our three social classes and the 25 Swedish regions as the region of residence or of birth. In order to isolate the potential role of communication processes in shaping the social and the spatiotemporal variation in the decline, our models are based on a *ceteris paribus* approach: we model the fertility decline as a communication process within a static society that is not subject to adjustment pressures⁷ or any other social change processes apart from the fertility transition itself. The individuals in our models also do not age or migrate.

In deriving information on existing communication links, we follow Hägerstrand (1965) and Rosero-Bixby and Casterline (1994) in using migration links between regions as proxies. To do this, we take data on the so-called "life-time net migration links", measured by the region of birth and the region of residence of the women in our sample. In our models, a specific woman can be influenced by agents residing and/or by processes occurring in her region of residence, her region of birth, or in regions to which agents who were born in her current region of residence migrated. Our decision to also assign relevance to the region of birth is based on the assumption that many women who left their birth region still had communication links to friends and family members living in that birth region.

In Fig. 13.3 we present the life-time net migration links for the three social classes in 1880. The displayed matrixes show the share of specific combinations of regions of birth and regions of residence in the different social classes (e.g., the share of elite women in Sweden who were born in Stockholm region and were living in Uppsala region in 1880). Figure 13.3 demonstrates that a large share of the elite

⁷The unaccounted adjustment incentives and pressures are likely to vary across space and time. They include not only variation in infant mortality and socio-economic factors that directly affect the costs of having children, but also variation in social norms that, for example, condemn the use of contraceptive techniques. The latter might create indirect costs, as in areas in which such social norms are widespread individuals who adopt a fertility control behaviour might face a loss of social capital, which Bourdieu and Wacquant (1992, p. 119) define as the resources that "accrue to an individual or a group by virtue of possessing a durable network of more or less institutionalized relationships of mutual acquaintance and recognition." Social capital losses might also have repercussions in terms of access to income opportunities.

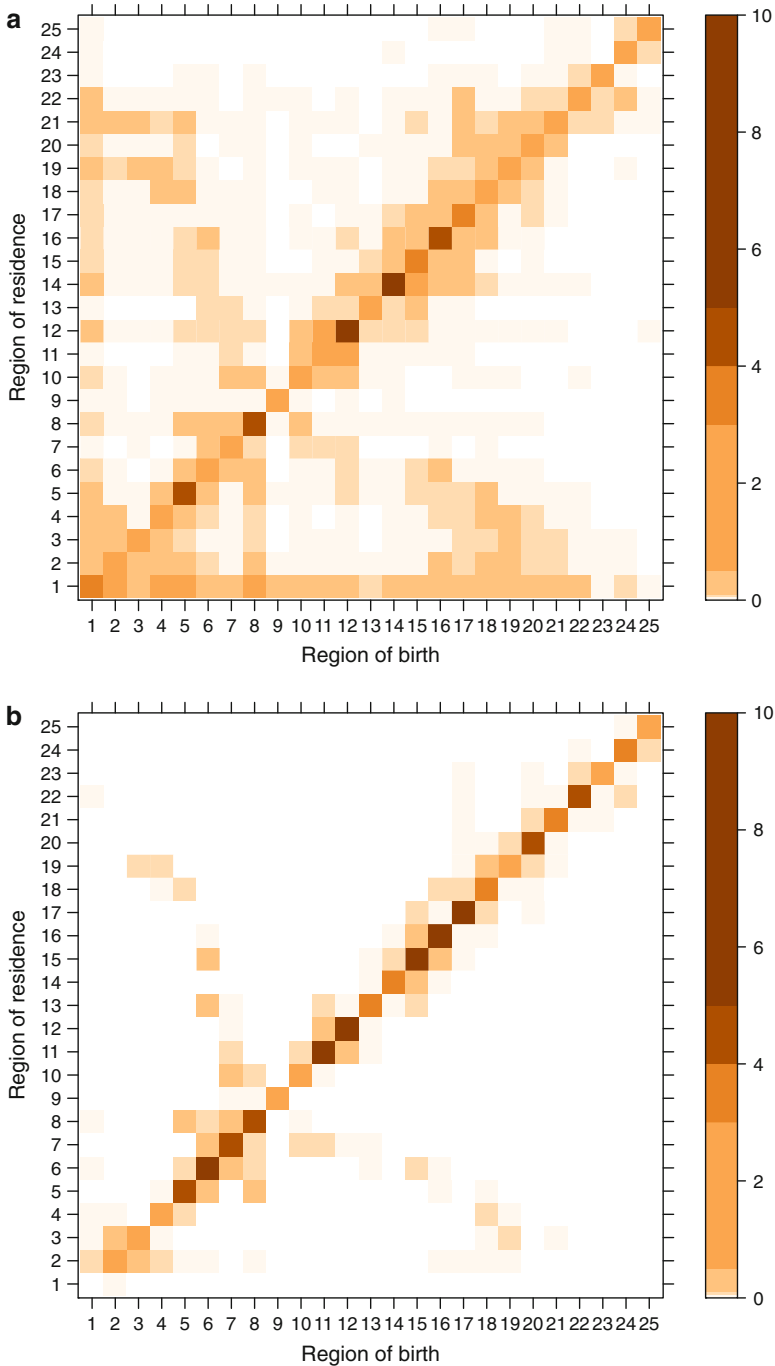


Fig. 13.3 Life-time net migration links of married women aged 20–49 by socioeconomic status (1880) (a) Elite (b) Farmers (c) Workers & others

Note: The graphs show for each social status group the percentage of women with a specific region-of-birth and region-of-residence combination in the total number of women of that group in Sweden. The city of Stockholm is region 1, the second biggest city, Gothenburg, is part of region 14,




Fig. 13.3 (continued) and the third-biggest city, Malmö, is part of region 12. The complete list of regions is given below. We only consider married women aged 20–49 who were born in Sweden with a spouse present in the household in the census of 1880, and for whom the social status could be detected

1. Stockholm city; 2. Stockholm county; 3. Uppsala county; 4. Södermanland county; 5. Östergötland county; 6. Jönköping county; 7. Kronoberg county; 8. Kalmar county; 9. Gotland county; 10. Blekinge county; 11. Kristianstad county; 12. Malmöhus county; 13. Halland county; 14. Gothenburg and Bohus county; 15. Älvsborg county; 16. Skaraborg county; 17. Värmland county; 18. Örebro county; 19. Västmanland county; 20. Kopparberg county; 21. Gävleborg county; 22. Västernorrland county; 23. Jämtland county; 24. Västerbotten county; 25. Norrbotten county
Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

women aged 20–49 were no longer living in their region of birth in 1880, while most of the farmers were living in or close to their region of birth. As we take this information as a proxy for communication links, this implies that in our models elite women are more likely to spread information across regional boundaries.

The plots for the elite and the workers and others are more similar than the plots for the elite and the farmers. However, the differences in the ways in which the elite and the workers and others were connected to the capital city of Stockholm are particularly striking. In both groups, a considerable share of the women who were born outside of Stockholm city had moved to the capital. Among the elite, however, we also find a substantial share of women who had been born in Stockholm city, but who were living in one of the other regions of Sweden in 1880 (first vertical line on the left of the elite graph in Fig. 13.3). This pattern appears to be attributable

to national development processes. For example, a Stockholm-born woman may have moved to a more remote area of Sweden because her husband was serving in the local elite as a civil administrator, a military officer, a doctor, a teacher, or a priest. It is likely that these women were still attempting to follow new developments in their hometown of Stockholm, and thus acted as a vanguard group who spread information about new lifestyles from the capital to other parts of Sweden. The farmers, on the other hand, were the least connected to the city of Stockholm. Along with differences in group sizes, these differences in migration/communication links by social class and region are the main mechanisms that are expected to cause social and spatial variation in our simulations of the fertility decline.

In the simulations, we do not attempt to model the full fertility histories of the observed women, as the censuses provide us with information only about the number of surviving children. Therefore, estimates of the “observed” fertility rates of women of different social classes across regions would need to be based on very bold assumptions. We decided instead to focus on the process of the adoption of a parity-specific fertility control strategy. Related to this we assume that prior to the fertility transition women generally believed that they had little influence over their total number of surviving children. The adoption of the parity-specific fertility control strategy therefore implies that women are starting to plan to have a target number of surviving children. This target number may be subject to change over the woman’s life course, but it is usually much lower than the number of children she would have if she has not adopted a fertility control strategy. We believe that a prerequisite for this shift in strategies is a substantial reduction in infant mortality, as having a small number of children would be a risky strategy if infant mortality levels were still high due to a high disease burden. Since this structural condition had been met in Sweden by 1880, it was no longer a potential obstacle to the fertility transition (see Dyson 2011). We believe that once a woman starts to consider having a target number of children, she is unlikely to abandon this way of thinking. We therefore follow González-Bailón and Murphy (2013) in their conceptualisation of an agent-based model of the fertility transition in France by modelling the adoption of a parity-specific fertility control strategy as an irreversible process.

In order to start our simulations, we need to have at least one agent in the population who has already adopted the new behaviour. We consider three different starting scenarios that comply with this starting condition: the *random start* scenario, the *diffusion from early decline countries* scenario, and the *diffusion from big cities* scenario. In addition, we define communication algorithms that are used to let the adoption of the new fertility behaviour diffuse across Swedish society. Here, we again consider three mechanisms: *social adaptation*, *social influence*, and *social learning* (see also Kohler 2001).⁸

⁸These mechanisms may have been operating simultaneously, but we will keep them separate in our model specifications.

13.5.1 *Three Different Starting Scenarios*

In the *random start* scenario, we randomly choose one woman somewhere in Sweden as the first adopter, and let the process diffuse from this woman. We repeat this simulation several times with a different randomly chosen “patient number one”. In choosing these first adopters, we use a stratified sampling that ensures that the social classes of the 100 different first adopters used in 100 iterations are representative of the respective sizes of these social classes in Swedish society in 1880. This type of scenario allows us to explore the question of whether the existing social and spatial variation in communication links is sufficiently deterministic that we would be able to reproduce major characteristics of the fertility transition patterns using our simple communication rules, even if the process started out of complete randomness.

In the second, *diffusion from early decline countries*, scenario we assume that information relevant to the adoption of a parity-specific fertility control strategy arrived in Sweden through communication links with other countries where this behaviour was already widespread. Around 1880 only two European countries had experienced substantial fertility declines as part of the demographic transition: namely, France and the French-speaking part of Belgium. As a proxy for regional and social variation in communication links with these countries, we add to our sample of 488,438 women born in Sweden all 32 women born in France and Belgium. These are mostly elite individuals who were living in one of the bigger Swedish cities (see Klüsener et al. 2016). We then assume that these 32 individuals were the first adopters, and let the fertility decline diffuse from them into Swedish society.

In the third, *diffusion from big cities*, scenario we believe that big urban centres are the early centres of the decline, as the populations are, relative to rural populations, either subject to greater economic incentives and pressures, or more open to new ideas. In simulations based on this scenario, the share of adopters in the big urban centres is set to a certain threshold at the beginning of the process (see details below). This scenario is the only one that refers to variation in adaptation incentives and pressures in the definition of the starting condition.⁹

⁹Initially, we also considered obtaining as additional starting scenario estimates of the share of individuals who had already adopted fertility control strategies from the CWR levels by social class and region of residence in 1880. However, Fig. 13.2 shows that there is substantial variation in the levels at which specific social classes in specific areas experienced the onset of the decline. We therefore believe that it would be a very strong assumption to claim that variation in CWR levels by social class and region at a single point in time would provide information about the share of individuals who had already adopted the new behaviour.

13.5.2 Three Different Communication Algorithms

The first diffusion algorithm we refer to as *social adaptation*. In this variant of the simulation model, in each time period all individual women (agents) who have not yet switched to a deliberate parity-specific fertility control strategy are at risk of adopting this new behaviour. The adoption risk varies across women, and is determined by the share of adopters in the “social surrounding”. The individual risk depends in part on the share of women of the same social class in the woman’s region of residence. This is based on the assumption that the risk that a woman will adopt the behaviour in a specific time period increases as the share of women of the same social class around her who have already adopted the behaviour increases. Thus, the risk of adopting the behaviour will be very low at the beginning, when there are hardly any women around who have already adopted it, but will rise as the behaviour is adopted by an increasing share of women. In addition to the share of adopters in the woman’s region of residence, we also believe that the share of adopters of the same social class in a woman’s region of birth has an effect on her risk of adoption. This is based on the assumption that women are likely to retain communication links to their region of birth through friends and family members. The inclusion of trends in the region of birth that can affect the risk of adoption allows us to let the behavioural change spread across regional boundaries.¹⁰

In addition, we also allow the behaviour to spread across social boundaries. To do this, we let the risk of adoption in the region of residence depend not only on the share of adopters in the same social group, but also to some degree on the share of adopters in the vanguard social group in this region with the highest share of adopters at a specific time.¹¹ In notation form, the individual risk of adoption RA by a woman of a specific social class s in a specific time period t is determined as follows:

$$RA_{s,t} = \frac{(SAR_{vs,t} * ws + SAR_{s,t} * (1 - ws)) * wr + SAB_{s,t} * (1 - wr)}{100} * x \quad (13.1)$$

in which SAR and SAB denote the share of women who adopted a fertility control behaviour (in percent) in the region of residence or, respectively, the region of birth at the beginning of that time period t , and vs the vanguard social class with the highest share of adopters at that time. Next, we have two weighting parameters, wr and ws , that allow us to specify to what degree diffusion is occurring across regional and social boundaries. For each of them we can chose values between 0 and 1. The parameter wr determines the weight that is given to the share of adopters

¹⁰For the women who were in 1880 still living in their region of birth, the share of adopters in the region of birth and the region of residence are identical.

¹¹We allow trickle-down effects from the vanguard group only in the region of residence, as we consider it rather unlikely that women would copy behaviour from the vanguard group in their region of birth if they were no longer living in that region.

in the region of residence in contrast to the region of birth. The higher w_r , the more weight is given to the share of adopters in the region of residence than to the share of adopters in the region of birth. The parameter w_s regulates the weight given in the region of residence to the share of adopters in the vanguard group in contrast to the share of adopters in the woman's own social class. The higher w_s , the more weight is given to the vanguard group. The last parameter x determines the maximum risk of adoption that will be approached when almost all of the women have adopted the behaviour. In order to introduce stochasticity in the model, we draw in each period for each woman who has not yet adopted parity-specific fertility control strategies a random number between 0 and 100 (all of the random numbers in our models are drawn from a uniform distribution). If this randomly drawn number is below the determined risk of adoption, the woman adopts the behaviour in this time period and will from then on be considered to be among the share of women who have adopted the new strategy.

The second algorithm we refer to as *social influence*. In this model variant, women who have already adopted a fertility control behaviour act as agents of change. For each time period for each woman y who has already adopted a fertility control behaviour, three random numbers between 0 and 100 are drawn. The first determines, based on a threshold $h.i_1$ (e.g., if a number below five is drawn) a social interaction (e.g., through conversation) with a randomly chosen other woman z of the same social class in the region of residence. In this social interaction, woman z is persuaded to adopt fertility control strategies, if she has not done so already. The second random number drawn for woman y regulates, based on a threshold $h.i_2$, a similar interaction with a randomly chosen woman z of the same social class in the region of birth. The third random number is considered only if woman y belongs to the vanguard group with the largest share of adopters. In that case, this third random number determines, based on a threshold $h.i_3$, the risk of a similar interaction with a woman z of any other social class in the region of residence. The advantage of this model variant is that it allows for interactions in which, for example, a woman who has moved from a village to a big urban centre communicates new ideas to social contacts who have remained in her home region (e.g., sisters, friends). These interactions could not be captured by our *social adaptation* specification, in which a woman who has remained in her birth region can only be influenced by developments in her home region.¹²

In the third algorithm, which we refer to as *social learning*, women who have not adopted fertility control strategies are copying behaviour from forerunners. For each time period for each woman y who has not yet adopted a fertility control behaviour, three random numbers between 0 and 100 are drawn. The first determines, based on a threshold $h.l_1$, a social interaction with a randomly chosen other woman z

¹²A reviewer has pointed out that we could implement backward influences in the *social adaptation* algorithm as well. We agree, but doing so would greatly increase the complexity of that model, as we would need to account for the share of women of a specific social class who had moved to other regions. We thus decided to stick with this simpler specification of our *social adaptation* algorithm, which requires just one equation.

of the same social class in the region of residence. If this other woman z has already adopted fertility control strategies, the interacting woman y adopts these strategies. The second regulates, based on a threshold $h.l_2$, a similar interaction with a randomly chosen woman z of the same social class in the region of birth, while the third determines, based on a threshold $h.l_3$, a similar interaction with a woman y of the vanguard group living in the region of residence. This copying of behaviour might have occurred more frequently than incidences of one woman persuading another woman to adopt the behaviour, as we assumed in the *social influence* algorithm. However, a disadvantage of this *social learning* algorithm is that a woman who has remained in her region of birth can—as in the *social adaptation* algorithm—only be influenced by processes that are occurring in her home region.

13.6 Results

In presenting the results of our simulation models, we will focus on the outcomes of our *diffusion from big cities* scenario and the *social adaptation* algorithm. The starting condition for this scenario requires us to refer to variation in adjustment pressures, as we assume that the first adopters were living in the cities where the adaptation incentives and pressures were probably the highest, and the inhabitants were the most open to new ideas. The outcomes for the *diffusion from early decline countries* and *random start* scenarios are presented in a second publication (Klüsener et al. 2016). We let each simulation run over 200 time periods.¹³

The assumption that the fertility decline started in big cities is realistic, as the city of Stockholm in particular was a forerunner in the decline (see Fig. 13.2). For this simulation model, we set as a starting condition the share of adopters in t_0 in Stockholm city in all three social classes to 2%. These first adopters were randomly sampled without taking other attributes into account.¹⁴ To assume that the elite, the farmers, and the workers and others had the same share of adopters in t_0 is rather unrealistic, as we know that the elite group experienced the onset of the decline in Stockholm city much earlier than the other two social groups (see Fig. 13.2).¹⁵ However, the same share of adopters was chosen for all of the social classes, as otherwise a faster adoption rate among the elite might have resulted from

¹³The time periods in the model outcomes presented in this contribution relate roughly to years, as most of the decline occurred within 50 time periods, which correspond with the 50 years between 1880 and 1930 in which Sweden experienced most of the fertility transition (Dyson 2011).

¹⁴As an alternative, we could have attempted to determine the likelihood that a woman adopted fertility control strategies based on her recent fertility history; i.e., by the number of children linked to that woman by the mother locator.

¹⁵The farmers in Stockholm city had the lowest levels throughout the period. This might be related to their small numbers, and the possibility that farmers who were living in the capital formed a very selective group.

the choice of a higher share of adopters in this group in the big cities in t_0 . Next to Stockholm city, the share of adopters in the counties of Gothenburg, Malmö, and Stockholm was set in all social classes to 1%. Stockholm county was also considered, as only a small group of farmers were living in Stockholm city. Thus, to ensure that for the farmers the process was spreading out of the Stockholm area with sufficient speed, we decided to include Stockholm county as an origin region of the decline as well. In total, we ran for each combination of parameter specifications (w_r , w_s , x) considered 100 iterations of our model. In order to minimise the impact of the sampling procedure with which we define the adopters in t_0 , we derived 100 different samples for these 100 iterations. These 100 different samples were kept the same for each combination of parameters considered.

The outcomes of the models are presented in Figs. 13.4, 13.5, 13.6, 13.7, 13.8, and 13.9. Figures 13.4, 13.5, 13.6, and 13.8 show the results of the model for different specifications of w_r and w_s and a maximum adoption risk x of 10.¹⁶ We believe that the share of women who have already adopted the behaviour in the region of residence has a greater influence on the adoption risk than the share of adopters in the region of birth. Thus, we considered w_r specifications of 0.5, 0.7, and 0.9. For w_s the decline patterns presented in Fig. 13.2 suggest that trickle-down effects were not very dominant, at least at the beginning of the process. We therefore chose w_s settings of 0.05, 0.1, 0.15, and 0.2.

Figure 13.4 displays the development of the national mean values by social class over time, Fig. 13.5 shows the trends in the standard deviation of the share of adopters in the 25 regions, and Fig. 13.6 provides the development of the Moran's I that measures the spatial clustering in the observed spatial decline patterns. In Fig. 13.7 we present maps of the simulated patterns, while Fig. 13.8 displays trends in Pearson's product moment correlation coefficients that compare the simulated and the observed regional patterns (by social class and for all regional values of the three social classes together). Figure 13.9 presents sensitivity checks for one specification of w_r and w_s that we consider to be quite plausible based on the observed patterns ($w_r = 0.9$; $w_s = 0.05$). The development of the mean values indicates that independent of the considered w_r and w_s specifications, the order in which the social classes experience the decline is always the same: i.e., the elite experiences the decline first, followed by the workers and others, while the farmers lag behind. When we increase our parameter w_s , which allows for trickle-down effects from the vanguard group, the temporal advantage of the elite decreases, particularly in the second half of the transition. If we reduce w_r and thus give a greater weight to developments in the birth region affecting the risk of adoption, the overall transition tends to occur in a shorter period of time, especially in the second half of the transition, when the process is spreading rapidly to the other regions of Sweden.

¹⁶We also ran all of the combinations of w_r and w_s presented here with $x = 15$ and $x = 20$, but the outcomes suggest that an increase of x predominantly just affects the speed of the process. Thus, we focus in the presentation of the results on the outcomes obtained with $x = 10$.

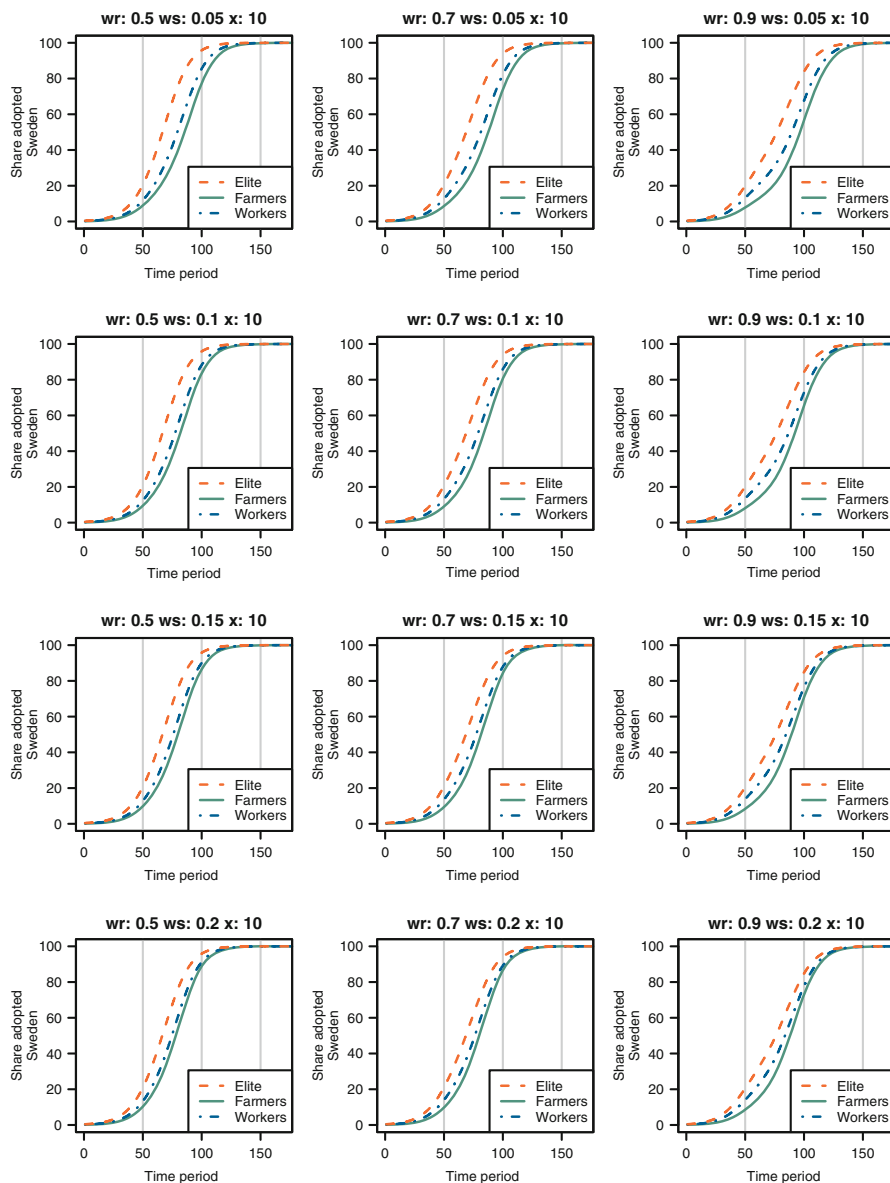


Fig. 13.4 Simulation outcomes – share of women who adopted a fertility control behaviour by socio-economic status

Scenario: Diffusion from big cities

Communication algorithm: Social adaptation

Note: Average of 100 simulations. The parameter w_r denotes the weight that is given to trends in the region of residence (w_r) or the region of birth ($1-w_r$). The parameter w_s indicates the weight within the region of residence that is given to trends in the vanguard social class (w_s) relative to the woman's own social class ($1-w_s$). The parameter x denotes the maximum adaptation risk in a specific time period, which is approached when almost all of the women have adopted the new behaviour

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

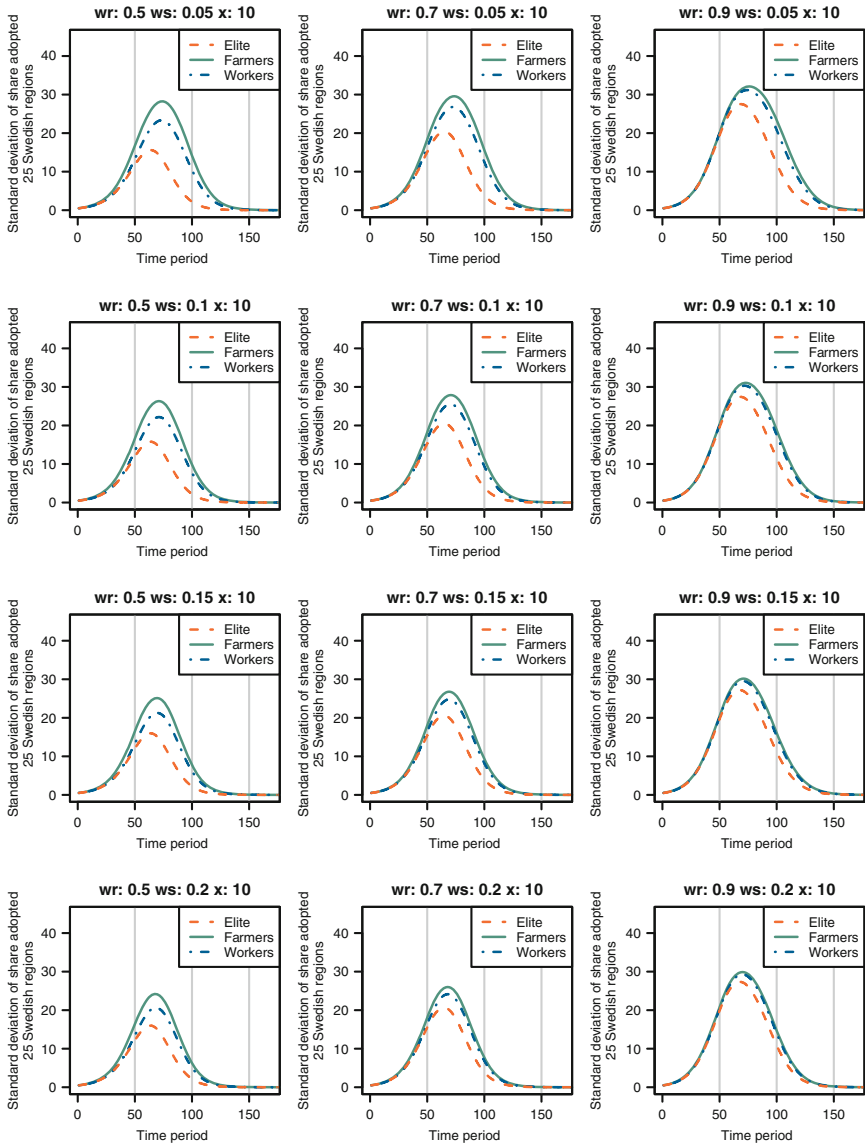


Fig. 13.5 Simulation outcomes – standard deviation in share of women who adopted a fertility control behaviour across the 25 Swedish regions by socio-economic status

Scenario: Diffusion from big cities

Communication algorithm: Social adaptation

Note: Average of 100 simulations. The parameter wr denotes the weight that is given to trends in the region of residence (wr) or the region of birth ($1-wr$). The parameter ws indicates the weight within the region of residence that is given to the trends in the vanguard social class (ws) relative to the woman's own social class ($1-ws$). The parameter x denotes the maximum adaptation risk in a specific time period, which is approached when almost all of the women have adopted the new behaviour

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

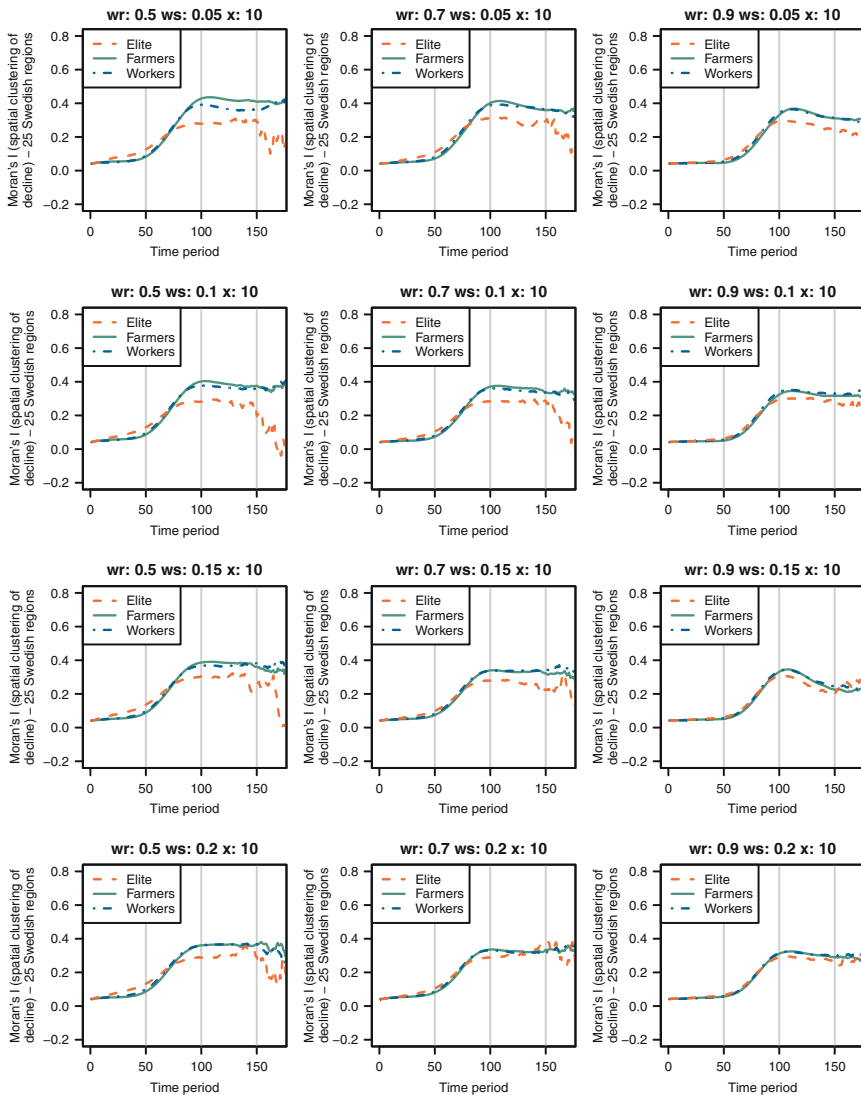


Fig. 13.6 Simulation outcomes – Moran’s I test of spatial autocorrelation (i.e., spatial clustering) in share of women who adopted a fertility control behaviour across 25 Swedish regions by socio-economic status

Scenario: Diffusion from big cities

Communication algorithm: Social adaptation

Note: Average of 100 simulations. The parameter w_r denotes the weight that is given to trends in the region of residence (w_r) or the region of birth ($1-w_r$). The parameter w_s indicates the weight within the region of residence that is given to the trends in the vanguard social class (w_s) relative to the woman’s own social class ($1-w_s$). The parameter x denotes the maximum adaptation risk in a specific time period, which is approached when almost all of the women have adopted the new behaviour. In deriving the Moran’s I we define the four nearest regions as neighbours. These neighbouring regions are identified by calculating the spherical distance between the geographical centroids of the 25 Swedish regions. Please note that the Moran’s I becomes very sensitive to small differences when all of the regional values converge to a value other than zero. This can cause some fluctuations when the regional values converge at 100 %

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

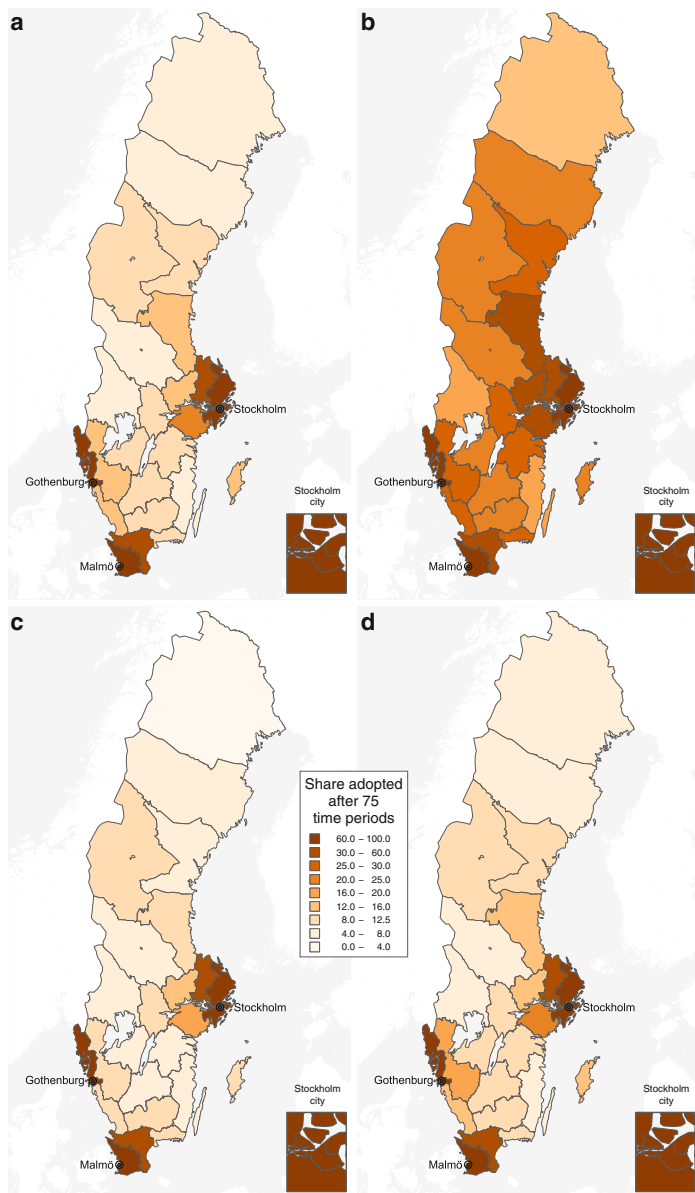


Fig. 13.7 Simulation outcomes – spatial patterns in share of women who adopted a fertility control behaviour after 75 time periods (w_r : 0.9 w_s : 0.05 x : 10) (a) Total (b) Elite (c) Farmers (d) Workers & others

Note: The maps represent averages of 100 simulations. The parameter w_r denotes the weight that is given to trends in the region of residence (w_r) or the region of birth ($1-w_r$). The parameter w_s indicates the weight within the region of residence given to the trends in the vanguard social class (w_s) relative to the woman's own social class ($1-w_s$). The parameter x denotes the maximum adaptation risk in a specific time period, which is approached when almost all of the women have adopted the new behaviour

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

Base Map: The Swedish National Archives, MPIDR Population History GIS Collection

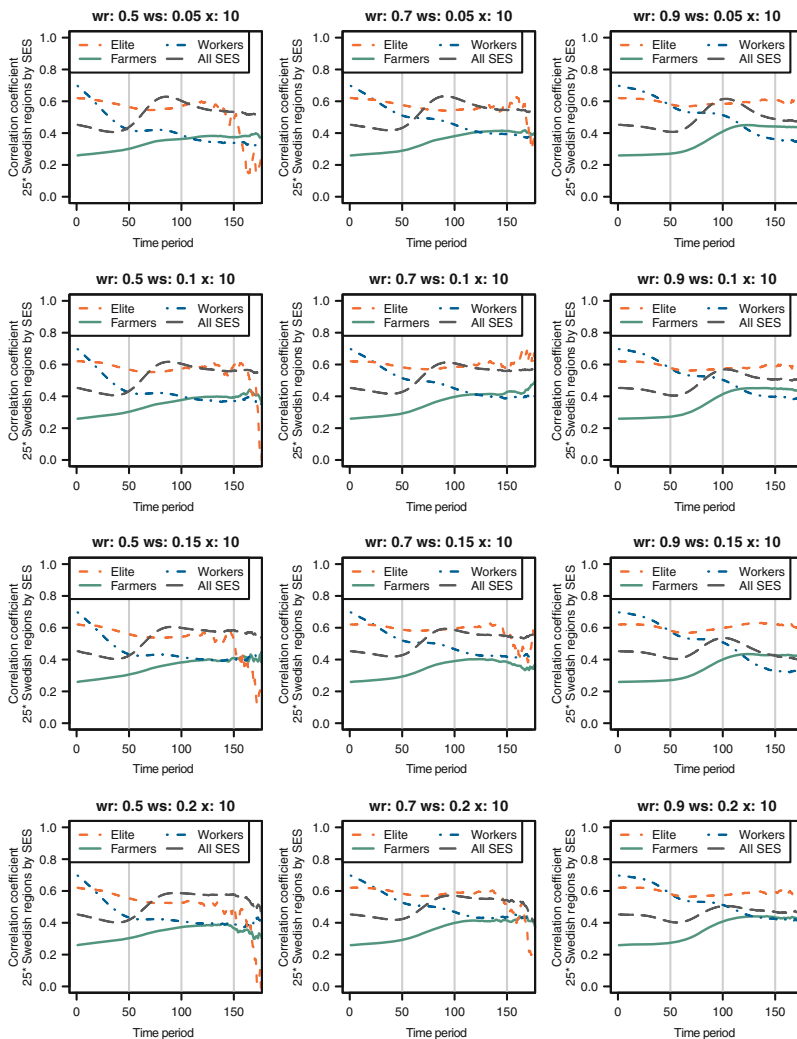


Fig. 13.8 Simulation outcomes – correlation between the simulated share adopted and observed fertility changes in the period 1890–1900 (in percent) across 25 Swedish regions by socio-economic status

Scenario: Diffusion from big cities

Communication algorithm: Social adaptation

Note: Average of 100 simulations. The parameter w_r denotes the weight that is given to trends in the region of residence (w_r) or the region of birth ($1-w_r$). The parameter w_s indicates the weight within the region of residence that is given to the trends in the vanguard social class (w_s) relative to the woman's own social class ($1-w_s$). The parameter x denotes the maximum adaptation risk in a specific time period, which is approached when almost all of the women have adopted the new behaviour. All SES refers to the correlation between the values for all groups (three social groups in 25 regions). In calculating the correlation statistics, we omitted farmers in Stockholm city due to their small group size and their highly specific fertility development (for that reason we mark "25* Swedish regions" in the y-axis title with an asterisk)

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

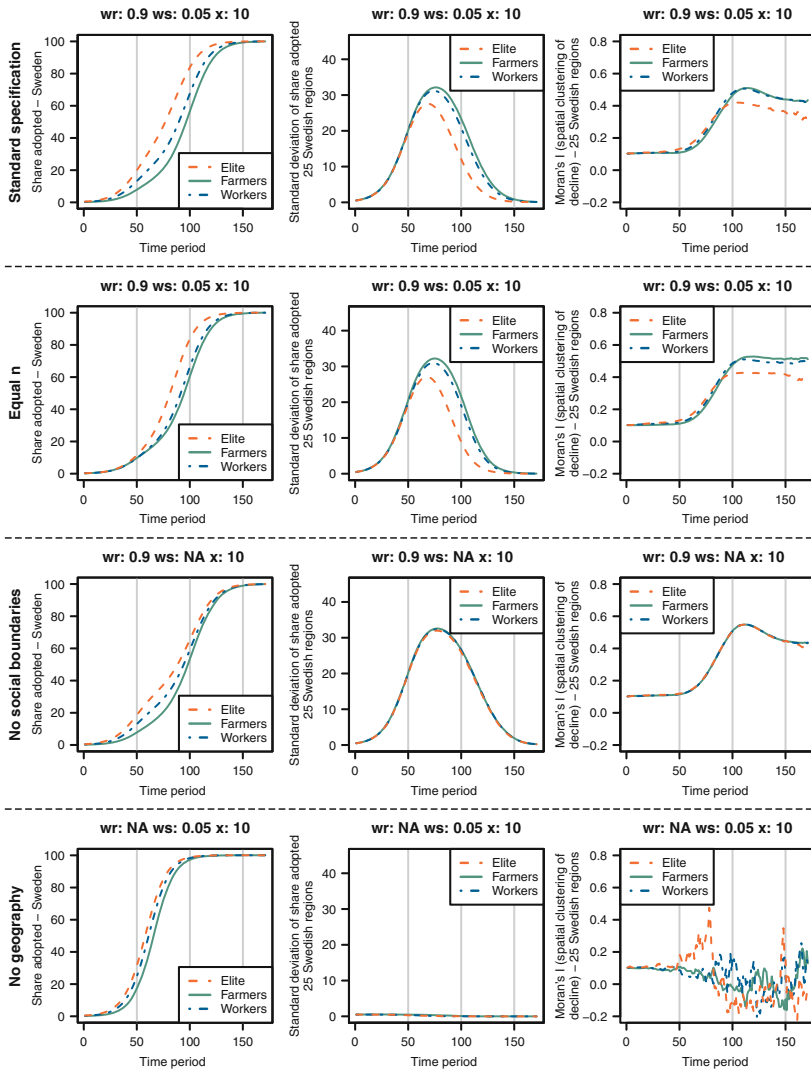


Fig. 13.9 Simulation outcomes – behaviour of the model under extreme conditions

Scenario: Diffusion from big cities

Communication algorithm: Social adaptation

Note: Average of 100 simulations. The parameter w_r denotes the weight that is given to trends in the region of residence (w_r) or the region of birth ($1-w_r$). The parameter w_s indicates the weight within the region of residence that is given to the trends in the vanguard social class (w_s) relative to the woman's own social class ($1-w_s$). The parameter x denotes the maximum adaptation risk in a specific time period, which is approached when almost all of the women have adopted the new behaviour. In deriving the Moran's I we define the four nearest regions as neighbours. These neighbouring regions are identified by calculating the spherical distance between the geographical centroids of the 25 Swedish regions

Source: Micro-level census data, SweCens, The Swedish National Archives; own calculations

Base Map: The Swedish National Archives, MPIDR Population History GIS Collection

In terms of the variation across regions, the development of the standard deviation of the regional share of individuals who already adopted the new behaviour displayed in Fig. 13.5 shows that in the initial phase of the decline, all three groups experience a very similar increase. But in all of the w_r and w_s specifications considered, the elite experiences at an earlier point in time a turn-around in this trend at a lower level of regional variance, while the regional variance among the other two groups continues to increase. The workers and others are the second group to register the turn-around, while the farmers are again the laggards. Thus, in our models the elite have the spatially most homogenous pattern of decline during the transition, while the farmers have the most heterogeneous pattern. When we decrease w_s and w_r , the social class differences become more pronounced. In addition, the outcomes for the Moran's I (Fig. 13.6) provide support for the view that the simulated decline pattern for the elite is more spatially homogenous than the simulated patterns for the other two groups in terms of the maximum Moran's I. Nevertheless, in many of the w_r and w_s parameter combinations considered, in the initial phase the elite has a higher Moran's I than the other two groups. However, this finding fits with the observed patterns in which the elite had the highest Moran's I in 1880. In addition, the differences between the Moran's I trends of the workers and others and the farmers are less pronounced compared to the development of the means and the standard deviations.

When we look at the spatial patterns we obtain with the *diffusion from big cities* scenario for the parameter combination $w_r = 0.9$, $w_s = 0.05$ and $x = 10$ at t_{75} (Fig. 13.7), we see that compared to the observed patterns shown in Fig. 13.1, they are a bit more focused on the big cities. Thus, to get a more realistic outcome using our *social adaptation* algorithm, we would either need to set the share of first adopters in big cities lower or reduce w_r in order to have the diffusion become less focused on the big cities. Overall, however there are many similarities between the simulated and the observed patterns. In the simulated map for all social classes (Fig. 13.7a), the decline is similar to the observed pattern focused on Stockholm. In addition, elevated levels of decline can be found in Gothenburg and the lake corridor linking Stockholm and Gothenburg, as well as in two regions in the southwest, where Malmö is located. Furthermore, in both the observed and the simulated patterns there is a corridor with elevated levels of decline moving upward from Stockholm to the northwest. In terms of social class differences, the models show that among the elite the decline has already spread to almost all parts of Sweden, except perhaps to the peripheral north; while the workers and others and the farmers are lagging behind. Among these two groups, the simulated decline is also more clustered on the areas around the early centres of the decline. Particularly remarkable is the fact that our models are replicating the elevated fertility decline in Södermanland county southeast of Stockholm among the farmers, the workers and others, and the total population.

The quite good fit of our simulations is also confirmed by the correlation graphs in Fig. 13.8, in which we contrast the simulated patterns with the observed patterns in the period between 1890 and 1900, when all social classes experienced at least some decline. For the elite and the workers and others we obtain for preferred

parameter combinations (e.g., $wr = 0.9$, $ws = 0.05$ and $x = 10$) at the beginning of our simulated transition correlations of around 0.6. For the farmers, however, the correlations are substantially lower, at around 0.25. This finding might be related to the fact that the farmers had not experienced significant fertility declines by 1900, which could affect the extent to which the regional fertility change pattern observed among this group between 1890 and 1900 is reflective of the pattern observed during the fertility transition.

In order to gain a better understanding of which factors play a dominant role in shaping specific aspects of the fertility decline patterns we observed in our model outcomes, we decided to run a number of additional consistency checks in which we moved the model in extreme conditions. In total, we developed three model variants in addition to our standard specification, which are called *equal n*, *no social boundaries*, and *no geography*. In the *equal n* variant we explore the effects that differences in group sizes have on our model outcomes, as our *social adaptation* algorithm lets the behaviour spread faster in smaller groups.¹⁷ To do this, we modified the group sizes in our 75 groups (three social groups in 25 regions) so that each group had the same number of women, and the total number of women in the simulation remained approximately the same.¹⁸ In the *no social boundaries* specification social class boundaries no longer act as moderators of the decline. The risk of adoption of a woman of a specific class s is simply dependent on the total share of adopters in her region of residence and her region of birth. For the *no geography* specification the regional boundaries are no longer a moderator of the decline, and we simply model the diffusion by social class for Sweden as a whole, but still take the trickle-down effect from the vanguard group into account.

The outcomes for the standard specification and the three variants for one parameter specification ($wr = 0.9$, $ws = 0.05$, $x = 10$) are presented in Fig. 13.9. The results for the *equal n* simulation differ particularly in the very early phase of the transition. In the first 50 time periods there are now hardly any differences between the three social classes. But in the middle and the last phase of the transition, the pattern for the *equal n* specification differs little from the pattern obtained based on our standard specification. This suggests that particularly in the early phase of the simulated decline, the elite group benefits from being smaller, while in the second and the last phases of the process the social class differences in the connectedness through space seem to dominate our model outcomes. This view is

¹⁷If, for instance, one individual out of 100 individuals in a social group adopts the new behaviour, this implies that one percent of all of the women has adopted the behaviour. By contrast, the adoption of the behaviour by one out of 10,000 women increases the adoption risks in this social group to a much lower degree. It is relevant to note that we believe that this is an inherent property of the diffusion process, and not just an unintended property of the *social adaptation* algorithm that we specified. Our *social influence* and *social learning* algorithms also have the same property.

¹⁸In order to generate a dataset with a total n that was similar to our standard dataset, we derived the target group n by dividing the total number of women in the simulation by the 75 groups (three social groups in 25 regions). We then sampled or duplicated observations to obtain the targeted group n in each group.

also supported by the outcomes for the *no geography* specification, in which the results are simply driven by variation in group sizes, while the social differences in connectedness through space are not taken into account. In the *no geography* specification, differences between the three groups emerge in the early phase of the transition, but no longer increase in the second and the final phases.

Another interesting aspect is the outcome for the *no social boundaries* specification. Even though social boundaries no longer act as moderators of the decline, for the first half of the transition we observe social class differences that are similar to those in our standard specification. This effect is, however, purely compositional, and is driven by the fact that the elite (and the workers and others) are more concentrated in the big cities than the farmers. Thus, in our standard variant, social class differences in the second phase of the lift-off phase (around time period 50) seem to be largely dominated by these compositional effects.

13.7 Discussion and Conclusion

The outcomes of our analysis provide support for the view that a substantial share of the major characteristics of the spatiotemporal fertility decline patterns by social status in Sweden can be reproduced with our simulations, in which the decline is modelled as diffusion process with simple communication rules, with migration links serving as proxies for social and spatial variation in communication links. These reproduced characteristics include that the elite were forerunners and the farmers were laggards, and that the elite had a more spatially homogenous pattern of decline. In the scenario we presented, we assumed that the process started in the big cities, and that the role of big cities as early centres of decline might be related to greater adaptation incentives and pressures. Inherent in this starting condition is the assumption that the decline would initially be concentrated mainly on big cities, and to a lesser extent on peripheral regions. However, as we will demonstrate in a second publication (Klüsener et al. 2016), we are also able to reproduce in our *diffusion from early decline countries* and *random start* scenarios virtually all of the major spatiotemporal characteristics of the fertility decline. Among these characteristics are that the big cities were early centres of decline, that there were diffusion clusters around these centres, and that peripheral regions were laggards. However, for the *random start* scenario, the outcomes are more volatile, and most of the characteristic patterns are only obtained all of the time if we average the outcomes of several simulations.

As a starting condition for our *diffusion from big cities* scenario we set the share of adopters in the three social classes to the same level. Based on the observed patterns, it would probably be more realistic to start the process with a higher share of adopters among the elite, and a lower share of adopters among the farmers. This would further widen the differences between the three social classes with regard to the time period in which the process gains momentum. It is also questionable whether it is realistic to use the same parameters for all three social classes in terms

of the weight w_r that is given to the share of adopters in the region of residence relative to the share of adopters in the region of birth, and of the weight w_s that is given in the region of residence to the share of adopters in the vanguard social class relative to the share of adopters in the woman's own social class. Based on Szreter's considerations regarding communication communities (Szreter 1996), we would assume that during the Swedish fertility transition the elite were more likely to have communicated across long distances, while the farmers mainly communicated locally. Hence, it might be warranted to give the elite a lower w_r than the farmers. This would allow the elite to communicate more across regional boundaries and to be less dependent on the share of adopters in the own region of residence. In addition, based on theoretical considerations, it is reasonable to assume that the workers and others would be more affected by trickle-down effects from the elite vanguard group than the farmers, given that the workers and others were more likely to reside in urban centres in which the elite were concentrated, and that some of the workers were employed as household servants for the elite. Thus, the model might be more realistic if the farmers were assigned a lower w_s than the workers. Based on our model outcomes, we assume that introducing social class variation in the w_r and w_s specifications would have further increased the differences between the three social classes in the temporal pace of the transition, and that the farmers in particular would have lagged further behind. We are also able to run the models with much more complex geographies, as we can disaggregate our data down to the level of the more than 2400 parishes of Sweden at that time. For reasons of simplicity we focused here on a relatively simple geography, but it would certainly be interesting to explore how using a more complex geography would affect our model results.

In interpreting the outcomes of our simulations it is important to point out that communication links created by migration decisions are also shaped by processes of socio-economic change and national development. Due to this, the existing communication links might serve to some degree as a proxy for these processes. It is also likely that migration patterns are related to spatial variation in social norms, as regional and local populations with a low share of migrants might be less open to social change than populations with a high share of migrants. But independent of these considerations, our outcomes demonstrate that even in a static society with no spatiotemporal and social variation in adaptation pressure, frequently observed fertility decline patterns can be reproduced with simulations that model the decline as an information diffusion process structured by social and spatial variation in communication links.

Our findings might be relevant for less developed countries that are still at the beginning of the fertility transition. If communication processes indeed play an important role in the process, then we might even see a decline under conditions of slow socio-economic development, once a certain number of people have adopted the new behaviour. However, we have to be very careful in making such assumptions, as unlike Sweden of the late nineteenth century, contemporary African societies might lack clearly structured social groups who share common values and ideas, and who could serve as communication communities (see Caldwell and

Caldwell 1987). Yet it should be noted that this view is not undisputed (Makinwa-Adebusoye 2007). We also cannot rule out the possibility that the Swedish case is rather a peculiar one. The fact that infant mortality in Sweden had declined to low levels several decades before the onset of the fertility decline might have helped to create the structural conditions for change across the country before the decline started. Given this temporal lag between the emergence of conditions for change and the change itself, it is possible that spatial and temporal variation in communication links shaped the pattern to a large extent. This process might unfold differently in societies in which the emergence of conditions for change and the fertility transition are occurring in parallel. For example, the continuation of high infant mortality rates among specific regions or groups might represent a bottleneck condition for the transition process (see also Coale 1973). However, since infant mortality in Africa has declined considerably in recent years (Storeygard et al. 2008), it is not unwarranted to assume that in Africa we might also be witnessing a time lag between changes in structural conditions and the adaptation of fertility behaviour to these new conditions.

Our outcomes are also of general relevance to researchers interested in identifying causal relationships between structural factors and present-day demographic change processes that are likely to be shaped by a mixture of communication processes and adaptation to changing circumstances. In recent times, such processes might include the spread of cohabitation and the diffusion of gender-egalitarian norms in terms of the division of household and childrearing tasks. If these processes are predominantly driven by communication processes, they are still likely to be structured by social and spatial variation in communication links, with the latter also being shaped by structural factors. In interpreting the statistical associations between structural factors and demographic outcomes, researchers should be aware that these associations do not have to evolve as a result of a direct causal relationship between structural factors and demographic outcomes. They can also emerge if a communication process is not directly causally linked to these conditions, but just moderated by social and spatial variation in communication links that are themselves shaped by these structural factors. Our *ceteris paribus* simulations demonstrate that the shaping of the communication links and the diffusion through communication do not necessarily have to occur simultaneously. Agent-based models seem to offer great potential for investigating these questions.

Acknowledgements The authors would like to thank Frans Willekens for discussions on the specifications of the models, and for his input and suggestions on earlier versions of this contribution. We also thank the anonymous reviewers and the editors for their comments and suggestions, and Miriam Hils for language editing. This work is part of the project “Towards the modern family. Socioeconomic stratification, family formation and fertility in a historical perspective”, funded by the Swedish Research Council and the Crafoord Foundation.

Appendix 1: Map of Swedish Counties in 1880



Base Map: The Swedish National Archives, MPIDR Population History GIS Collection

References

- Bengtsson, T., & Dribe, M. (2014). The historical fertility transition at the micro level: Southern Sweden 1815–1939. *Demographic Research*, 30, 493–534.
- Bivand, R. S., Pebesma, E., & Gómez-Rubio, V. (2013). *Applied spatial data analysis with R* (2nd ed.). New York: Springer.
- Bourdieu, P., & Wacquant, L. J. D. (1992). *An invitation to reflexive sociology*. Chicago: University of Chicago Press.
- Caldwell, J. C., & Caldwell, P. (1987). The cultural context of high fertility in Sub-Saharan Africa. *Population and Development Review*, 13(3), 409–437.
- Carlsson, G. (1966). The decline of fertility: Innovation or adjustment process. *Population Studies*, 20(2), 149–174.
- Castroline, J. B. (Ed.). (2001). *Diffusion processes and fertility transition: Selected perspectives*. Washington, DC: National Academy Press.
- Cleland, J., & Wilson, C. (1987). Demand theories of the fertility transition: An iconoclastic view. *Population Studies*, 41(1), 5–30.
- Coale, A. J. (1973). The demographic transition reconsidered. In International Union for the Scientific Study of Population (IUSSP) (Ed.), *Proceedings of the International Population Conference 1973 1* (pp. 53–73). Liège: International Union for the Scientific Study of Population.
- Coale, A. J., & Watkins, S. C. (Eds.). (1986). *The decline of fertility in Europe: The revised proceedings of a conference on the Princeton European Fertility Project*. Princeton: Princeton University Press.
- Costa, R. (2014). The fertility transition in Belgium: A contagious process? *Poster presented at the European Population Conference, Budapest, Hungary. 25–28 June 2014*.
- Dribe, M. (2009). Demand and supply factors in the fertility transition: A county-level analysis of age-specific marital fertility in Sweden, 1880–1930. *European Review of Economic History*, 13(1), 65–94.
- Dribe, M., & Scalone, F. (2014). Social class and net fertility before, during and after the demographic transition: A micro-level analysis of Sweden 1880–1970. *Demographic Research*, 30, 429–464.
- Dribe, M., Hacker, J. D., & Scalone, F. (2014a). The impact of socio-economic status on net fertility during the historical fertility decline: A comparative analysis of Canada, Iceland, Sweden, Norway, and the USA. *Population Studies*, 68(2), 135–149.
- Dribe, M., Oris, M., & Pozzi, L. (Eds.). (2014b). Socioeconomic status and fertility before, during, and after the demographic transition: An introduction. *Demographic Research*, 31, 161–182.
- Dribe, M., Helgertz, J., & van de Putte, B. (2015). Did social mobility increase during the industrialization process? A micro-level study of a transforming community in Southern Sweden 1828–1968. *Research in Social Stratification and Mobility*, 41(1), 25–39.
- Dyson, T. (2011). The role of the demographic transition in the process of urbanization. *Population and Development Review*, 37(suppl. 1), 34–54.
- Galloway, P. R., Hammel, E. A., & Lee, R. D. (1994). Fertility decline in Prussia, 1875–1910: A pooled cross-section time series analysis. *Population Studies*, 48(1), 135–158.
- Goldstein, J. R., & Klüsener, S. (2014). Spatial analysis of the causes of fertility decline in Prussia. *Population and Development Review*, 40(3), 497–525.
- González-Bailón, S., & Murphy, T. E. (2013). The effects of social interactions on fertility decline in nineteenth-century France: An agent-based simulation experiment. *Population Studies*, 67(2), 135–155.
- Hägerstrand, T. (1965). A Monte Carlo approach to diffusion. *European Journal of Sociology*, 6(1), 43–67.
- Haines, M. R. (1992). Occupation and social class during fertility decline: Historical perspectives. In J. R. Gillis, L. A. Tilly, & D. Levine (Eds.), *The European experience of declining fertility, 1850–1970: The quiet revolution* (pp. 193–226). Cambridge: Blackwell.

- Hofsten, E., & Lundström, H. (1976). *Swedish population history: Main trends from 1750 to 1970*. Stockholm: National Central Bureau of Statistics.
- Klüsener, S., Scalone, F., & Dribe, M. (2013). Spatial vs. social distance in the diffusion of fertility decline: Evidence from Sweden 1880–1900. *Paper presented at the XXVII IUSSP International Population Conference, Busan*. 26–31 August 2013.
- Klüsener, S., Dribe, M., & Scalone, F. (2016). *Fertility transition as a communication process*. Unpublished manuscript.
- Kohler, H.-P. (2001). *Fertility and social interaction: An economic perspective*. Oxford: Oxford University Press.
- Lindström, M. (2001). Conservative voting and the church: The religious dimension in the electoral geography of the conservative (moderate) party in Sweden 1921–1998. *Statsvetenskaplig Tidskrift*, 104(4), 305–328.
- Livi-Bacci, M. (1986). Social-group forerunners of fertility control in Europe. In A. J. Coale & S. C. Watkins (Eds.), *The decline of fertility in Europe: The revised proceedings of a conference on the Princeton European Fertility Project* (pp. 182–200). Princeton: Princeton University Press.
- Makinwa-Adebusoye, P. (2007). Sociocultural factors affecting fertility in Sub-Saharan Africa. In United Nations (Ed.), *Prospects for fertility decline in high fertility countries* (Population bulletin of the United Nations, Special issue nos. 46/47, pp. 55–69). New York: United Nations.
- MPIDR [Max Planck Institute for Demographic Research]. (2014). *MPIDR Population History GIS Collection*. Rostock: MPIDR.
- Rosero-Bixby, L., & Casterline, J. B. (1994). Interaction diffusion and fertility transition in Costa Rica. *Social Forces*, 73(2), 435–462.
- Ruggles, S., Roberts, E., Sarkar, S., & Sobek, M. (2011). The North Atlantic Population Project: Progress and prospects. *Historical Methods*, 44(1), 1–6.
- Scalone, F., & Dribe, M. (2012). Socioeconomic status and net fertility in the demographic transition: Sweden in 1900 – A preliminary analysis. *Popolazione e Storia*, 11(2), 111–132.
- Schmertmann, C. P., Assunção, R. M., & Potter, J. E. (2010). Knox meets Cox: Adapting epidemiological space-time statistics to demographic studies. *Demography*, 47(3), 629–650.
- Shyrock, H. S., & Siegel, J. S. (1980). *The methods and materials of demography*. Washington, DC: U.S. Department of Commerce, Bureau of the Census.
- Storeygard, A., Balk, D., Levy, M., & Deane, G. (2008). The global distribution of infant mortality: A subnational spatial view. *Population, Space and Place*, 14(3), 209–229.
- Szreter, S. (1996). *Fertility, class and gender in Britain 1860–1940*. Cambridge: Cambridge University Press.
- Van Bavel, J. (2004). Diffusion effects in the European fertility transition: Historical evidence from within a Belgian town (1846–1910). *European Journal of Population/Revue Européenne de Démographie*, 20(1), 63–85.
- van Leeuwen, M. H. D., & Maas, I. (2011). *HISCLASS: A Historical International Social Class Scheme*. Leuven: Leuven University Press.
- van Leeuwen, M. H. D., Maas, I., & Miles, A. (2002). *HISCO: Historical International Standard Classification of Occupations*. Leuven: Leuven University Press.
- van Poppel, F. W. A. (1985). Late fertility decline in the Netherlands: The influence of religious denomination, socioeconomic group and region. *European Journal of Population/Revue Européenne de Démographie*, 1(4), 347–373.
- van Poppel, F., Reher, D. S., Sanz-Gimeno, A., Sanchez-Dominguez, M., & Beekink, E. (2012). Mortality decline and reproductive change during the Dutch demographic transition: Revisiting a traditional debate with new data. *Demographic Research*, 27, 299–338.

Chapter 14

Feedback Mechanisms in the Postponement of Fertility in Spain

Daniel Ciganda and Francisco Villavicencio

14.1 Introduction

The idea that reproductive preferences (intentions) represent a key element of fertility change was already present in classic demographic transition theory. Notestein (1953) recognized that the set of social norms, values, and beliefs sustaining a particular economic and demographic regime are “deeply woven into the social fabric and are slow to change.” For decades, demographers have attributed the lag in the decline in fertility during the demographic transition to people’s resistance to adapting to a new survival scenario.

Some of the most prominent contemporary fertility theories argue that normative shifts, either in the form of secularization (Lesthaeghe and Van de Kaa 1986) or changing gender values (McDonald 2000; Esping-Andersen and Billari 2015), are the main engine of change in fertility levels. However, most of the time norms seem to be resisting, rather than promoting, demographic change.

Norms have prevented not only more rapid fertility declines during the demographic transition, but also expected increases in the mean age at childbirth in Eastern European countries (Perelli-Harris 2005; Mynarska 2010), the adoption

D. Ciganda (✉)

Department of Political and Social Sciences, Pompeu Fabra University, Barcelona, Spain
e-mail: daniel.ciganda@upf.edu

F. Villavicencio

Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark

Max-Planck Odense Center on the Biodemography of Aging, Odense, Denmark
e-mail: villavicencio@imada.sdu.dk

of modern contraceptive methods (Munshi and Myaux 2006), and the spread of alternative family formation behaviors among Japanese men and women (Rindfuss et al. 2004).

But norms are too complex for most conventional modeling techniques and most of the time they remain at the theoretical level. In this chapter we explicitly model the dynamic relationship between norms and behavior and we argue that the evolution of (age) norms can also explain a substantial part of the postponement of fertility in Spain.

The explanation we propose has four central elements: the expansion of tertiary education, the increase in economic uncertainty, the role of social interaction as a multiplier of these structural changes, and the dynamic relationship between intentions and behavior. Our objective is to go beyond the assessment of the relative contributions of each of these elements, and to analyze the mechanisms at the individual level that explain the emergence of an aggregate trend.

We structure the chapter as follows. First, we discuss some of the approaches used to analyze the dynamics that have pushed fertility to later ages, and their limitations. Next, we describe the model target and the model itself, and present the empirical data used to calibrate it to the Spanish case. Finally, we present several simulated scenarios designed to help us gain insight into the mechanisms at play in the rise in the mean age at first birth (MAFB).

14.2 The Dynamics of Fertility Postponement

14.2.1 *Social Interaction*

Social interaction became a mainstream concept in demography after the results of the Princeton Project highlighted the geographic and linguistic boundaries surrounding the onset and the pace of fertility decline. Thereafter, the spread of information and attitudes through social networks was a key element of most models of fertility change during the demographic transition.

According to Casterline (2001), the initial attempts to incorporate into fertility theories the notion of diffusion were based on an eminently practical goal: namely, the acceleration of the spread of contraceptive techniques in developing countries. This might explain why birth control has been the main focus of most of the empirical applications of diffusion models (Entwisle et al. 1996; Kohler 1997; Kohler et al. 2001; Montgomery and Casterline 1993; Munshi and Myaux 2006; Rosero-Bixby and Casterline 1993).

The most recent wave of studies on social interaction and fertility have relied on the availability of detailed datasets and new methods to empirically analyze the role of social networks in family formation decisions (Aparicio Diaz et al. 2011; Balbo and Mills 2011; Balbo and Barban 2014; Lyngstad and Prskawetz 2010; Mathews and Sear 2013). These studies focused less on fertility *change* and more on the

question of how family and friends influence fertility attitudes and behaviors in the transition to parenthood. They defined the mechanisms through which interaction affects behavior: i.e. *social learning*, a self-initiated process through which agents obtain information and knowledge from others; or *social influence*, a process through which other people and their behaviors exert pressure and control over the individual.

From a dynamic perspective, social interaction also operates through different mechanisms. Of those proposed by Kohler et al. (2002), two are particularly relevant in our analysis: *social feedback mechanisms* and *status quo enforcement*.

Social feedback is the process through which social interaction increases the pace or the extent of the original change in fertility behavior triggered by socioeconomic changes. In this case, imitation or the intensification/relaxation of social pressure generates an effect that can become independent of the initial change in material conditions.

Status quo enforcement refers to the mechanism, described above, through which social norms generate resistance to innovative behavior. So far, there have been few empirical analyses of this process (for an example, see Munshi and Myaux 2006). But as we will try to show in the remainder of this chapter, status quo enforcement might be one of the fundamental dimensions of fertility change.

14.2.2 *Feedback Loops Between Preferences and Behavior*

The decision to have a child has been analyzed using multiple behavioral frameworks, like the “ready, willing, and able” approach developed by Coale (1973), or the rational action theory embedded in most economic models. Recently, attempts have been made to promote and incorporate the theory of planned behavior (TPB) (Ajzen 1985) into the analysis of reproductive decision-making (Testa et al. 2011).

Demographers became interested in the TPB primarily because of the role intentions play as close predictors of observed behavior within the TPB framework. In fact, the main focus of empirical studies on intentions has been to test their capacity to predict future fertility trends (for a review, see Morgan 2001).

However, the TPB framework has been criticized for its static nature, as the theory does not adequately take into account the recursive loop between intentions and behavior (Sniehoff et al. 2014; Morgan and Bachrach 2011). The focus of the TPB on the synchronic perspective blurs the process through which intentions themselves change as a result of previous behaviors.

The discussion surrounding the TPB evokes the longstanding sociological debate between those who see agency (behavior) as the key mechanism in the explanation of social processes, and those who highlight the role of structures (norms, institutions). Other disciplines have also referred to this debate as “voluntarism vs determinism”. Scholars who have attempted to overcome this dichotomy have generally focused on the recursive dynamics between agency and structure (Giddens 1984; Bourdieu and Wacquant 1992).

This perspective implies the existence of a mechanism through which new behaviors contribute to a change in the prevailing norms, which in turn feed back to the individual level, thereby shaping the intentions of subsequent generations in a micro-macro loop.

While they are harder to model, these mechanisms can illuminate more complex, and potentially more interesting dynamics like the “downward spiral of fertility” described by Goldstein et al. (2003) (who argued that the *apparent* decline in the ideal family size in Germany and Austria was a product of the experiences of cohorts living in low-fertility regimes), or the emergence of multiple equilibria regarding the timing of fertility described by Kohler et al. (2002). In the latter, feedback generated through social interaction explains rapid and substantial changes in the MAFB.

14.2.3 *Equilibrium Between Intentions and Behavior?*

Even if we imagine that intentions and behavior are mutually dependent, the question of what this relationship looks like remains. Previous postponement models have assumed that intentions and behavior converge, at least in the long run.

In the model by Aparicio Diaz et al. (2011), the decision to have a child at time t results in a new aggregated probability (intention) at $t+1$. Here the effect of previous behavior in the updating of people’s preferences is linear and cumulative. Intentions and behavior remain perfectly aligned. The model by Kohler et al. (2002) allows for a temporary mismatch of intentions (the desired MAFB) and behavior (the observed MAFB), although in the long run they also converge to an equilibrium. However, the available empirical evidence on intentions suggests that this assumption might be misleading.

Regarding the quantum of fertility, Bongaarts (2001, p. 261) argued that in pre-transitional societies the ideal family size tends to be below the observed total fertility rate (TFR), but that this relationship shifts in post-transitional settings. He concluded that “a declining desired family size is indeed one of the principal forces driving fertility transitions, but in reality levels of fertility often deviate substantially from stated preferences.”

Unfortunately, the amount of available data on preferences regarding the timing of the transition to parenthood is very limited, and comparisons of the evolution of the ideal relative to the observed mean ages are difficult to make. However, data from 2006 showed that, in European countries, people’s preferences and behavior were far from being in equilibrium (Testa 2006). For example, women surveyed in Spain stated that the ideal age for becoming a mother was around 25.5, while the observed mean age at first birth for that year was 29.3.

The scarce evidence available suggests then that the evolution of the *ideal mean age at first birth* relative to the *observed mean age* will follow a similar pattern to the one observed for the ideal family size relative to the TFR: i.e. above the ideal before a certain amount of control over fertility has been achieved, and below the ideal after the mean age at motherhood has been pushed beyond a certain threshold;

as appears to have been the case in most European countries in the mid-2000s (Testa 2006).

14.2.4 Educational Change

The positive association between education and the age of the transition to parenthood has long been recognized in demography (Marini 1984; Rindfuss et al. 1980). Research on the topic has identified two distinct dimensions of the effect of education on fertility: enrollment and post-enrollment. The first effect refers to the difficulties which can arise in balancing the roles of student and mother, while the second is related to the higher opportunity costs of childbearing for highly educated women.

According to recent estimates of the contribution of increasing education to the postponement of fertility in three European countries (Britain, France, and Belgium), the joint effects of enrollment and post-enrollment are responsible for most of the rise in the mean age at first birth (Neels et al. 2014; Ní Bhrolcháin and Beaujouan 2012). However, the findings of another set of studies suggest that education might not be the main driver of fertility postponement (Rendall et al. 2010; Rindfuss et al. 1996). The results presented later in this chapter lend support to the second perspective, as they indicate that educational expansion explains only a modest fraction of the total delay in marriage/parenthood.

14.2.5 Unemployment

The relationship between economic constraints and fertility decisions has received substantial attention in recent years after a series of studies suggested that the long-standing negative correlation between prosperity and fertility levels had changed its sign (Adsera 2004; Ahn and Mira 2001; Kohler et al. 2002; Myrskylä et al. 2009). Although a vast body of literature on this topic has been generated, it is still difficult to obtain some stylized facts regarding the size and direction of the effect of unemployment on fertility decisions.

A series of studies have reported no significant effects (Kravdal 2002; Kreyenfeld 2010; Ozcan et al. 2010), while other studies have found strong negative effects that range from a 25 % reduction in the risk of having a first birth in France (Pailhé and Solaz 2012), to a 60 % reduction in Germany (Kreyenfeld 2005), and a 40 % reduction in Spain (Baizán 2006).

The considerable diversity of the institutional settings studied and the difficulties researchers face in distinguishing between income and substitution effects partially

explain the ambiguity of some of the results found in the literature.¹ A potentially fruitful strategy for disentangling income from substitution effects is to distinguish women in traditional male-breadwinner arrangements from women in dual-earner households. Unfortunately, the information needed to make this distinction is often missing from surveys, and only a few studies have provided results accounting for the employment status of both members of the couple. Most of these studies have confirmed the assumption that substitution effects prevail when women are (exclusively) caregivers, while unemployment tends to depress or delay fertility when both members of the household work (Baizán 2006; Vignoli et al. 2012).

To account for the effect of unemployment in our model, we had to reconstruct the historical series of unemployment rates in Spain from the earliest available figures in the 1930s (Fig. 14.1). The shortest series corresponds to the unemployment rates computed from the information provided by the Spanish Labor Force Survey (EPA, *Estadística de Población Activa*), which is considered the most reliable source in Spain for labor market indicators, including unemployment rates. The second series shows the figures obtained from the official employment offices. Although it covers a longer period of time, this indicator only consider workers in the formal economy, which could lead to an underestimation of the unemployment




Fig. 14.1 Unemployment rates, Spain 1933–2012: (1) Unemployment measured by the Spanish Labor Force Survey (EPA), and (2) Unemployment registered (Statistical Yearbooks of Spain) (Source: Spanish Statistical Office (INE 2015))

¹According to Becker (1981) an income rise will not only increase the demand for children, but also the indirect costs of forming a family; i.e. the potential income and career opportunities that parents have to give up in order to spend time with their children. An *income* effect is observed when the demand for children is positively affected by an increase in resources, and a *substitution* effect is observed when the effect is negative.




Fig. 14.2 Mean age at first birth (MAFB), Spain 1975–2013 (Source: Spanish Statistical Office (INE 2015))

rates. The other crucial difference between the two indicators is that the registered unemployment rate is computed over the working-age population (16–64), whereas the EPA rate considers only the economically active population, and thus provides higher estimates. However, our goal in presenting both series is not to highlight their differences, but to show that Spain seems to have enjoyed a period of very low unemployment until the 1970s, when the rate increased dramatically, coinciding with the increase in the MAFB, presented in Fig. 14.2.

14.3 Model Target

The main targets of our model are the evolution of the mean age at first birth (MAFB) and the evolution of the schedule of age-specific fertility rates (ASFR), which for Spain are available only from 1975 onward.

Although most European countries registered significant postponements in the MAFB in the last decades of the twentieth century, in Spain the increase was particularly intense. As we can see in Fig. 14.2, the mean age increased by about 5 years over the observed period.

An initial exploration of the curve suggests an effect of unemployment in shaping the trend: a similarly steep increase until the mid-1990s, a deceleration up to the end of the first decade of the twenty-first century, and another peak coinciding with the most recent economic crisis.

Figure 14.3 provides more information about the nature of this change. The evolution of the age-specific fertility rates shows that the increase in the mean age




Fig. 14.3 Age-specific fertility rates, Spain 1975–2012 (Source: Spanish Statistical Office (INE 2015))

shown in Fig. 14.2 has been the result of a reduction in fertility rates at younger ages, but also of the increase in births at older ages. From 1975 to 2012 the peak of the distribution shifted from around age 23 to age 30. It is also interesting to note the small bump in the distribution from age 17 to age 24, which likely represents the contribution of migrant women, and which prevented an even greater increase in the mean age at the transition to motherhood by the end of the period.

14.4 Model Layout

In our model, each woman older than age 15 has the intention to form a union or to have a child. These intentions are represented by transition probabilities which depend on individual characteristics (age and education) and on whether the agent has already found a partner or not.

The final probability with which an agent makes a decision to get into a union is the result of the agent's original intention (u), her employment status, and the influence coming from her network of friends (social influence). Analogously, the final probability of having a child inside the union is the result of the couple's original intention (f), their employment status, and the social influence of their network of friends. We assume, however, that having a child outside the union is not influenced by the same intervening factors, as most of these births will not be the result of a well thought-out decision; so even though we consider these events, they are not affected by unemployment or social influence. Each of these elements

will be described in greater detail in Sect. 14.5; here we provide a brief description of the central mechanisms.

The initial intentions u and f are obtained from empirical data, as described in Sect. 14.5.1. They capture both information on the behavior of previous cohorts and the agent’s ideas regarding the maximum ideal age to marry/become a parent. The behavior of previous cohorts gives the agent (couple) an idea of what others can expect regarding her (their) family formation decisions given the age and education level. In this regard, u and f also capture one of the levels of social influence; the second one, a more local level, is captured through the agent’s network. These expectations, however, are bounded by social (and biological) limits, which means norms determine to a certain extent how much and how fast intentions change.

Each year, before making a decision, agents consider not only their intentions but also their employment status and the behavior of their friends. An *unemployment multiplier* is introduced to capture the reduction in the probability of experiencing the transition to marriage (parenthood) when agents experience involuntary spells out of the labor market (Sect. 14.5.5).

Analogously, a *social influence multiplier* captures the influence from the closest network of friends and at the same time provides agents with information about how conservative/innovative the behavior of their friends is in comparison with the behaviors of previous generations. Social interaction positively reinforces those behaviors that are becoming increasingly acceptable/common in the population (Sect. 14.5.6).

The effect of educational expansion is introduced at the level of the education-specific intentions by changing the composition of the population by educational level as described in Sect. 14.5.4.

Figure 14.4 presents an example of a 25-year-old woman with tertiary education. Let’s denote by $u(t, 25, 3)$ —25 referring to the age and 3 to the educational level—her intention to get into a union in a given year t , which depends originally on how common or acceptable it is for a university-educated woman to marry (or cohabit)

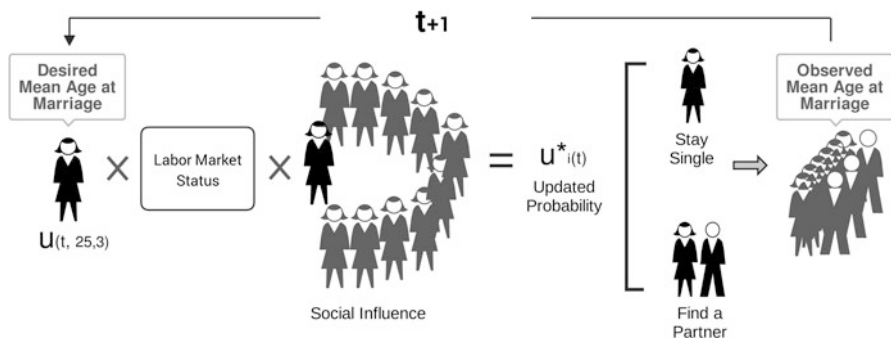


Fig. 14.4 A model of fertility postponement: Union formation. The example of a 25-year-old and university-educated woman

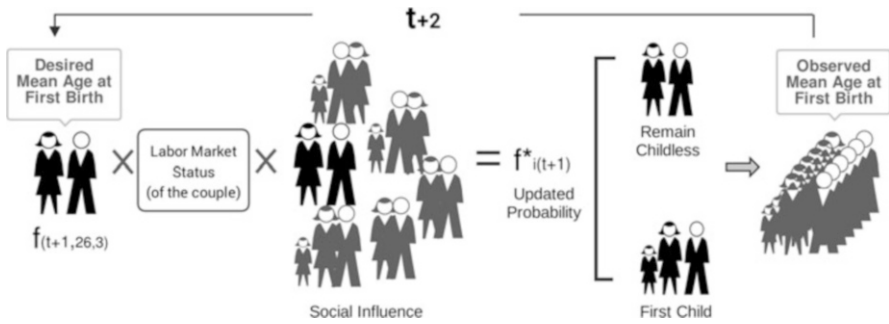


Fig. 14.5 A model of fertility postponement: First birth. The example of a couple formed by a 26-year-old and university-educated woman and her partner

at age 25. The desirability/acceptability is given by the proportion of university-educated women of that age who got married in the previous year.

After considering her own employment situation and the behavior of her peers (social influence), our woman updates her original intention to marry and makes a decision based on her individual updated intention $u_i^*(t)$. The individual decisions of each of the 25-year-old and university-educated women in the population modify the norm regarding marriage at that age and educational level. The new information on how common or acceptable it is to enter a union for that population group is used in the next time step $t + 1$ to form the baseline intentions at age 25 of the following cohort of university-educated women— $u(t + 1, 25, 3)$ —in a micro-macro loop.

When a woman decides to marry based on the prevailing norms, her background characteristics, her employment situation, and the influence from her network, the agent in the model that represented that woman now represents a couple, as depicted in Fig. 14.5. In the following year $t + 1$, the new couple—formed by a 26-year-old and university-educated woman and her partner—start with the intention to have a child $f(t + 1, 26, 3)$, which is shaped by the behavior of other couples the year before. In this case, the relevant background characteristics (age and education) are those of the woman, which are not necessarily the same as those of her partner.

As in the decision to marry, the couple update their original intention based on the influence of their network of friends, and after taking into consideration their joint situation in the labor market. The final updated intention to have a child will be $f_i^*(t + 1)$. This process resembles that of the Bongaarts (2001) model, in which a series of intervening factors prevent couples from realizing their fertility intentions.

The decisions of all couples in the population affect people’s perceptions regarding the desirability/acceptability of such behavior. These perceptions in turn modify the fertility intentions of subsequent generations of couples, and, consequently, the desired mean age at first birth.

However, the interesting question here is how exactly people come to update their expectations and form their preferences by taking into account the experiences of previous cohorts. In other words, how norms adapt to people’s behavior and how in turn they shape their future actions.

The existence of ideal ages as well as age *deadlines* for the transition to parenthood has been well documented (Settersten and Hägestad 1996; Billari et al. 2011; Van Bavel and Nitsche 2013). It is likely that these markers, especially those related to the *proper* age for having a child, change during the postponement transition as a result of the mechanism described above. It is also likely, however, that the ideal age for having a child does not increase indefinitely. A certain threshold must exist after which people start resisting the push toward later childbearing ages.

Our assumption is that people will follow the behavior of previous cohorts and update their expectations as long as the threshold that marks the upper limit of the ideal age for having a first child had not been crossed. The threshold marks the point at which individuals will start resisting further increases of the MAFB even if the material incentives push in the opposite direction. We believe that this resistance is triggered by the proximity of the biological limit but also by a series of social norms, for which we model different thresholds according to the education level of agents.

In addition, we believe that the postponement of motherhood encounters some resistance at the beginning of the process due to a similar path dependence process generated by prevailing norms. It takes time for individuals to realize that conditions are changing and to start adapting their preferences accordingly. In Sect. 14.5.8 we provide details on how these dynamics are introduced in the model.

To summarize, we understand the process of postponement as the result of the interaction of the four main factors introduced above. It is in response to *rising economic uncertainty* and increased opportunity costs associated with the *expansion of higher education* that couples in Spain start postponing marriage and having children. The expansion of education and the increasing uncertainty in the labor market provide the original push to the MAFB. However, this original change is amplified and sustained via *social interaction* as young men and women start imitating the behavior of their peers and friends. People's beliefs about the ideal age for marriage and for becoming a mother *social norms* play a crucial role both at the beginning and at the end of this process. Initially, norms generate resistance associated with the time it takes for individuals to realize that conditions and expectations are changing. Towards the end of the process, the resistance is generated by the proximity to the social and biological limits of reproduction.

14.5 Technical Description

Our simulation runs for 70 years, from 1944 until 2014, and we base our simulations on Spanish data (more details about the data are provided in Sect. 14.6). The ages of the initial population are randomly assigned according to the female population structure of Spain from the 1940 census (INE 2015). Starting the simulation in 1944 ensures that all of the initial women of reproductive ages (15+) will be out of their reproductive period when our analysis of the MAFB begins in 1975.

The model contains five procedures for agents which are carried out at each time step: aging, partnership formation, reproduction, entry into/exit from the labor market, and the building of a network. Each time step (or iteration) corresponds to 1 year. At each new iteration the agents age, and they may die off according to the corresponding age- and year-specific mortality rates of Spain (HMD 2015).

As they enter the simulation, the agents are assigned the final educational level they will achieve: primary, secondary, or tertiary. At age 15, an individual becomes an adult who can find a partner (marriage or cohabitation), who might reproduce, and who builds her own social network by choosing a maximum of ν contacts from a larger pool of potential friends based on their social distance with respect to education. From age 16 agents can become unemployed according to observed age- and sex-specific probabilities. The agents do not remain in the simulation beyond age 45.

The population is composed of female and couple agents. In addition to the information of the female partner, couple agents have information on the male member, as shown in Table 14.1. Table 14.2 summarizes the global parameters used in the simulation.

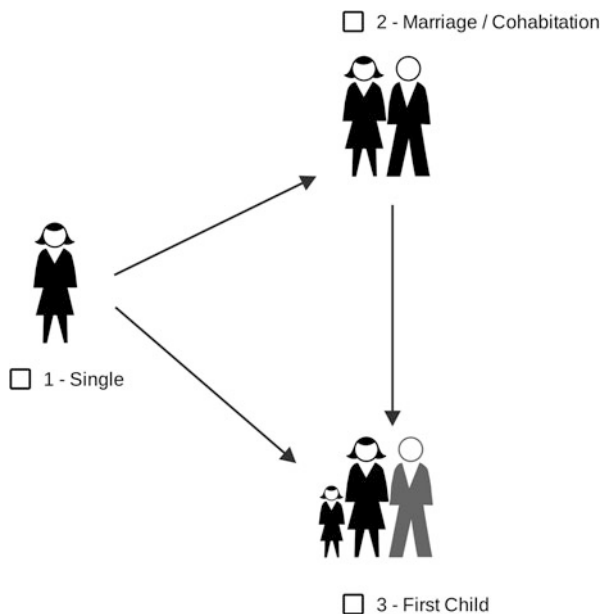
Table 14.1 Agents’ characteristics

Agent variables	Variable name	Values
Identity number	<i>id</i>	1, 2, 3, ...
Age	<i>x</i>	0–45
Age partner	<i>xp</i>	15–53
Age at first birth	<i>xb</i>	15–45
Education level	<i>edu</i>	1: “primary” 2: “secondary” 3: “tertiary”
Marital status	<i>ms</i>	0: “single” 1: “married/cohabitation”
Employment status	<i>es</i>	0: “agent employed” 1: “agent unemployed”
Employment status partner	<i>esp</i>	0: “partner employed” 1: “partner unemployed”
Network	<i>net</i>	# individuals in the network

Table 14.2 Global parameters

Global parameters	Parameter name	Value
Starting year	<i>iniYear</i>	1944
Final year	<i>finYear</i>	2014
Sex ratio at birth	<i>SRB</i>	0.515
Minimum age at birth	<i>minAge</i>	15
Maximum age at birth	<i>maxAge</i>	45
Recurrence in unemployment	σ	0–0.99
Maximum network size	ν	1, 2, 3, ...
Fertility rate for parity 1+	<i>f2</i>	0–0.99

Fig. 14.6 Multistate model:
(1) Single, (2) Marriage/
Cohabitation, and (3)
First Child



14.5.1 Transition Probabilities

As was mentioned above, the intentions on which our agents' behaviors are based are represented by empirical transition probabilities. To obtain these probabilities, we first estimate the original union and fertility rates using the multistate model presented in Fig. 14.6. The model resembles the classic illness-death model without recovery used in medical research (Beyersmann et al. 2011). All of the individuals start at stage 1-*Single* (without children), and have the potential to stay in this state or to move to either state 2-*Marriage/Cohabitation*, or to the absorbing state 3-*First Child*. After reaching state 2, individuals can either stay or leave the state to enter state 3. Each of these transitions is governed by a cause-specific hazard from which we obtain the Nelson-Aalen estimators of the cumulative hazard for each event.

From the estimators of the cumulative hazard, we derive the sets of age- and education-specific fertility probabilities for a first birth inside of a union $f^o(x, edu)$, fertility probabilities for a first birth outside of a union $fs^o(x, edu)$, and union probabilities $u^o(x, edu)$. These sets of probabilities represent both the observed and the intended fertility behaviors at the beginning of our period of interest, as, for the sake of simplicity, we assume that behaviors and intentions are in equilibrium at that time. For parity one or higher, the fertility rate is f_2 for all individuals.²

²As we model the effect of labor market exits exogenously, we need a set of initial probabilities that is net of the effect of unemployment to avoid an overestimation of this effect. Unfortunately, as the dataset we use for the estimation of the original probabilities does not contain information on

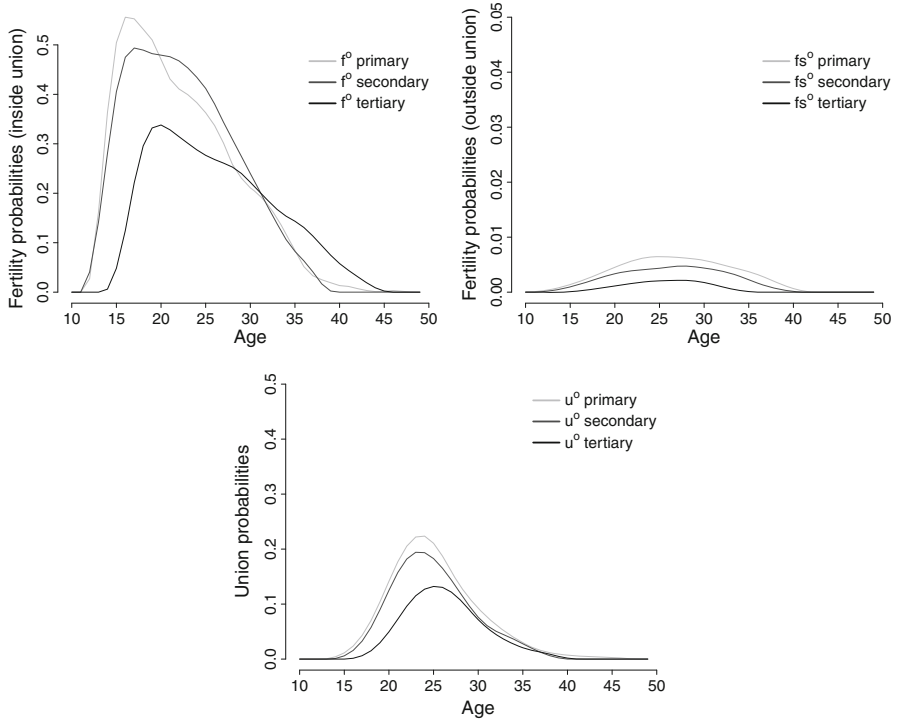


Fig. 14.7 Age- and education-specific probabilities from multistate model. Cohorts 1940–1960, Spain. From left to right: (1) Fertility probabilities inside union, (2) Fertility probabilities outside union, and (3) Union probabilities. Note: A different vertical scale is used for fertility probabilities outside the union

14.5.2 Initialization of the Model

The observed fertility and union rates correspond to the cohorts of females born between 1940 and 1960. We restrict our sample to this period for two reasons. First, we want to avoid mixing a large number of cohorts with different labor market and educational experiences. Second, the fertility schedules of these cohorts shaped the period mean ages at first birth in the mid-1970s, before the beginning of the increase we are trying to model (see shadowed area in Fig. 14.8).

For the computation of the MAFB, we need all of the agents to be exposed to the entire set of intensities presented in Fig. 14.7. Hence, starting the simulation

the employment histories of the interviewees, we have to provide a rough estimate of the effect of unemployment. As we noted in Sect. 14.2.5, Spain did not have high levels of unemployment until the mid-1980s, which means that the effect of unemployment on our cohorts born in 1940–1960 would have been relatively mild. We assume a decreasing effect by age: compared with the original probabilities, the final probabilities are about 15 % higher at age 15, only around 5 % higher at age 30, and about the same by the end of the reproductive period at age 45, as shown in Fig. 14.7.




Fig. 14.8 Lexis diagram of the female cohorts used in the initialization of the model

in 1944 ensures that all of the women who entered the initial population being older than 15 (solid black line in Fig. 14.8) and were not exposed to the entire set of probabilities from age 15 to 45 are not considered in the computation of the simulated MAFB in 1975. Moreover, this initialization procedure (1944–1974) prevents us from assigning an initial parity to the agents and an age at first birth to those with parity one or higher, for which we have no empirical reference.

14.5.3 Age of Partners

When an agent enters a union, a random age and the corresponding age-specific unemployment rate are assigned to her partner. The age is obtained from a truncated normal distribution, using the age of the agent plus two years ($x_i + 2$) as the mean value, and $a = x_i - 4$ and $b = x_i + 8$ as the lower and upper limits of the age of the partner:

$$xp_i \sim N_{[a,b]}(x_i + 2, 1). \tag{14.1}$$

14.5.4 Evolution of Educational Attainment

Tertiary education has expanded rapidly among women in Spain: the share of women who completed tertiary education rose from 5 % of those born in the late 1930s, to one-third of those born in the 1970s, to around 45 % of the more recent cohorts (Castro-Martín and Martín-García 2013).




Fig. 14.9 Observed and predicted proportion of newborns by achieved education level. Female cohorts, Spain 1934–1998 (Source: 2011 Spanish Census, Spanish Statistical Office (INE 2015))

As we can see in Fig. 14.9, since the late 1930s the proportion of female newborns who complete tertiary education has been increasing almost linearly. The 2011 census provides reliable figures for the generations born up to 1980; after that point, and based on the figures presented above, we assume a continuation of the linear increase in the proportion of women who access a tertiary level education, a linear declining trend in the share of women who only have secondary education, and a stagnation in the share of women who never progress beyond the primary level of education (at under 10%).

In our model, education is defined according to the three levels mentioned above, which correspond to the number of years of formal schooling: fewer than six (*primary*), from six to 13 (*secondary*), and more than 13 (*tertiary*). Each year we assign the newborns in our model a level of education matching the proportions that each of these levels represent in the total female population, as shown in Fig. 14.9.

14.5.5 Unemployment Effects

Starting from the observed unemployment rates, we obtain the proportion of the population who were unemployed in each year of our simulation, by age group and sex. From these series we then model the exits from the labor market, while assuming that a proportion σ of those who were unemployed in the previous year will stay in that state.

We assume that unemployment affects an individual's decisions about whether and when to enter a union and to have a first child within a union. The strength of this

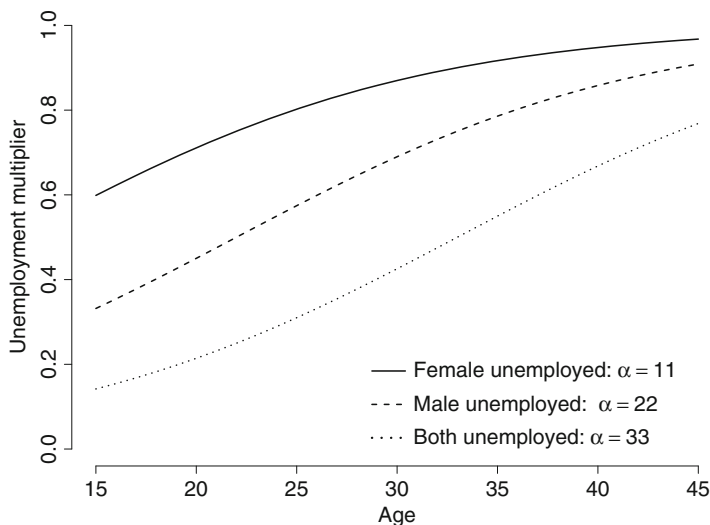


Fig. 14.10 Effects of unemployment by age on the original fertility and union rates

effect depends on the employment status of the agent or the employment status of the couple after union formation, with three different scenarios: (1) the agent is unemployed ($es_i = 1; esp_i = 0$); (2) her partner is unemployed ($es_i = 0; esp_i = 1$); and (3) both members of the couple are outside of the labor market ($es_i = esp_i = 1$). For single agents $esp_i = 0$. In practice, that means that the original fertility and union rates of each agent i are modified by an unemployment multiplier um_i defined as

$$um_i(x_i; \alpha) = \frac{1}{1 + \exp(-0.1(x_i - \alpha))}, \quad (14.2)$$

where x_i is the age of the agent, and α a parameter that depends on the employment status of the couple and determines the strength of the effect. Figure 14.10 shows the unemployment multiplier over the union rates $u(t, x, edu)$ for the different combinations of the employment status of both members of the couple. These sets of α values are the ones we used in our final model in Sect. 14.7.1, which provides a good fit for our target.

In all three cases the effect is more pronounced at younger ages. For instance, for those couples in which only the female partner is unemployed the probability of entering a cohabiting union or having a first child (after being in a union) is about 80% of the original probability if the female partner is 25 years old, about 65% if the male partner is unemployed, and about 30% if both partners are out of the labor market.

The function of the multiplier for the fertility probabilities is the same, although the effects are slightly smaller than the multiplier for the union probabilities. We assume that the margin to postpone the formation of a new household in response to economic uncertainty is larger than the margin to postpone the decision to have children given biological constraints.

14.5.6 Social Influence Effects

As mentioned before, agents do not base their decisions about the timing of marriage/cohabitation and the transition to parenthood exclusively on their immediate material conditions (unemployment) and their perceived opportunities (education); their choices are also influenced by their network of friends. As the decisions of the members of a given agent’s network are themselves determined by the labor market conditions and their education, the strength and the direction of the influence the network exerts will be shaped by its members’ material conditions and their perceived opportunities. In other words, a social feedback mechanism is triggered by a change in these conditions, which in turn reinforces and amplifies the original effect.

At age 15, each agent forms a network composed of a maximum of ν members of the same age. The agents randomly choose these contacts from a pool of potential friends based on a social distance function that depends on their educational level. The social distance between two agents i and k is defined as

$$sd_{ik} = \exp(-\beta(|edu_i - edu_k| + 1)^2), \tag{14.3}$$

where edu_i and edu_k are the respective educational levels, and β is a parameter that controls the level of educational homophily in the agent’s network.

Equation 14.4 shows that the social influence si_i that an agent i of age x_i receives from her network is based on the distance between the proportion of members in her network who are already mothers ρ_i , and the average proportion of mothers of age $x = x_i$ in all networks in the previous generation (10 years before) ρ_x^* . This means that the degree of influence on the agent to have a child at any given age is based on how common/acceptable it is to have a child at that age relative to how common/acceptable it was to have a child at that age 10 years ago. The reference to the past allows agents to know not only which behavior is accepted/expected but also how behavior is *changing*, and to follow innovative behavior as long as her (their) friends are adopting it:

$$si_i(x_i; \gamma, \kappa) = \begin{cases} \frac{1}{1 + \exp(-\kappa_1(x_i - \frac{\gamma_1}{1 + \rho_i - \rho_x^*}))} + 1, & \text{if } \rho_i - \rho_x^* > 0.05 \\ 1, & \text{if } |\rho_i - \rho_x^*| \leq 0.05 \\ \frac{1}{1 + \exp(-\kappa_2(x_i - \frac{\gamma_2}{1 + \rho_i - \rho_x^*}))}, & \text{if } \rho_i - \rho_x^* < -0.05, \end{cases} \tag{14.4}$$

where $\gamma = (\gamma_1, \gamma_2)$ and $\kappa = (\kappa_1, \kappa_2)$.

For the transition to marriage/cohabitation, the social influence function works in a similar way, although instead of considering the proportion of mothers in




Fig. 14.11 Social influence by age

the network we use the proportion of members of the network who were already married/cohabiting π_i , and the average proportion of agents of age $x = x_i$ who were already married/cohabiting in all networks in the previous generation π_x^* . In both cases the strength of the effect is given by parameters $\gamma = (\gamma_1, \gamma_2)$ and $\kappa = (\kappa_1, \kappa_2)$.

When the difference between ρ_x^* and ρ_i or π_x^* and π_i exceeds the 5 % threshold, the social influence multiplier si_i either augments ($\rho_x^* < \rho_i, \pi_x^* < \pi_i$) or reduces ($\rho_x^* > \rho_i, \pi_x^* > \pi_i$) the original rates $u(t, x, edu)$ and $f(t, x, edu)$. If the absolute difference is less than 5 %, the social influence has no effect. Figure 14.11 illustrates the effect for a set of γ and κ values that offer a good model fit (see Sect. 14.7.1).

We assume that agents’ decisions become progressively independent of the behavior of their friends as they approach the upper limit of the family formation period, and that other influences become stronger. Hence, we model a decreasing negative effect of social interaction as age increases. Conversely, we assume that the positive effect will increase with age as it joins other influences (proximity to the biological limit, family influences) in pushing forward the transition to parenthood.

The parametrization of these effects is not an easy task given the lack of previous empirical references. The values we present here were obtained after several exercises with the calibration of the model. For example, given these parameters the probability of marriage/childbearing for an agent who is 30 years old is about 90 % of the original probabilities if there is one fewer marriage/mother in her network (assuming the networks have an average size of 20) than in the average of all networks of agents of that age 10 years before ($si_i = 0.9$). By contrast, if there is one additional marriage/mother in her network, her probability of marriage/childbearing increases by about 40 % ($si_i = 1.4$).

14.5.7 Model Equations

Following the explanations introduced in Sect. 14.4, the unemployment and the social influence effects are captured by the individual multipliers um_i and si_i to obtain for each agent i the individual updated probabilities (intentions) of union formation $u_i^*(t)$ and of having the first child within marriage $f_i^*(t)$:

$$u_i^*(t) = u(t, x_i, edu_i) \times um_i(x_i; \alpha) \times si_i(x_i; \gamma, \kappa), \tag{14.5}$$

and

$$f_i^*(t) = f(t, x_i, edu_i) \times um_i(x_i; \alpha) \times si_i(x_i; \gamma, \kappa). \tag{14.6}$$

As was mentioned above, the fertility probabilities of women who have children outside of marriage or cohabitation are not affected by unemployment and the local social influence, as given the cultural framework in Spain we assume that this type of event is frequently unplanned, and is therefore less conditioned by labor market or peer behavior considerations.

14.5.8 Micro-Macro Loop

The effects of unemployment and social influence at the micro level result in a series of decisions about marriage and childbearing which modify the ideal ages at these events at the macro level, and hence the intentions at the micro level of subsequent cohorts. This process starts in 1975, once the initialization process (1944–1974) is complete. At $t = 1975$, $u(t, x, edu) = u^o(x, edu)$, $f(t, x, edu) = f^o(x, edu)$, and $fs(t, x, edu) = fs^o(x, edu)$, where u^o , f^o , and fs^o are the probability transitions obtained from the multistate model discussed in Sect. 14.5.1 and Fig. 14.7.

Equation 14.7 describes how this process of reciprocal dependence works:

$$u(t + 1, x, edu) = \begin{cases} \frac{\overline{u^*}(t, x, edu)}{\max(\frac{\overline{u^*}(t, x, edu)}{\overline{u}(t, x, edu)}, \theta)}, & \text{if } \frac{\overline{u^*}(t, x, edu)}{\overline{u}(t, x, edu)} < 1 \\ \overline{u^*}(t, x, edu), & \text{otherwise,} \end{cases} \tag{14.7}$$

where $u(t + 1, x, edu)$ represents the baseline intentions at time $t + 1$ by age and educational level, $\overline{u^*}(t, x, edu)$ is an *average* by age and educational level of the updated probabilities defined in Eq. 14.5, and $\overline{u}(t, x, edu)$ is an *average* by age and educational level of the original probabilities (without effects) at time t . $\overline{u^*}(t, x, edu)$ recovers the effects of unemployment and social influence on the original intentions of the previous year. An analogous mechanism applies for the computation of the transition probabilities to parenthood inside unions $f(t + 1, x, edu)$.




Fig. 14.12 Desired mean age thresholds by education level

The parameter that governs the loop between intentions and behavior is θ , which depends on the observed MAFB and the educational level, as shown in Fig. 14.12. It controls the position of the threshold after which agents are no longer willing to postpone family formation. When θ is closer to one then all of the effects of unemployment and social influence are recovered in the new baseline, and the desired and the observed mean ages grow at a similar speed. As intentions near the ceiling imposed by people's preferences, then θ approaches zero, and the new baseline does not consider a portion or all of last year's effects; people thus resist further increases in the ages at marriage and childbirth. A similar resistance is observed at the beginning of the postponement process as information about the transformation of the age at family formation reaches everybody in the population.

In Sect. 14.7.2 we show different scenarios resulting from a series of assumptions about the dynamics between the observed and the desired timing of family transitions.

14.6 Data and Tools

Our model could be described as a semi-artificial population model, a particular type of agent-based model (ABM) which, according to Bijak et al. (2013), results from the introduction of ABM techniques into a predominantly empirical discipline like demography. Semi-artificial population models are characterized by the combination of empirical and simulated data.

For the computation of the original union and fertility rates, we use the 1991 Sociodemographic Survey (INE 2015). It provides a large representative sample of the Spanish population (age > 10), with 159,154 observations.

To obtain the initial age structure of the population we use the 1940 census. For the initial distribution of the population by education, we use information from the 1970 census, which is the first to present disaggregated population figures. Both censuses are from the Spanish Statistical Office (INE 2015).

The age- and sex-specific mortality rates from 1944 to 2014 were obtained from the Human Mortality Database (HMD 2015).

For the reconstruction of the long unemployment series, we made use of various sources. The numbers of people who were registered as unemployed came from the Statistical Yearbooks of Spain published by the Spanish Statistical Office (INE 2015). We also used an interpolation of the censuses from 1930 to 2011 to obtain the number of working-age individuals. The Spanish Labor Force Survey (EPA) series 1960–1978 came from Carreras and Tafunell (2006). Finally, the series for the period 1979–2014 came from the Spanish Statistical Office (INE 2015).

Simulations were run in R (R Core Team 2015) and NetLogo (Wilensky 1999), using the `RNetLogo` extension (Thiele et al. 2012). To obtain the estimates from the multistate model we used the `survival` (Therneau 2015) and the `mvna` (Allignol et al. 2008) R packages. The code is optimized to take advantage of parallel computing using the `snowfall` R package (Knaus 2013).

Running four parallel simulations with an initial population of 3,000 agents each already produces useful results. The results reported here, however, were obtained with initial populations of around 30,000 agents for smoother trends.

14.7 Results

We begin this section by presenting the fit to our target of the model, as described above. In a second step, we present scenarios for different assumptions of some of the key mechanisms. Finally, we try to assess the individual role of each of the components of the model by presenting simulation results in which we alternately omit each of these effects.³

14.7.1 *Original Model*

Figure 14.13 shows the observed and the simulated MAFB. This fit corresponds to the non-linear specification of the model presented in Eqs. 14.5 and 14.6. Non-

³All these simulations were carried out with the following values of the global parameters described in Table 14.2: recurrence in unemployment $\sigma = 70\%$, maximum network size $\nu = 20$, and fertility rate for parity $1+f_2 = 0.15$.




Fig. 14.13 Observed vs simulated MAFB Spain, 1975–2013




Fig. 14.14 Observed (*left*) vs simulated (*right*) ASFR over time. Spain 1975–2012

linearity here refers to the fact that the results the model provide cannot be predicted by adding up each of its individual components.

The parameters for the unemployment and social influence multiplier correspond to the values shown in Figs. 14.10 and 14.11. This specification closely reproduces the observed trend.

Although the previous figure gives us an idea of how well the model approximates the data, the real challenge lies in matching the evolution of the distribution of ASFRs. As shown in Fig. 14.14, the model also reproduces this trend relatively well, especially the resulting distribution and the shift in the peak from around age 23 to around age 30.




Fig. 14.15 Observed vs simulated MAFB. Results from linear models. Spain, 1975–2012. Norms = Behavior (*left*) and Norms < Behavior (*right*)

14.7.2 Alternative Models

As described in Sect. 14.5.8, the bridge between the individual and the aggregate level in our model is provided by the influence that age norms regarding marriage and childbearing exert on the agents. Hence, the pace and the shape of fertility change will depend greatly on how these social norms shift in response to changes in behavior.

In our original model, we assume that people’s preferences are initially resistant to change. As mentioned before, this resistance is attributable to both the inertia of cultural norms and the time it takes for information regarding new material incentives to reach and be processed by individual agents. After the change in socioeconomic incentives picks up speed and individuals adapt their expectations, the age norms start to catch up with behaviors, but only until they reach the threshold created by people’s beliefs about the upper limit of the ideal age range for marrying or having children.

In this section we compare our original model with two other linear specifications in which we remove the thresholds and the effect of norms is constant over time.

Figure 14.15 shows two different scenarios. In the left graph, age norms evolve at the same speed as behaviors. Each year individuals adjust their expectations by taking into account all of the information generated in the previous year regarding the material incentives and obstacles to marrying and reproducing, as well as the behavior of their peers ($\theta = 1$). In this case norms offer no resistance. On the other hand, the right graph presents a scenario in which norms are highly resistant to change, and individual preferences defy the most immediate changes in the socioeconomic incentives for marriage and childbearing ($\theta = 0$).

These scenarios result in significant overestimation and underestimation, respectively, of the postponement process. It seems reasonable to assume that during the




Fig. 14.16 Observed vs simulated MAFB. Results from non-linear model with omitted effects. Spain, 1975–2012. Model with education only

postponement process age norms adapt in relation to changes in material conditions; but it also seems clear that this adaptation does not simply mirror the changes at the structural level.

14.7.3 Net Effects

In this section we present the results of a series of exercises in which we try to isolate the different effects considered in our model. Figure 14.16 shows the results of a specification of the model in which we only consider the effect of educational expansion, as expressed in Eq. 14.8. The only driver in this case is the compositional change of the population by education level:

$$u_i^*(t) = u(t, x_i, edu_i) \text{ and } f_i^*(t) = f(t, x_i, edu_i). \tag{14.8}$$

Figure 14.17 presents a scenario in which the effect of unemployment is omitted (see Eq. 14.9), and one in which the effects of social influence are omitted (see Eq. 14.10):

$$u_i^*(t) = u(t, x_i, edu_i) \times si_i(x_i; \gamma, \kappa) \text{ and } f_i^*(t) = f(t, x_i, edu_i) \times si_i(x_i; \gamma, \kappa), \tag{14.9}$$

and

$$u_i^*(t) = u(t, x_i, edu_i) \times um_i(x_i; \alpha) \text{ and } f_i^*(t) = f(t, x_i, edu_i) \times um_i(x_i; \alpha). \tag{14.10}$$




Fig. 14.17 Observed vs simulated MAFB. Results from non-linear models with omitted effects. Spain, 1975–2012. Model without unemployment effect (*left*) and model without social influence effect (*right*)

The effect of education (the one that results exclusively from the compositional change in the population by education level) is substantially lower than we had originally expected, accounting for less than a fifth of the total increase in the MAFB by the end of the period. Although this modest effect of education questions the consensus regarding the causes of postponement, the result is not unexpected if we look at the relatively narrow gap between educational levels in the original distributions of age-specific fertility probabilities we use in the model (see Sect. 14.5.1). However, it is important to keep in mind that the result does not refer to the total effect that education exerts on the MAFB but to its most direct mechanism.

Regarding the multiplier effect of social influence (left panel of Fig. 14.17), the first element worth noting is its large contribution to the process, even though in this case its amplifying effect refers only to the delays caused by educational expansion. It's also interesting to note that in a world without unemployment the trend does not show the more recent increase which is likely associated with the latest Spanish economic crisis which resulted in a steep increase of unemployment rates.

The second model, which only accounts for the compositional change with respect to education plus the effect of unemployment (right panel of Fig. 14.17), results in a relatively small increase in the MAFB. This difference is attributable in part to the fact that while unemployment exclusively affects the proportion of people unemployed in that year (20–25 % in years of high unemployment), the effect of social influence spreads through the networks, reaching most, if not all agents in the population.

The other interesting element is that the difference between the two is amplified by the fact that the model with social influence pushes the MAFB enough to overcome the initial resistance provided by social norms, as defined in Eq. 14.7, while the model without social influence doesn't. This helps to illustrate the relative

futility of trying to isolate effects, and supports our assumption that the pace of the process cannot be explained solely by the measure of the strength of each of its individual components.

14.8 Conclusions

Using an agent-based model we showed how the postponement of fertility in Spain can be explained by a set of relatively simple mechanisms at the individual level: rising opportunity costs of childbearing associated with higher education, growing economic uncertainty from an increasingly unstable labor market, and the multiplier effect of social interaction.

While the direct effect associated with educational expansion was substantially lower than we expected, our results showed that the feedback effects from social interaction were far more relevant than we had originally imagined. The extent to which networks amplify and spread an original force was particularly apparent when we compared the influence of social interaction with other factors that affect only a subset of the population, like unemployment. The evidence we present here provides support for the claim that the echo generated by social interaction can exceed the impact of the original effects that triggered it. Nevertheless, the existence of non-linear dynamics makes it impossible to describe the postponement process as a simple aggregation of each of its individual components, or to assess precisely the contributions of each of these components. But by using an ABM we were able to go beyond the assessment of the presence/absence of effects and explore the question of exactly how some of these mechanisms push forward the decision to form a family.

We found, for example, that for the shift of the peak of the distribution of age-specific fertility rates the increase in the *positive* influence on marrying/having children at later ages is as important as the increase in the *negative* influence at younger ages.

The other key element in the explanation of the postponement process is the micro-macro feedback loop through which past behaviors trigger normative changes, and which in turn translate into updated expectations for succeeding cohorts of men and women.

We tested different hypotheses and confirmed that the assumption of an equilibrium between preferences and behaviors leads to simulation results which deviate substantially from the observed trends. The key dynamic here seems to be the existence of a threshold after which people are reluctant to accept further delays in the age at which they start having children. Thus, it appears that norms are not so much *converging* or *lagging behind* as they are encouraging *resistance* to structural changes.

Understanding how age norms change is therefore essential to understanding the pace and the extent of fertility change. The formation of people's expectations and preferences is shaped by a larger set of elements than those we explored here. But while we did not directly address all of these factors, we do not intend to treat norms as black boxes. The strength with which people resist further increases in the

timing of family formation depends on their expectations regarding the conditions for childbearing. In addition to their perceptions of the present and future dynamics of the labor market or the educational system, these conditions include elements such as the availability of affordable childcare or the existence of support systems that facilitate work-family balance.

On the methodological side, we tried to provide another illustration of how computational models present a great opportunity to add complexity and dynamism to our representations of human behavior.

In this chapter we have provided only a very general description of the postponement process, leaving many potentially interesting dynamics to be explored in future work. For example, researchers may want to investigate the influence of the size, the composition, and the level of homophily of the social networks; or the changes in the effects of economic uncertainty as the number of dual-earner households increases. Moreover, modeling the changing role of women with regard to paid and unpaid work would undoubtedly shed more light on our conclusions. The classic schema of the innovators versus the followers of demographic change could also be tested by modeling different thresholds and ceilings (resistance and limits of normative change) for different subgroups of the population.

Finally, as we consider future scenarios, we believe that some of the forces that have been pushing the timing of family formation will continue moving in the same direction, at least in the medium term. There is room for further educational expansion in Spain, and a high degree of economic uncertainty is likely to be a feature of people's lives in the near future. These trends may be met with some resistance, but, as we noted above, there is no guarantee that behaviors will naturally converge with preferences. Thus, a large gap between behaviors and preferences could become a permanent feature in the coming years. The action taken to reduce this gap—or the failure to take action—will certainly shape the ongoing development of the timing of family formation in Spain.

References

- Adsera, A. (2004). Changing fertility rates in developed countries. The impact of labor market institutions. *Journal of Population Economics*, 17(1), 17–43.
- Ahn, N., & Mira, P. (2001). Job bust, baby bust?: Evidence from Spain. *Journal of Population Economics*, 14(3), 505–521.
- Ajzen, I. (1985). *From intentions to actions: A theory of planned behavior*. Berlin/Heidelberg: Springer.
- Allignol, A., Beyersmann, J., & Schumacher, M. (2008). mvna: An R package for the Nelson-Aalen estimator in multistate models. *R News*, 8(2), 48–50.
- Aparicio Diaz, B., Fent, T., Prskawetz, A., & Bernardi, L. (2011). Transition to parenthood: The role of social interaction and endogenous networks. *Demography*, 48(2), 559–579.
- Baizán, P. (2006). El efecto del empleo, el paro y los contratos temporales en la baja fecundidad española de los años 1990. *Revista Española de Investigaciones Sociológicas*, 115(1), 223–253.
- Balbo, N., & Barban, N. (2014). Does fertility behavior spread among friends? *American Sociological Review*, 79(3), 412–431.

- Balbo, N., & Mills, M. (2011). The influence of the family network on the realisation of fertility intentions. *Vienna Yearbook of Population Research*, 9, 179–205.
- Becker, G. S. (1981). *A treatise on the family*. Cambridge: Harvard University Press.
- Beyersmann, J., Allignol, A., & Schumacher, M. (2011). *Competing risks and multistate models with R*. New York: Springer
- Bijak, J., Hilton, J., Silverman, E., & Cao, V. D. (2013). Reforging the wedding ring: Exploring a semi-artificial model of population for the United Kingdom with Gaussian process emulators. *Demographic Research*, 29, 729–766.
- Billari, F. C., Goisis, A., Liefbroer, A. C., Settersten, R. A., Aassve, A., Hagestad, G., & Spéder, Z. (2011). Social age deadlines for the childbearing of women and men. *Human Reproduction*, 26(3), 616–622.
- Bongaarts, J. (2001). Fertility and reproductive preferences in post-transitional societies. *Population and Development Review*, 27, 260–281.
- Bourdieu, P., & Wacquant, L. J. (1992). *An invitation to reflexive sociology*. Chicago: University of Chicago press.
- Carreras, A., & Tafunell, X. (Eds.) (2006). *Estadísticas históricas de España: Siglos XIX–XX* (2nd ed.). Bilbao: Fundación BBVA.
- Casterline, J. B. (2001). Diffusion processes and fertility transition: Introduction. In J. B. Casterline (Ed.), *Diffusion processes and fertility transition: Selected perspectives* (Chapter 1, pp. 1–38). Washington, DC: National Academy Press.
- Castro-Martín, T., & Martín-García, T. (2013). The fertility gap in Spain: Late parenthood, few children and unfulfilled reproductive desires. In G. Esping-Andersen (Ed.), *The fertility gap in Europe: Singularities of the Spanish case* (Number 36 in Social studies collection, Chapter 2, pp. 45–81). Barcelona: “la Caixa” Welfare Projects.
- Coale, A. J. (1973). The demographic transition reconsidered. In *Proceedings of the International Population Conference 1973*, Liège: Ordina Editions.
- Entwisle, B., Rindfuss, R. R., Guilkey, D. K., Chamratrithirong, A., Curran, S. R., & Sawangdee, Y. (1996). Community and contraceptive choice in rural Thailand: A case study of Nang Rong. *Demography*, 33(1), 1–11.
- Esping-Andersen, G., & Billari, F. C. (2015). Re-theorizing family demographics. *Population and Development Review*, 41(1), 1–31.
- Giddens, A. (1984). *The constitution of society: Outline of the theory of structuration*. Berkeley: University of California Press.
- Goldstein, J., Lutz, W., & Testa, M. R. (2003). The emergence of sub-replacement family size ideals in Europe. *Population Research and Policy Review*, 22(5–6), 479–496.
- HMD (2015). Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de. (Data downloaded in May 2015).
- INE (2015). Instituto Nacional de Estadística (Spanish Statistical Office). <http://www.ine.es/>. (Data downloaded in March–June 2015).
- Knaus, J. (2013). *snowfall: Easier cluster computing (based on snow)*. R package version 1.84–6.
- Kohler, H.-P. (1997). Learning in social networks and contraceptive choice. *Demography*, 34(3), 369–383.
- Kohler, H.-P., Behrman, J. R., & Watkins, S. C. (2001). The density of social networks and fertility decisions: Evidence from South Nyanza District, Kenya. *Demography*, 38(1), 43–58.
- Kohler, H.-P., Billari, F. C., & Ortega, J. A. (2002). The emergence of lowest-low fertility in Europe during the 1990s. *Population and Development Review*, 28(4), 641–680.
- Kravdal, Ø. (2002). The impact of individual and aggregate unemployment on fertility in Norway. *Demographic Research*, 6, 263–294.
- Kreyenfeld, M. (2005). Economic uncertainty and fertility postponement. Evidence from German panel data. *MPIDR Working Paper WP 2005-034*.
- Kreyenfeld, M. (2010). Uncertainties in female employment careers and the postponement of parenthood in Germany. *European Sociological Review*, 26(3), 351–366.

- Lesthaeghe, R., & Van de Kaa, D. J. (1986). Twee demografische Transitie's? In R. Lesthaeghe & D. J. Van De Kaa (Eds.), *Bevolking: Groei en Krimp* (Mens en Maatschappij, pp. 9–24). Deventer: Van Loghum-Slaterus.
- Lyngstad, T. H., & Prskawetz, A. (2010). Do siblings' fertility decisions influence each other? *Demography*, 47(4), 923–934.
- Marini, M. M. (1984). Women's educational attainment and the timing of entry into parenthood. *American Sociological Review*, 49(4), 491–511.
- Mathews, P., & Sear, R. (2013). Does the kin orientation of a British woman's social network influence her entry into motherhood? *Demographic Research*, 28, 313–340.
- McDonald, P. (2000). Gender equity in theories of fertility transition. *Population and Development Review*, 26(3), 427–439.
- Montgomery, M. R., & Casterline, J. B. (1993). The diffusion of fertility control in Taiwan: Evidence from pooled cross-section time-series models. *Population Studies*, 47(3), 457–479.
- Morgan, S. P. (2001). *Should fertility intentions inform fertility forecasts?* US Census Bureau Conference "The Directions of Fertility in the United States", Alexandria, 2–3 Oct 2001.
- Morgan, S. P., & Bachrach, C. A. (2011). Is the Theory of Planned Behaviour an appropriate model for human fertility? *Vienna Yearbook of Population Research*, 9, 11–18.
- Munshi, K., & Myaux, J. (2006). Social norms and the fertility transition. *Journal of Development Economics*, 80(1), 1–38.
- Mynarska, M. (2010). Deadline for parenthood: Fertility postponement and age norms in Poland. *European Journal of Population/Revue européenne de démographie*, 26(3), 351–373.
- Myrskylä, M., Kohler, H.-P., & Billari, F. C. (2009). Advances in development reverse fertility declines. *Nature*, 460(7256), 741–743.
- Neels, K., Murphy, M., Ní Bhrolcháin, M., & Beaujouan, É. (2014). Further estimates of the contribution of rising educational participation to fertility postponement: A model-based decomposition for the UK, France and Belgium. Paper presented at the Population Association of America 2014 Annual Meeting, Boston, May 1–3.
- Ní Bhrolcháin, M., & Beaujouan, É. (2012). Fertility postponement is largely due to rising educational enrolment. *Population Studies*, 66(3), 311–327.
- Notestein, F. W. (1953). *Economic problems of population change*. London: Oxford University Press.
- Ozcan, B., Mayer, K. U., & Luedicke, J. (2010). The impact of unemployment on the transition to parenthood. *Demographic Research* 23, 807–846.
- Pailhé, A., & Solaz, A. (2012). The influence of employment uncertainty on childbearing in France: A tempo or quantum effect? *Demographic Research*, 26, 1–40.
- Perelli-Harris, B. (2005). The path to lowest-low fertility in Ukraine. *Population Studies*, 59(1), 55–70.
- R Core Team (2015). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing.
- Rendall, M., Aracil, E., Bagavos, C., Couet, C., DeRose, A., DiGiulio, P., Lappegard, T., Robert-Bobée, I., Rønsen, M., Smallwood, S., et al. (2010). Increasingly heterogeneous ages at first birth by education in Southern European and Anglo-American family-policy regimes: A seven-country comparison by birth cohort. *Population Studies*, 64(3), 209–227.
- Rindfuss, R. R., Bumpass, L., & St. John, C. (1980). Education and fertility: Implications for the roles women occupy. *American Sociological Review*, 45(3), 431–447.
- Rindfuss, R. R., Choe, M. K., Bumpass, L. L., & Tsuya, N. O. (2004). Social networks and family change in Japan. *American Sociological Review*, 69(6), 838–861.
- Rindfuss, R. R., Morgan, S. P., & Offutt, K. (1996). Education and the changing age pattern of American fertility: 1963–1989. *Demography*, 33(3), 277–290.
- Rosero-Bixby, L., & Casterline, J. B. (1993). Modelling diffusion effects in fertility transition. *Population Studies*, 47(1), 147–167.
- Settersten, R. A., & Hägestad, G. O. (1996). What's the latest? cultural age deadlines for family transitions. *The Gerontologist*, 36(2), 178–188.

- Sniehotta, F. F., Presseau, J., & Araújo-Soares, V. (2014). Time to retire the theory of planned behaviour. *Health Psychology Review*, 8(1), 1–7.
- Testa, M. R. (2006). Childbearing preferences and family issues in Europe. Number 253 in Special Eurobarometer. European Commission.
- Testa, M. R., Sobotka, T., & Morgan, S. P. (2011). Reproductive decision-making: Towards improved theoretical, methodological and empirical approaches. *Vienna Yearbook of Population Research*, 9, 1–9.
- Therneau, T. M. (2015). *A Package for Survival Analysis in S*. version 2.38. <http://CRAN.R-project.org/package=survival>
- Thiele, J. C., Kurth, W., & Grimm, V. (2012). RNETLOGO: An R package for running and exploring individual-based models implemented in NETLOGO. *Methods in Ecology and Evolution*, 3(3), 480–483.
- Van Bavel, J., & Nitsche, N. (2013). ‘The proper age for parenthood’ and second birth rates in Europe. *European Sociological Review*, 29(6), 1149–1161.
- Vignoli, D., Drefahl, S., & De Santis, G. (2012). Whose job instability affects the likelihood of becoming a parent in Italy? A tale of two partners. *Demographic Research*, 26, 41–62.
- Wilensky, U. (1999). *NetLogo*. Evanston: Center for Connected Learning and Computer-Based Modeling, Northwestern University.

Part V
Health, Mortality, and Support in Old Age

Chapter 15

Linking Income, Transfers, and Social Support in an Agent-Based Family Exchange Model

Fanny A. Kluge and Tobias C. Vogt

15.1 Background

While there is clear evidence that wealthier individuals tend to live longer, there is no consensus among scholars about why this is the case. Health economists have argued that wealth is a proxy for health itself, as healthier individuals have the capacity to earn higher incomes over their life course (Grossman 1972). Meanwhile, sociologists have found that income is strongly associated with education, and have argued that better educated individuals live longer and generate more income (Kroh et al. 2012; Deaton and Paxson 2001; Deaton 2003). Income may also promote survival because of the amenities money can buy. Wealth has been shown to be an important determinant of mortality (Preston 1975). Wealthier individuals can afford health-relevant goods and services, like better housing and nutrition, and out-of-pocket payments for certain health services.

In our analysis, we test the assumption that income affects survival both directly through the level of pension benefits, and indirectly through intergenerational transfers. The use of increased pension income to make intra-familial transfers represents an indirect and mediating link between income and improved old-age survival. Transfers play a major role in explanations for the rise in the average life span of humans and other species (Carey and Judge 2001; Lee and Chu 2012). It appears that transfers have contributed not only to the survival chances of offspring, but also to improved longevity (Gurven et al. 2012). Thus, raising the income of the elderly could lead to increases in the amounts transferred. These transfers may contribute to old-age survival by encouraging adult children to provide their parents with social support. Studies have shown that social isolation is a reliable

F.A. Kluge (✉) • T.C. Vogt
Max Planck Institute for Demographic Research, Rostock, Germany
e-mail: kluge@demogr.mpg.de; vogt@demogr.mpg.de

predictor of old-age health and mortality (Stephoe et al. 2013; Berkman 1995). Berkman (1995) emphasised that emotional support and social connectivity are more important determinants of old-age survival than practical help. For example, Ellwardt et al. (2013) showed that emotional support can have a protective effect on cognitive decline, whereas instrumental support can have none or even negative effects.

Negative associations have been found not only for all-cause mortality, but also for several health impairments and different causes of death (Stringhini et al. 2012; Eng et al. 2002). Holt-Lunstad et al. (2010) showed in a meta-analysis that the effects on the risk of death among older people of a lack of social relationships are comparable to the effects of smoking and excess alcohol consumption. Emotional and instrumental support from family members has been found to be protective against depressive symptoms and an increased risk of dying (Zunzunegui et al. 2001). However, it appears that not all forms of intergenerational support are beneficial for the mental and physical health of older people (Seeman 1996). It is far more important that the intergenerational relationships are motivated by the notion of exchange. The absence of reciprocity in the relationship between older parents and their adult offspring seems to change the character of the exchange from being helpful to being strenuous (Moorman and Stokes 2014).

The ability to comply with the societal norm of reciprocity depends on each generation's capacity to provide resources. Older parents provide their children with financial support, and children provide their parents with emotional or functional support. This pattern is typical of Western welfare states with a system of generous public pensions, whereas in East Asian countries intergenerational family transfers are primary sources of income at younger and at older ages (Lee and Mason 2011). In this context, an increase in income might enable an older family member to start making transfers, or to increase their transfer amounts. An increase in the transfer amounts may intensify the expectation that the adult children will reciprocate by providing support. The return transfers are not necessarily equivalent, and they may not occur at the same point in time. Children may also give different levels of material, functional, or emotional support over their life course depending on their ability to provide transfers. Likewise, children may decide to increase their support of their older parents because they expect to receive larger transfers in return. In either case, both generations can benefit from the exchange of resources: the older generation benefit from feeling integrated into the family and important to their offspring, and adult children benefit from receiving subsidies in times of need and from contributing to family solidarity.

The discussion of the literature on social support presented here is not a complete review of the literature on social support. There are many different aspects to social support, and the mechanisms through which social support affects survival are not exclusive, and are sometimes ambiguous. Nevertheless, most studies have found that social connectivity is an important determinant of well-being, subjective health, and survival. In our analysis, we measure the level of support by the number of contact hours per week between adult children and their parents, as reported in the Time Use Surveys for Germany. This is a very broad concept, as it includes personal visits,

telephone calls, and instrumental support, such as help with bureaucratic issues or household tasks. While we cannot determine the quality of the contact, we assume that any time spent together is beneficial for survival.

In our model, income increases enable older individuals to buy health-relevant amenities and services, or to make transfers to their children. Children may respond to these transfers by providing more support, as measured in hours of contact per week. This support may be expected to help older individuals maintain their health status and survival chances. In the model, we assume that wealth can affect old-age survival also via factors like health behaviours, medical care, and environmental pollution. We apply our model to the natural experiment setting of post-reunification Germany, and use the structural break to investigate whether changes in income and transfer patterns among older eastern Germans are associated with improvements in health. The results of this investigation allow us to answer the question of whether increases in income improve health and survival through changes in the exchange relationship in the family. The chapter is structured as follows. First, we provide a short overview of the natural experiment character of the German reunification. Subsequently, the model is outlined in detail, followed by a description of the data used. Finally, we present our results and discuss their implications.

15.1.1 The German Reunification: A Natural Experiment

When the Berlin Wall fell in 1989, Germans in the east and west had lived for four decades under very different social, economic, and political conditions. The reunification of the country one year later introduced the western German political system to the east, and eastern Germans gained access to the generous social security system of the west. Living conditions in the east improved quickly, and soon approached western German standards. Life expectancy responded to this transformation, and quickly converged with the western German level. Improvements in mortality among the elderly accounted for nearly 90 % of this life expectancy convergence while younger age groups contributed “only” 10 % (Vogt 2013).

In discussions of the potential triggers of this rapid catch-up, a number of factors have been mentioned, such as changing health behaviours, improving medical infrastructure, and reduced environmental pollution. However, rising income levels have been shown to be the most important factor (Diehl 2008). Vogt and Vaupel (2013) found that mortality among pensioners whose incomes rose converged very quickly with mortality among their western German counterparts, and that pensioners with children benefited from these income increases even more than childless pensioners. Meanwhile, intergenerational monetary transfers increased (Kohli et al. 2002), and the amount of time children spent with their elderly parents increased (Time Use Surveys 1991 and 2001). Based on these findings, we will investigate the question of whether increased pensions had an impact on eastern German old-age mortality via intra-familial transfers.

15.1.2 An Individual's Behaviour in an Agent-Based Model

In order to investigate the association between increasing pension income and improved old-age survival on the individual level, we need to understand the decision-making process of an individual. In our analysis, the main question is whether the exchange relations in families (i.e., transfers of money and of care) contribute substantially to macro-level health and mortality outcomes. From an economic perspective, there are several motives for private transfers, ranging from pure altruism (Becker 1974), to accidental bequests (Yaari 1965), to the desire for an exchange (Cox 1987; Henretta et al. 1997; Norton and Van Houtven 2006; Koh and MacDonald 2006). Yet these motives fail to take into account the social embeddedness of an individual, and of his or her need to comply with certain norms and social expectations of behaviour. The theory of planned behaviour can account for these economic, rational choice motivations, but it adds a social aspect to explanations of an individual's decision-making process (Ajzen 1991).

Ajzen's theory of planned behaviour focuses on three aspects: an individual's views on a given behaviour, the social and subjective norms regarding the behaviour, and the level of perceived control. All three aspects are incorporated into our model.

An intention to act in a certain way is formed by an individual's attitude or expectation regarding the outcome of this behaviour. The individual may expect to derive certain benefits from his or her prospective behaviour. Children may expect to receive larger intergenerational transfers if they spend more time with their parents. In return for transfers, parents may expect to have more visits from or to spend more time with their children.

At the same time, individuals cannot be totally selfish in their actions. There are normative constraints that prevent parents from using all of the resources for themselves, and from failing to support their offspring. Likewise, social and individual normative beliefs about how children should behave in a family context make it difficult for the adult children to receive transfers without returning resources. This underlines the importance of family solidarity and reciprocity as drivers in an exchange relationship.

Finally, Ajzen's theory includes another important aspect of the intention-forming process: namely, the perception of having control over a certain behaviour. In this process, an individual must feel that he or she has the necessary means to perform an action. Despite having a general desire to make transfers and a belief that such transfers are socially acceptable, older parents may not be able to give money to their children if they lack the necessary resources. This aspect is key for our model. Pension income increases do not necessarily change the willingness of older people to make transfers, but the increases may alter their ability to provide support. We observe empirically that increases in pension income are accompanied by increases in private downward transfers. Thus, it appears that parents voluntarily increase the amounts they transfer to their children if they can afford to do so. That is why we assume that this link holds, and that having the necessary resources will lead to higher transfers.

15.2 Model Implementation

We develop an agent-based model that we initialize with realistic demographic and economic rates for eastern Germany. Although the potential time horizon for computation is much longer, we focus on the transition period between the years 1980 and 2000 to test our hypotheses. We use NetLogo (Wilensky 1999) to model the family relationships. RNetLogo is used to estimate the demographic rates for each individual depending on his or her age. The RNetLogo package offers an interface to embed the agent-based modelling platform NetLogo into the R environment (Thiele 2014).

The initial population size is set at 5,000 individuals. Each agent is assigned a numerical identifier, an age, a sex, a marital status, a number of children, and a partner identifier if he or she is married (the identifier is set to false if the agent is not married). The list of characteristics for each agent is given in Table 15.1. The model simulates eastern Germany's relevant population characteristics from 1952 to 2051. The initial population age distribution approximates that of eastern Germany in 1952. The individuals' sex is randomly chosen, with the sex ratio being 0.5. Each woman has an initial number of children based on her age. If a woman is under age 15, the number of children is zero. If she is between ages 15 and 25, the number of children is Poisson distributed with a mean of one; and if she is over age 25, the initial number of children is Poisson distributed with a mean of two. When a woman gets married, her partner takes over her family characteristics.

Figure 15.1 summarises the six subroutines of the NetLogo model in a detailed graphical way. The first part of the model refers to the pure demographic events. In each period individuals age by one year, and can die, form a union, have a child, or migrate. A woman can form her own household with a unique family identifier at age 15. Individuals age until they reach their age at death. This age is determined by the death rates from the Human Mortality Database (HMD) (2013, see details in Sect. 15.3) for men and women for the respective year and age. For people aged 60 and above, an additional adjustment factor is estimated from the model. This factor is based on their pension level, whether they have children, and a care indicator (the adjustment factor is explained later in this section where the respective subroutine is described in detail). If a child dies, the number of children is corrected downwards. If an individual receives the information that he or she has reached the age at death, the about-to-die indicator is set to one, and the agent dies at the beginning of the following period. Individuals are able to form a union when they reach age 15. They determine whether an appropriate mate exists: i.e., they seek out a partner who is not married, is of the opposite sex, is within 10 years of their age, and is not related. The mate receives the same family information, such as the number of children and the family identifier as the female. If one partner dies, the status of the remaining partner is set to widowed.

Children can only be born within marriage. Females between ages 15 and 49 receive their age-specific fertility rates for the respective year from the German Federal Statistical Office data for the past. After 2010, age-specific fertility rates

Table 15.1 List of agent's characteristics (Note: random numbers are drawn from a uniform distribution unless otherwise indicated)

Variable	Description
id	Numerical identifier for each individual
age	Age of individual
sex	Sex of individual, female = 1, male = 0
fr	Age-specific fertility rate depending on individual's age
dr f adj	Age-specific death rate, females, depending on values from the Human Mortality Database (HMD) (2013, see details in Sect. 15.3) and individual adjustment factors
dr m adj	Age-specific death rate, males, depending on values from the Human Mortality Database (HMD) (2013, See details in Sect. 15.3) and individual adjustment factors
married	True or false value depending on marital status
mate	Contains id number of mate
mother	Contains id number of mother to identify siblings
children	Reports actual number of children
birth	New birth in model, estimated with asfr between 1952 and 2050
birthinmodel	0 = no child in model, 1 = child born within model
famid	Family identifier (id of mother/wife)
hhinc	Reports pension income of household combined
transfers	Amount of pension income transferred to children
transfer increase	Transfers relative to pre-1990 levels
pension age	Pension entry age, random between 60 and 65
pension	Amount of monthly pension income
pension increase	Percentage increase in pensions relative to pre-1990 level
years contributed	Number of years contributed for pension estimation
childlost	Number of years deducted for women, 3 years for each child
occupation	Adjustment factor for different occupations, random number between 0.7 and 2
unemp	Years of unemployment, random number between 0 and 3
widowed	True or false variable depending on partner's death
support	Hours of support provided by children
availability	Availability = 1 if at least one child is available for care
support received	Support received by elderly, mean of all children
care	Hours of contact per week, depending on transfers
le adj	Life expectancy adjustment factor for the elderly, depending on pension level and support
about to die	Indicator set to 1 if age at death is reached

Source: own considerations

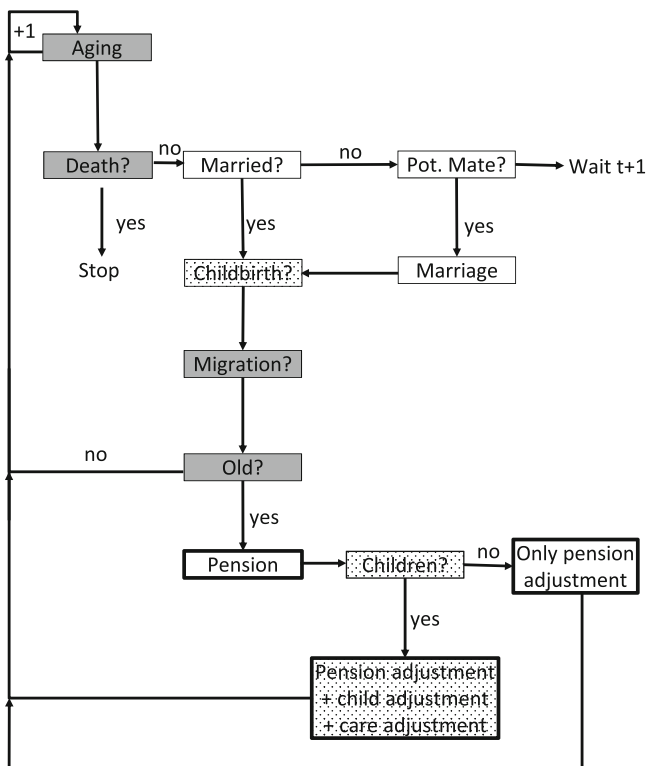


Fig. 15.1 The familial exchange model in a flow diagram (Note: *White* boxes refer to the mating and marriage market routine, *dotted* boxes to fertility and the existence of children, *grey* refers to all other demographic subroutines and *bold* frames to the pension and transfer routine)

are held constant for the remaining years until 2051. This should not be a problem, as our model focuses on the changes between 1980 and 2000 in eastern Germany, and we do not seek to predict future developments. If a woman gives birth, her number of children increases and the information is also sent to her mate.

Younger agents can move away from their family after age 18. We assume that moving stops after age 50. The moving procedure is simple and random, and the idea is to control for the availability of children who are able to provide care and who live in the neighbourhood. We estimate the distance between a parent and his or her child. If the parent and child live close together (within a radius of 0.75 times the maximum link length), we assume that the child is available to care for his or her elderly parent.

The remaining subroutines describe the economic parameters used in the model. We refrain from modelling a labour market for the young, as this would add a dimension to our model that is irrelevant for our purposes. Individuals have a randomised pension age between 60 and 65. Before 1990, retirees receive a pension of around 200 euros. This value corresponds to the per capita pension values for

the GDR, estimated within the National Transfer Accounts (NTA) (see details in Sect. 15.3) project for 1988. After 1990, the pension values increase significantly, which is in line with real data. The pension is estimated for the year in which each individual decides to retire. This estimate is made by multiplying the actual pension value of the corresponding year by the number of years each individual contributed, as this is the simplest way to calculate German pension benefits. To increase the variance in pension benefits, and thus to account for different income levels, we added an occupation adjustment that varies between 0.7 and 2. The number of years each individual contributes depends on his or her unemployment history (each individual can have up to 3 years of unemployment that are deducted). Because in the GDR periods of unemployment tended to be short, we have chosen the relatively low deduction level of up to 3 years. In addition, women lose 3 years of contributions for each child.

The elderly transfer a fraction of their pension to their descendants. The fraction is estimated as a yearly percentage of mean transfers to the next generation from the NTA project. Before 1990, the elderly were giving around 1 % of their pension income to their children or grandchildren, according to the micro-survey from 1988. Shortly after reunification, the share increased to around 3 %, and peaked in 2003 at a value of almost 6 %. The elderly shared their newly acquired wealth within the family. For the years in between, when no micro-survey and no NTA were available, the data are interpolated. We estimate the household income of each elderly couple and determine the amounts of the transfers they made to their children in the respective year. The children base the level of support they provide to their elderly parents on the size of the transfers they receive. The individual amount received is compared to the mean of transfers of the young. If the child receives more than the average amount, he or she intensifies his or her care efforts from 1 to 2 h per week. In a second step, we compare the transfers to the pre-1990 values. If transfers hardly increased, no additional hours are spent with the parents. If the transfer amounts are more than three times larger than the amounts transferred before 1990, children spend in total 3 h per week with their parents; similarly, if the transfer amounts are more than five times larger than the amounts transferred before 1990, children spend 5 h per week with their parents (the hours correspond to the values of the Time Use Survey, in which we find that on average children and parents spend about 3 h per week together). The amount of support as a mean of their own children in return determines the adjustment for care for the corresponding death rates of the elderly individuals.

The last adjustment to the death rates is estimated based on the income group of the individual. The pension level is determined relative to pre-1990 levels. If the pension increase is more than threefold, the survival of the elderly individual is enhanced.

Thus, based on his or her income group and family constellation, each individual receives an adjustment factor for his or her own death rates. Table 15.2 shows some examples for adjustment factors used in the model. These adjustment factors are estimated relative risks from the public pension insurance data.

Table 15.2 Adjustment factors depending on income level and parental status

Income level	Children	Adjustment factor
Above average	Yes	0.68
Above average	No	1.58
Average and below	Yes	0.79
Average and below	No	1.85

Source: own calculations based on public pension insurance data

We provide a brief example of a female agent born in 1955 in eastern Germany. She ages one year for each tick, and at every time step she receives the age-specific death rate for the respective year and age as baseline information. If the agent receives the information that she is about to die, she will leave the model. If she is alive, she can form her own household from age 15 onwards. If she is not married, the agent is searching for a potential partner. If she is married or has found a partner, she can give birth to a child with the probability of her age-specific fertility rate. Between the ages of 18 and 50, the individual can migrate, although this only affects her ability to care for her parents, and not her other family characteristics. When the agent reaches retirement age, her pension amount and her number of children are determined. Depending on her own pension wealth, her parental status, and the care efforts of her children, she receives an adjustment factor to her own death rates before she enters the next period.

We store the data for all of the individuals for each period between 1952 and 2051, and the most important variables – such as age, sex, family identifier, number of children, fertility, adjusted mortality, availability, and support – are used in the RNetLogo environment for further analysis.

15.3 Data

As our analysis focuses on the natural experiment setting of post-reunification Germany, we use in our model realistic economic and demographic rates from eastern Germany during the transition period. Information on survival is obtained from the Human Mortality Database (2013). We use age- and sex-specific death rates for the years 1952–2010 as a baseline for the respective male and female age groups.

For a more refined analysis of survival among individuals of pension age, we use the transition rates of individuals in the public pension insurance dataset, who differ in terms of their socio-economic status and their number of children. We explicitly use the scientific use file “Demografiedatensatz Rentenwegfall/-bestand 1993–2007” from the Forschungsdatenzentrum der Rentenversicherung (public pension insurance). This dataset covers more than 90 % of all eastern German pensioners, and is used to estimate the adjustment factors that reflect the individual’s income and number of children.

The age-specific fertility rates for eastern Germany between 1952 and 2010 are obtained from the Federal Statistical Office. The data up to 1989 are available from the special issue on population statistics 1946–1989 (Statistisches Bundesamt 1999). From 1990 onwards, age-specific fertility rates are covered in the publication on population and employment (Statistisches Bundesamt 2000).

The economic variables are provided by the National Transfer Accounts (NTA) database for eastern and western Germany (Kluge 2010).¹ Most importantly, we used pension and transfer data estimated for the years 1988, 1993, 2003, and 2008; separately for eastern and western Germany. The values for 1988 are based on data from one of the rare representative individual-level data sources that exist for the GDR era: namely, the income survey of blue-collar and white-collar worker households conducted in 1988. The survey, which covers 30,000 households with around 80,000 individuals, contains information on income levels, income composition, and changes in income. Important socio-economic variables, such as household size, gender, age, and highest level of education, are also included. The survey is representative for around 92 % of all individuals in the GDR in 1988 (Staatliche Zentralverwaltung fuer Statistik 1988).

The German Income and Expenditure Survey in 1993, 2003, and 2008 (Einkommens- und Verbrauchsstichprobe, or EVS) serves as the micro-foundation of the relative age shares for the later years. The EVS is based on a representative quota sample of Germany's private households, and provides information on income, consumption, transfers, savings, and assets. For example, the EVS for 2003 includes around 50,000 households, and is representative of households with a monthly net income of less than 18,000 euros. In addition to income and consumption, the EVS includes all of the relevant public transfers to households, and allows for estimations of private familial flows. To estimate inter-household transfers, we use the amounts given and received from other households in the respective year reported in the survey. Using this method, we are able to estimate the overall shares of transfers from the elderly to the successive generations for the respective year. For the years in between, no survey was available and the shares were imputed using linear interpolation.

In addition to demographic and economic variables, we used the Time Use Surveys 1991/1992 and 2001/2002 (latest available data in May 2015) to estimate the share of individuals who had frequent or occasional contact with family members. The scientific use files for Germany include around 5,000 households with approximately 12,000 individuals and 37,000 diary entries. Respondents fill in a calendar for three representative days. In the calendar they list all of their activities that take at least 10 min, and report with whom they are spending their time.

Depending on the individual characteristics of the model agent, the agent receives a set of demographic and economic parameters within each time step of the computation.

¹The theoretical framework of NTA builds upon Samuelson (1958), Diamond (1965), and Lee (1994). A detailed overview of the construction of NTAs is available at www.ntaccounts.org.

15.4 Results

In the baseline scenario, the model captures key variables of interest, such as the main demographic patterns and the pension amounts. Using relatively few parameters, we are able to replicate the major macro developments for the eastern German case study. We show a representative simulation result of the overall model that includes the main economic and demographic parameters. The results, which are estimated using an initial population size of 5,000 individuals, follow the eastern German rates of interest. To assess the relative importance of different factors, we estimate reduced versions of the model in alternative scenarios shown in the next section. We separately show the impact of varying pension and transfer levels on the mean ages at death.

15.4.1 Adequacy of the Model

First, we display a visualisation of the results for some key demographic parameters. We illustrate the model fit for a representative year of interest to show the detailed age patterns. As an example, Fig. 15.2 shows the real and estimated age profiles for eastern Germany in 1995. The average age of the population was lower before reunification than it was in subsequent years: the mean age increased from 38.1 in




Fig. 15.2 Age profiles eastern Germany 1995, real data and simulation results (Source: Federal Statistical Office, simulation results)




Fig. 15.3 Age-specific fertility rates, eastern Germany 1995, real data and simulation results (Source: Federal Statistical Office, simulation results)

1985 to 43.5 in 2005. Two major developments led to the ageing of the population. First, fertility declined sharply after 1989, with the total fertility rate reaching a low of 0.7. Thus, a generation of children was missing. Second, the number of older individuals who were reaching high and very high ages increased significantly. In particular, the age group 60–80 expanded. The initial age distribution used to forecast the longevity of individuals was based on the life expectancy in eastern Germany in 1952. The dips apparent in the graph reflect the real-world developments of two world wars and the Spanish Flu.

The use of age-specific fertility rates allows us to estimate the corresponding total fertility rate at the macro level. Our findings indicate that in 1985 there were around 38 women in each fertile age group (ages 15–49), and the average number of births was 1.77. In 1995, there were around 35 women in each fertile age group, and the average number of births fell to 0.9. The parameters subsequently recovered and reached the observed value of 1.3 in 2005. The real and estimated age-specific fertility rates for 1995 are shown in Fig. 15.3. The real-world pattern of postponing childbearing to higher ages is reflected in the estimates, which also show any fluctuations.

The remaining results deal with the relevant economic parameters; in our case, pensions and transfers. Figure 15.4 shows the time series of mean pension income for eastern Germany 1975–2008, and compares the model estimates with the observed data. Before reunification, most pensioners received slightly less than 200 euros a month. Thus, compared to the working-age population, pensioners were worse off. Within just 5 years after reunification the picture had changed




Fig. 15.4 Pension estimates (average), eastern Germany 1975–2008, real estimated data and simulation results (Source: Statistical Yearbooks GDR 1975–1990, FRG 1990–2008, simulation results)

dramatically, as the average pension income had more than quadrupled. Eastern German pensioners were integrated into the western German social security system, which rewarded them for their long work histories and high labour force participation rates. This holds also true for women as the east German female labour market participation was very high. The mean public pension benefit payment values were therefore even higher in the east than they were in the west. The simulated pension benefits are in line with real observations, including the significant increase observed after 1989. In its current form, however, the model is not suitable for analysing future developments. The eastern Germans who reached retirement age immediately after reunification were much better off than they would have been in the GDR, as their relatively long and consistent work histories were favourable for the calculation of benefits. However, younger eastern German cohorts have been experiencing longer periods of unemployment and frequent job shifts, and will thus receive lower pension benefits than the current generation of retirees (Simonson et al. 2012). It would therefore be necessary to include more realistic labour market parameters in the model.

Private transfers from the elderly to successive generations are depicted in Fig. 15.5. The results are reliable for the transition period of 1988–2008; before and after this period we simply keep the results constant. From the beginning of the 2000s–2008, the average share of pension income transferred to younger generations increased from around 1% to nearly 6%. This corresponds to around 50 euros of additional monthly income for young people. The years in which the




Fig. 15.5 Private transfers, as a percentage of pension income and total amount per month, eastern Germany 1952–2051, simulation results (Source: simulation results)

transfers were highest were also the years in which the pension values were highest. Thus, we can see that the elderly were sharing their newly acquired wealth with other family members.

15.4.2 *Different Scenarios on Changes in Income, Care, and Transfers*

In the remainder of this section, we show the simulation results for the mean age at death if key parameters of the model are changed. We estimate a disadvantageous baseline scenario and a scenario with realistic rates to show the bandwidth of life expectancy. We then separately show the impact of different pension and transfer levels on mean ages at death. Figure 15.6 displays the results of our first scenario. In the baseline estimates, including pensions and transfers, the real observed death rates for eastern Germany are used for computation. These results are based on the assumption that all of the covariates that were responsible for the significant decrease in death rates after 1990 are covered in the model. The simulation fluctuates around the real observed values: e.g., in 2010 these values were 73.5 for German men and 81.0 for German women. While it is clear that this model will not produce exactly the same results, the results are very similar. It therefore seems appropriate to use the model for further simulations.





Fig. 15.6 Mean age at death, baseline scenario, eastern Germany 1970–2020 (Source: simulation results with different adjustment factors)

The scenario without increasing pensions and transfers estimates the mean ages at death with the lowest possible survival benefit. This is done by adjusting the death rates to a level that is twice as high as the level in the baseline scenario, which corresponds to the results of the German pension data on the relative disadvantage of poor individuals without children. In this simulation, the mean age at death also increases, but at a slower pace. This accounts for the fact that the life expectancy convergence was triggered not only by income, but by other factors that we cannot quantify in our model, such as health care provision, environmental effects, or individual life-style factors.

Until 1990 the two scenarios are comparable, and hover around the same values. After reunification the adjustment factors are taken into account, and an advantage for the baseline scenario quickly emerges. Compared with the baseline scenario, we find that the mean age at death is around 3 years lower in 1990–2010. This could be the additional combined effect of higher pension income and transfers for the elderly. The results for the mean ages at death also include the ages at death of individuals under age 60. Although it is rarely the case in our simulation, the results are still affected by developments in the early 1990s, when mortality was exceptionally high among young men.

In scenario two (depicted in Fig. 15.7), the pension levels and mortality adjustment factors are varied after 1989. The scenario that is closest to reality is the one with the real observed pensions and the high mortality effects (dotted line). In this scenario we assume that individuals receive the full eastern German pension amounts in the 1990s and 2000s. Their pension levels over this period are compared

Fig. 15.7 Mean age at death, pension scenario, eastern Germany 1970–2020 (Source: simulation results with different adjustment factors and varying pension levels)



to their pension levels before 1990. If the post-1990 pension is less than three times the size of the pre-1990 pension, the mortality rate is adjusted by a factor of 1.3. If the post-1990 pension is between three and five times larger than the earlier pension (as was generally the case), the mortality rate is adjusted by a factor of 0.8. This ratio corresponds to the advantage wealthier pensioners have relative to pensioners who are less well off (based on estimates from official German pension data). If the pension increase is even larger, the adjustment factor is set to 0.4. The dashed line shows the results for the same pension levels, but unfavourable adjustment factors (2, 1, and 0.6). In the last scenario, the pension levels are cut to half of the real amount, and the unfavourable adjustment factors are used. This is the least beneficial scenario in terms of old-age survival. Between the highest and the lowest scenarios the average difference in the mean ages at death is 3 years. A comparison of the high and the low pension scenarios shows that the higher pension scenario has an average advantage of around one year. We find that the model is more responsive to changes in the adjustment factors than to variations in pension levels. The extent to which this is the case also depends on the pension brackets used to calculate the different adjustment factors, which could be changed in future analyses. The results confirm the findings of earlier studies of eastern Germany, which showed the importance of income and income differentials for survival (Gaudecker and Scholz 2007; Kibele et al. 2013).

The pension scenario was driven by the individual pension wealth of an agent. This direct link is important, because the literature cites income as a driver of survival. In the last scenario, we add the effect of intra-familial transfers to our analysis. The joint income of the elderly couple now determines the income level of the household and, in a second step, their transfers to their children. The children then base their care efforts on the transfer amounts they receive. This is a twofold

mechanism. First, the children compare the transfer amounts they are receiving to the amounts they received before 1990. Depending on the size of the increase, they provide 0, 3, or 5 h of support per week. Second, the children compare the transfer amounts they are receiving to the amounts other young individuals are receiving in the same time period. Based on the literature prior to reunification, we assume that the children have always provided some basic level of care. This basic level is set at 1 h per week, even if the children receive less than the average transfer amount. If they receive more than the average transfer amount, we assume that they double their care efforts. If the parents receive less than 1 h of care per week, the death rate adjustment factor is the same as that of a wealthier individual without children, or 1.5. If they receive between one and 3 h of care per week, the adjustment factor is 0.7. This ratio represents the advantage parents have relative to childless individuals. If they receive more than 3 h of additional care per week, their adjustment factor is a highly favourable 0.4.

The increase in *intervivo* transfers and expected bequests leads to an increase in some form of emotional or functional support. This reciprocity can be reproduced in our agent-based model setting. Figure 15.8 summarises the results for different transfer levels. The dotted line refers to the scenario without transfers. Here, all individuals (regardless of whether they have a family) have an adjustment factor of 1.5. During the transition years of 1990–2010, this converts to a disadvantage in the mean age at death of around 2 years. Between the two scenarios with high (solid line) and low (dashed line) transfers, we find differences in the short run only. In the scenario in which the transfers quadruple, the mean age at death increases much




Fig. 15.8 Mean age at death, transfer scenario, eastern Germany 1970–2020 (Source: simulation results with different transfer levels)

more quickly, but in the long run the mean age at death is not affected by the size of the observed transfers. The transfers simply accelerate the increase. According to the Time Use Survey 2001/02 for eastern Germany, adult children spend an average of 3.5 h per week in contact with their elderly parents. When the transfer amounts are higher, the average level of support provided per week is, at 4 h, slightly higher than the norm. When the transfer amounts are lower, the average level of support provided per week is, at 3 h, slightly lower than the norm.

15.5 Discussion

In this chapter, we tested a theoretical approach for examining the links between income, familial networks, and survival. After reviewing the literature, we identified three parameters relevant for our study: namely, pension income and familial reciprocal exchanges in the form of private transfers and care. The interdependencies of these parameters were tested using an agent-based model, which proved to be highly suitable for estimating the relationships. A major advantage of our study is that we were able to model familial relationships and the importance of networks on the micro level. This approach helped us discern realistic macro patterns. The observed relationships are especially relevant for elderly individuals. While the survival of younger individuals depends on different factors, the survival of the elderly is particularly dependent on giving and receiving transfers.

We first created a baseline scenario, which used the real observed death rates that lead to an average survival advantage of 3 years, and compared it to a scenario in which pensions and transfers are excluded. In our model, we find that higher pension levels are associated with an average survival advantage of one year. Transfers to children are an important driver of care, as they accelerate the increase in the mean age at death. In the pension scenario, the improvements in survival are mainly driven by the shares of richer and poorer individuals. The network scenario also depends on income levels, but the dynamic comes from the relative survival advantage of parents relative to childless individuals.

The illustrated mechanisms cannot fully account for the rapid convergence of eastern and western German death rates after reunification. It is important to keep in mind that health care quality, environmental, and life-style factors also played important roles. Still, this is an inductive example of how much survival can be affected through indirect pathways, such as social interactions. This approach is well suited for inductive studies. Case (2004) and colleagues found a similar mechanism in a case study in South Africa. This study indicated that the increased pension wealth of elderly coloured South Africans was shared across generations, and led to improvements in survival.

We theoretically showed the mechanism through which increases in public pensions could be converted into improved longevity. The data we used, which came from a wide range of sources, support this link. But because the data included in

our agent-based model did not come from a single source, the adjustment was not directly observable. This is both a strength and a weakness of our approach. In the future, the predictive power of the model could be improved through the inclusion of a realistic labour market parameter and more economic parameters for younger individuals.

Acknowledgements The authors wish to thank two anonymous reviewers for their valuable comments on an earlier version of the manuscript. We wish to thank the participants of the agent-based modelling winter course 2013/14 at the Max Planck Institute for Demographic Research for their support. In addition, we are thankful for valuable comments from the participants of the agent-based modelling workshop held in Leuven, Belgium in September 2014.

References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211.
- Becker, G. (1974). A theory of social interactions. *The Journal of Political Economy*, 82(6), 1063–1093.
- Berkman, L. F. (1995). The role of social relations in health promotion. *Psychosomatic Medicine*, 57(3), 245–254.
- Carey, J. R., & Judge, D. S. (2001). Life span extension in humans is self-reinforcing: A general theory of longevity. *Population and Development Review*, 27(3), 411–436.
- Case, A. (2004). Does money protect health status? Evidence from South African pensions. In David A. Wise (Ed.), *Perspectives on the economics of aging* (pp. 287–312). Chicago: University of Chicago Press.
- Cox, D. (1987). Motives for private income transfers. *The Journal of Political Economy*, 95(3), 508–546.
- Deaton, A. (2003). Health, income and inequality. *National Bureau of Economic Research Reporter: Research Summary*. Retrieved August, 15, 2009.
- Deaton, A. S., & Paxson, C. (2001). Mortality, education, income, and inequality among American cohorts. In David A. Wise (Ed.), *Themes in the economics of aging* (pp. 129–170). Chicago: University of Chicago Press.
- Diamond, P. (1965). National debt in a neoclassical growth model. *The American Economic Review*, 55(5), 1126–1150.
- Diehl, K. (2008). Mögliche Ursachen für die rasche Reduktion der ostdeutschen Übersterblichkeit nach der Wiedervereinigung. *Zeitschrift für Bevölkerungswissenschaft*, 33(1), 89–110.
- Ellwardt, L., Aartsen, M., Deeg, D., & Steverink, N. (2013). Does loneliness mediate the relation between social support and cognitive functioning in later life? *Social Science & Medicine*, 98, 116–124.
- Eng, P. M., Rimm, E. B., Fitzmaurice, G., & Kawachi, I. (2002). Social ties and change in social ties in relation to subsequent total and cause-specific mortality and coronary heart disease incidence in men. *American Journal of Epidemiology*, 155(8), 700–709.
- Gaudecker, H.-M., & Scholz, R. (2007). Differential mortality by lifetime earnings in Germany. *Demographic Research*, 17, 83–108.
- Grossman, M. (1972). On the concept of health capital and the demand for health. *The Journal of Political Economy*, 80(2), 223–255.
- Curven, M., Stieglitz, J., Hooper, P. L., Gomes, C., & Kaplan, H. (2012). From the womb to the tomb: The role of transfers in shaping the evolved human life history. *Experimental Gerontology*, 47(10), 807–813. Biodemographic determinants of lifespan.

- Henretta, J. C., Hill, M. S., Li, W., Soldo, B. J., & Wolf, D. A. (1997). Selection of children to provide care: The effect of earlier parental transfers. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 52(Special Issue), 110–119.
- Holt-Lunstad, J., Smith, T. B., & Layton, J. B. (2010). Social relationships and mortality risk: A meta-analytic review. *PLoS Med*, 7(7), e1000316.
- Human Mortality Database. (2013). *Human Mortality Database*, volume March 11th. University of California, Berkeley and Max Planck Institute for Demographic Research, Rostock.
- Kibele, E. U., Jasilionis, D., & Shkolnikov, V. M. (2013). Widening socioeconomic differences in mortality among men aged 65 years and older in Germany. *Journal of Epidemiology and Community Health*, 67, 453–457.
- Kluge, F. A. (2010). How East and West Germans finance their lifecycle consumption: evidence from NTA. *MPIDR Working Paper*, 027.
- Koh, S.-K., & MacDonald, M. (2006). Financial reciprocity and elder care: Interdependent resource transfers. *Journal of Family and Economic Issues*, 27(3), 420–436.
- Kohli, M., Künemund, H., Motel, A., & Szydlik, M. (2002). Families apart? Intergenerational transfers in East And West Germany. In S. Arber, C. Attias-Donfut (Eds.), *The myth of generational conflict. The family and the state in ageing societies* (pp. 88–99). London: Routledge.
- Kroh, M., Neiss, H., Kroll, L., & Lampert, T. (2012). Menschen mit hohem Einkommen leben länger. *DIW-Wochenbericht*, 79(38), 3–15.
- Lee, R. (1994). The formal demography of population aging, transfers, and the economic life cycle. In L. G. Martin & S. H. Preston (Eds.), *Demography of Aging* (pp. 8–49). Washington: National Academy Press.
- Lee, R., & Chu, C. C. (2012). The evolution of transfers and life histories. *Experimental Gerontology*, 47(10), 803–806. Biodemographic determinants of lifespan.
- Lee, R., & Mason, A. (Eds.) (2011). *Population aging and the generational economy – A global perspective*. Cheltenham: Edward Elgar.
- Moorman, S. M., & Stokes, J. E. (2014). Solidarity in the grandparent-adult grandchild relationship and trajectories of depressive symptoms. *The Gerontologist*, 1–13. doi:10.1093/geront/gnu056. Research Article, online first.
- Norton, E. C., & Van Houtven, C. H. (2006). Inter-vivos transfers and exchange. *Southern Economic Journal*, 73(1), 157–172.
- Preston, S. (1975). The changing relation between mortality and level of economic development. *Population Studies*, 29(2), 231–248.
- Samuelson, P. A. (1958). An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money. *The Journal of Political Economy*, 66(6):467–482.
- Seeman, T. E. (1996). Social ties and health: The benefits of social integration. *Annals of Epidemiology*, 6(5), 442–451.
- Simonson, J., Kelle, N., Gordo, L. R., Grabka, M. M., Rasner, A., & Westermeier, C. (2012). Ostdeutsche Männer um 50 müssen mit geringeren Renten rechnen. *DIW-Wochenbericht*, 79(23), 3–13.
- Staatliche Zentralverwaltung fuer Statistik. (1988). Arbeitsrichtlinie für das Bezirkszählbüro zur Vorbereitung und Durchführung der Einkommensstichprobe 1988 in Arbeiter- und Angestelltenhaushalten. *Drucksache E 4. Ministerrat der DDR, Berlin*.
- Statistisches Bundesamt. (1999). Bevölkerungsstatistische Übersichten 1946–1989. *Sonderreihe mit Beiträgen für das Gebiet der ehemaligen DDR*, 32.
- Statistisches Bundesamt. (2000). Bevölkerung und Erwerbstätigkeit. *Gebiet und Bevölkerung*, Fachserie 1, Reihe 1.
- Steptoe, A., Shankar, A., Demakakos, P., & Wardle, J. (2013). Social isolation, loneliness, and all-cause mortality in older men and women. *Proceedings of the National Academy of Sciences*, 110(15), 5797–5801.
- Stringhini, S., Berkman, L., Dugravot, A., Ferrie, J. E., Marmot, M., Kivimaki, M., & Singh-Manoux, A. (2012). Socioeconomic status, structural and functional measures of social

- support, and mortality: The British Whitehall II cohort study, 1985–2009. *American Journal of Epidemiology*, 175(12), 1275–1283.
- Thiele, J. C. (2014). R marries NetLogo: Introduction to the RNetLogo package. *Journal of Statistical Software*, 58(2), 1–41.
- Vogt, T. C. (2013). How many years of life did the fall of the Berlin wall add? A projection of East German life expectancy. *Gerontology*, 59(3), 276–282.
- Vogt, T. C., & Vaupel, J. W. (2013). Does higher income yield higher life expectancy? Findings from East Germany after reunification. Unpublished manuscript.
- Wilensky, U. (1999). NetLogo: Center for connected learning and computer-based modeling. *Northwestern University*. <http://ccl.northwestern.edu/netlogo>.
- Yaari, M. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2), 137–150.
- Zunzunegui, M.-V., Bland, F., & Otero, A. (2001). Support from children, living arrangements, self-rated health and depressive symptoms of older people in Spain. *International Journal of Epidemiology*, 30(5), 1090–1099.

Chapter 16

Agent-Based Modelling to Inform Health Intervention Strategies: The Case of Severe Acute Malnutrition in Children in High-Burden Low-Income Countries

Hedwig Deconinck, Carine Van Malderen, Niko Speybroeck, Jean Macq, and Jean-Christophe Chiem

16.1 Introduction

Severe acute malnutrition (SAM) affects 17 million children worldwide (World Health Organization et al. 2012) and kills over half a million children annually (Black et al. 2013). Children with SAM have a high risk of death, exceeding nine times higher than that of children without the condition, and require intensive medical support (World Health Organization and United Nations Children's Fund 2009). The health and socio-economic implications of SAM are of great concern for countries committed to reaching the World Health Assembly Global Target 6 (World Health Organization et al. 2014) of reducing and maintaining childhood wasting to less than 5 % by 2025.

Until 2000, children with SAM were managed as inpatients with low coverage and high case-fatality. The innovation of ready-to-use therapeutic food allowed children with uncomplicated SAM to be treated at home and made decentralised outpatient care possible. The outpatient approach was piloted in nutrition emergencies in Sudan, Malawi and Ethiopia where evidence showed its potential to reach many malnourished children and improve their recovery and survival (Zinszer et al. 2014; Collins et al. 2006). The package of health interventions to improve the outpatient

Electronic supplementary material The online version of this chapter (doi: [10.1007/978-3-319-32283-4_16](https://doi.org/10.1007/978-3-319-32283-4_16)) contains supplementary material, which is available to authorized users.

H. Deconinck (✉) • C. Van Malderen • N. Speybroeck • J. Macq • J.-C. Chiem
Institut de recherche santé et société, Université catholique de Louvain,
Clos Chapelle-aux-champs 30 bte30.13, 1200 Brussels, Belgium
e-mail: hedwig.deconinck@gmail.com; carine.vanmalderen@uclouvain.be;
Niko.Speybroeck@uclouvain.be; jean.macq@uclouvain.be; jean-christophe.chiem@uclouvain.be

management of SAM aims to identify children under 5 years of age with SAM early in the community, and refer and treat them in primary health care before complications develop and hospitalisation is needed. SAM interventions include facility-based interventions for improving operational efficiency and quality of care, as well as community-based interventions to activate and inform the population. The endorsement of this approach by United Nations agencies in 2007 facilitated the introduction and scale-up of SAM in primary health care as a routine child health intervention (World Health Organization et al. 2007). The cost-effectiveness of the management of SAM in primary health care (Black et al. 2013; Scaling Up Nutrition Movement 2010) has facilitated the introduction of SAM services in over 60 high-burden low- and middle-income countries to date (United Nations Children's Fund 2013). To foster scale-up, a robust intervention strategy is needed to enhance both the coverage and quality of decentralised SAM services to improve child survival.

16.1.1 Challenges of Designing SAM Intervention Strategies

The effectiveness of SAM interventions encounters a number of methodological and practical challenges related to contextual and human factors. SAM intervention strategy designs may overcome these challenges by accessing relevant information and applying a systems lens. To date, effectiveness studies of SAM management analysed performance of service delivery and health outcomes but not health system functioning (Deconinck et al. 2015). Moreover, health system functioning data for the management of SAM are difficult to measure and interpret, and study methods have not yet been adapted. There are several reasons for this. The management of SAM has only recently been introduced in the basic package of national health services as SAM's recognition as a disease and primary cause of death was delayed (World Health Organization et al. 2014). Second, the management of SAM is a complex intervention including several components with many actors at various delivery platforms (Deconinck et al. 2015) and comorbidities (World Health Organization 2013; Jones and Berkley 2013). Third, contextual and human factors, such as SAM affecting the poorest populations living in rural and remote areas and motivation of health workers (Joint Learning Initiative 2004), played a critical role in the effectiveness of SAM interventions. Such effects may be hard to assess because either health system performance data and analytical tools are not available, or the dynamics of health system functions are not understood. While information specific to other health interventions may be utilized as a proxy indicator for SAM interventions, its use may lead to weak strategy formulations, and its insights may not be generalizable for strengthening the overall health system. For example, studies that identified factors boosting or impairing SAM efficiency did not examine the effects of interventions on reducing barriers. Instead, the studies identified demand- and supply-side barriers to SAM management access, use and quality that constrained effective coverage and proposed intuitive one-to-one solutions (Rogers et al. 2015). Common demand-side barriers to SAM service access and use – distance to the site




Fig. 16.1 Demand- and supply- side barriers to health interventions

of care and lack of awareness of the disease and its care path – may be addressed by community-based interventions (Fig. 16.1). Common supply-side barriers that affect SAM service performance – poor quality of care and low motivation of health workers – may be addressed by health facility-based interventions.

A system perspective lens examines SAM interventions within the overall health system performance rather than as disease-specific isolated events. To date, study analyses (Deconinck et al. 2015) ignored the effect of synergies, interactions or feedback loops among barriers, interventions and health systems as illustrated in a simplified way in Fig. 16.2. Intervention strategies derived from this linear thinking may therefore underperform. In addition, the complexity of the SAM health condition and interventions require addressing causal factors across multiple delivery platforms, sectors and actors. At the individual level, interactions between service providers and users shape their behaviour and the behaviour of the system. Longitudinal data would not be able to assess these system effects. For example, individual human factors such as awareness, motivation, preferences and perceptions drive health seeking behaviour, and are difficult to study and quantify. There is consensus on the critical role of intricate interactions between system functions (Luke and Stamatakis 2012) that drive the health system dynamics, leading to effects such as feedback loops and delays. However, health actors (e.g. policy makers, planners, managers, service providers and users) may not always understand what works for whom and why, and what system effects positively or negatively influence and interact. Therefore, methodological approaches that capture systems effects and model the critical role of contextual and human factors impacting the health system may be lacking (Macq and Chiem 2011; Witter et al. 2013). An unstructured methodological approach often leads to the unproductive conclusion that all factors influence each other. In this context, health interventions that fail to encompass a system perspective do not detect unintended results and may have inefficient or counterintuitive results (Forrester 1971). For example, increasing access to a health service (a demand-side factor) while maintaining a low quality (a supply-side factor) may increase mistrust in the health services and lead to disintegration of its overall performance. The latter adverse effect was unintended and resulted from both service users', and providers' behaviours (Fig. 16.2).

Strategies for improving the management of SAM include interventions that act on both demand- and supply-side levels. On the one hand, health workers need the




Fig. 16.2 System perspective on demand- and supply-side factors of health interventions

competence and ability to provide quality health care, motivation and accountability, all variables that are intrinsically linked. Interventions that improve supply-side factors include involvement in design and management, career development opportunities, continual education, on-the-job mentoring, supportive supervision and access to essential supplies. On the other hand, service users need to be informed and activated to stimulate demand and trust. Interventions that improve demand-side factors include raising awareness, involving communities in service delivery, and bringing services closer to the users. The multitude of contextual and human factors influencing and interacting suggests that intervention strategies successful in one context cannot be easily transferred and scaled up in another context without sound understanding.

16.1.2 Markov Versus Agent-Based Modelling


To enhance the understanding of SAM interventions and complexity, this study explored the role of modelling encouraged by the following example and tool. The management of SAM is comparable to the management of serious mental illness with comorbidities, which Silverman et al. (2015) described as a complex system problem. The study found that the many layers of the health system were themselves networked systems (human physiology, family, community) with different parts having their own motivations that behaved probabilistically, leading to unexpected patterns (e.g. lowering the number of patients followed by a nurse lowered patient readmission and reduced costs) (Silverman et al. 2015). When applying the SIMULATE (System, Interactions, Multi-level, Understanding, Loops, Agents, Time, Emergence) checklist (Marshall et al. 2015) to the management of SAM, it was determined that dynamic simulation modelling was appropriate. Indeed, management of SAM involves processes made up of multiple events (detection, health seeking, diagnosis, treatment initiation and progression, and recovery, defaulting or death) and relationships between demand- and supply-side variables (System). These relationships are non-linear, making outcomes difficult to anticipate (Interactions). Moreover, the management of SAM can be modelled

from different strategic or operational perspectives; it may be demand- or supply-side oriented (Multilevel). The complex features already mentioned (e.g. individual factors, such as changing motivation) make it impossible to solve the problem analytically (Understanding). Moreover, the management of SAM involves feedback loops, for example, admission → good outcomes → positive opinion → admission (Rogers et al. 2015) (Loops), interaction of the behavioural properties of multiple agents (children, health workers) (Agents), and time-dependant transitions from different health states (SAM, recovered, died) (Time). Finally, the problem requires considering the intended and unintended consequences of SAM interventions to address policy resistance and achieve target outcomes (Emergence).

In a first simulation, we applied Markov modelling to illustrate how this approach may contribute to a better understanding of the evolution of the SAM process in children. Markov models are widely used to evaluate health care interventions (Marshall et al. 2015). They assume that individuals are in one of a finite number of discrete health states, and transition from one state to another by an event (Sonnenberg and Beck 1993). These models address some points of the SIMULATE list, in particular those of repeated events and time-dependent transitions across different health states. As the analysis of a complex health care delivery system such as the management of SAM requires a more flexible method (Marshall et al. 2015), in a second simulation, we applied agent-based modelling (ABM). This approach makes modelling of the remaining elements of the SIMULATE list (non-linear and feedback relationships, demand and supply motivations, behaviours and strategic perspectives) and the emergence of intended and unintended consequences possible. ABM defines an agent as a set of variables; several types of agents can be created and their rules of interactions can be defined (Fig. 16.3). These rules of interaction can be expressed using a logical language such as ‘if-then’ statements (Collopy and Armstrong 1992; Gilbert and Terna 2000). This model definition can then be applied in a computational model, simulating the dynamic actions and interactions of agents as individuals or collective entities. The power and flexibility of simulation can help model multiple scenarios, which can possibly integrate diverse data such as volatile contextual and human factors, but also detailed geographical maps mimicking real environments, for example, of people seeking health care. The introduction of this level of detail helps to create more realistic diffusion patterns of disease.

These features make ABM an appealing methodological tool to inform health intervention strategies to improve the management of SAM. First, data that are

Fig. 16.3 Agent-based model with agents, their characteristics (variables 1 and 2) and interactions defined as if-then rules



simulated could compensate to some extent for the paucity of information regarding SAM health system functions. Indeed, the role of contextual and human factors such as awareness of service users and motivation of service providers that influence health outcomes could be modelled. In addition, the simulation can generate longitudinal data under various scenarios adding a time dimension by modelling for example the short- and long-term impact of different types of interventions. Second, ABM can represent a complex system of agents that interact through feedback loops and with delay effects. In fact, the complexity of the SAM interventions and the interaction of health system functions call for applying such a systems lens. By creating a dynamic representation of a complex system and comparing scenarios, an agent-based model can test how community- and facility-based interventions improve the effectiveness of SAM services by acting on health-seeking behaviour and quality of care (Gilbert and Terna 2000; Swanson et al. 2012). Finally, the structure underlying ABM is particularly well suited to capturing expert knowledge. Tacit and empirical knowledge have given experts a good understanding of elements and actors that play a role in a health system. They are therefore well placed to provide relevant and evidence-based information to define agents and their variables and interactions. Moreover, the 'if-then' rules that define the interactions do not need mathematical formulas.

ABM has already supported successful participatory research (Bousquet et al. 2007; Edmonds 2010; Salerno et al. 2010). In these exercises, participating health actors formulated rules and subsequently compared and interpreted results from simulated scenarios. ABM scenarios may thus enhance group learning for understanding problems and formulating solutions for improved health interventions (Atun 2012) that are context adapted and apply a systems lens.

16.1.3 Aim of the Case Study

Features of ABM that could improve the design of health intervention strategies such as those targeting SAM have been listed in the previous section. However, to date, the number of applications of ABM to inform such intervention strategies has been limited (Lempert 2002; Brailsford 2005). The purpose of this paper is to illustrate the actual benefits and limitations of ABM to inform complex health intervention strategy designs. We therefore developed a simulation model to improve understanding of how a package of health interventions that address demand- and supply-side barriers may improve the effectiveness of SAM service delivery and ultimately survival of children with SAM.

16.2 Methods and Results

Based on both expert knowledge and scientific evidence, we designed a conceptual framework of a model defining agents and their rules of interactions. In a first exploration, interactions were based only on transition probabilities, leading to a

simple Markov model. We then introduced elaborate rules taking advantage of the specific features of ABM. Finally, we defined and analysed several intervention scenarios, and analysed benefits and limitations of ABM.

16.2.1 Material

The successive models developed in this study resulted from structured discussions on barriers to the implementation of SAM interventions. These discussions involved two researchers in simulation utilisation who designed an agent-based model and a SAM expert with extensive knowledge of the evidence base and experience in SAM intervention strategy implementation.

Scientific evidence supported the selection and formulation of the variables, rules and parameter values in the final model. Simulation models were implemented using NetLogo 5.0.4. Source codes of subsequent models are in Appendix 1 in the electronic supplementary material. Appendix 2 describes the simulation as formatted within the standardized overview, design concepts, and details (ODD) protocol (Railsback and Grimm 2011). Appendix 3 includes pseudo-codes and a flow chart, and Appendix 4 includes sensitivity analysis results. Appendices 5 and 6 provide the NetLogo programs of the Markov and ABM models.

16.2.2 Conceptual Framework

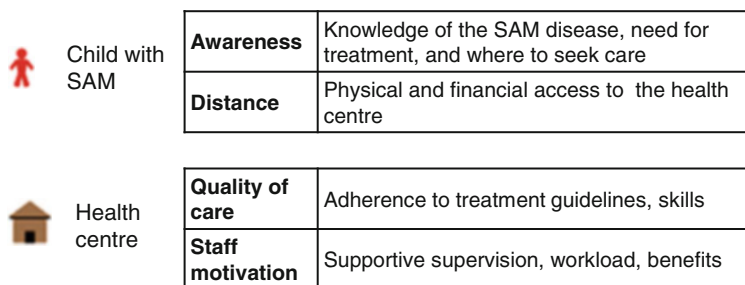
This section first defines SAM in children under 5 and the package of health interventions for improving the management of SAM before selecting and defining agents and their variables. The interactions between the intervention components are then specified. Finally, outcomes of interest are defined, together with assumptions inherent to the modelling exercise.

16.2.2.1 SAM and Health Interventions

The definition of SAM in children under 5 details the clinical symptoms and diagnostic measures that establish the condition, as supported by the literature (World Health Organization et al. 2014) (Table 16.1). The package of health interventions for improving the management of SAM consists of intervention components to improve service uptake (demand-side) and service performance (supply-side) (Table 16.1).

Table 16.1 Definition of severe acute malnutrition (SAM) in children under 5 and health interventions

Severe acute malnutrition (SAM) in children under 5 years of age	SAM is a result of recurrent infections, poor dietary intake or poor nutrient absorption. SAM is defined by weight-for-height below -3 or more standard deviations of the median value of the WHO 2006 Child Growth Standard, mid-upper arm circumference less than 115 mm, or presence of nutritional oedema (World Health Organization 2013)
	This study focused on uncomplicated SAM, i.e. with appetite, no other severe illness and no need for hospitalisation
Health interventions for improved management of uncomplicated SAM	Community-based interventions to improve community involvement and demand creation for SAM may include sensitizing community members and groups; selecting, training and supervising community workers; screening in communities for early detection and referral; establishing referral systems and transportation; visiting homes for follow-up of SAM problem cases and counselling
	Health centre-based interventions to improve the outpatient management of SAM without medical complications may include diagnosing, referring and treating cases; counselling carers; training and supervising health workers; managing medical equipment and supplies; and providing technical support and tools for knowledge and skills development of clinical care and organisational management (Deconinck et al. 2015)

**Fig. 16.4** Selected agents with variables related to severe acute malnutrition (SAM) interventions

16.2.2.2 Agents

The model of SAM interventions that address specific questions of supply- and demand-side barriers involves two types of ‘agents’ with a corresponding set of attributes, or ‘variables’: a child with SAM and a carer (e.g. mother), and a health centre (Fig. 16.4).

Definitions of the agents’ variables are provided in Table 16.2, and for simplicity the variable labels are used in the text.

Table 16.2 Definition of agents' variables of severe acute malnutrition (SAM) interventions and possible corresponding interventions

Variable	Definition, threshold and example of interventions
Quality of care (0–100)	<p>Definition:</p> <p>Quality of care is defined by the degree to which health services for individuals and populations increase the likelihood of desired health outcomes that are consistent with current professional knowledge (Mainz 2003). Quality of care also includes perceptions of service users and providers, as well as other qualitative elements (World Health Organization 2006a). This study evaluated quality of care as adherence to the national treatment protocol on a continuous scale from low to high</p> <p>Examples of interventions to improve the quality of care are providing training; disseminating guidelines and other job aids; and providing timely essentials supplies (preventing stockouts of antibiotics and therapeutic food)</p>
Staff motivation (0–100)	<p>Definition:</p> <p>Staff motivation (or satisfaction) is defined as the level of effort and desire to perform well (World Health Organization 2006b). It may be influenced by financial and non-financial incentives, and indirectly evaluated by, e.g. unjustified absenteeism, and low staff performance (ONEil and Reimann 2013). This study evaluated staff motivation on a continuous scale from low to high</p> <p>Examples of interventions to improve staff motivation are providing supportive supervision that includes appraisal and mentoring; improving knowledge and skills through training; creating a positive working environment; providing career development opportunities; and providing an additional pay, such as transport, hardship or education allowances</p>
Awareness (0/1)	<p>Definition:</p> <p>Awareness is defined by the carer's (e.g. mother's) knowledge of why SAM is a problem, how it can be solved, and where to seek care. This study evaluated awareness as a binary variable; awareness is acceptable when the carer knows the problem and solution for the child's illness and where to seek care</p> <p>Examples of interventions to improve awareness are providing health and nutrition counselling for social and behaviour change and self-referral; and involving communities in designing, planning and implementing service delivery for improved accountability and trust</p>
Distance (0/1)	<p>Definition:</p> <p>Distance is defined by average hours (or km) to walk to the site of care (in our case the health centre) (World Health Organization 2008). This study evaluated distance as a binary variable; distance is acceptable when the site is less than a 1 h walk or 4 km from where the carer and child live</p> <p>Examples of interventions to improve distance are providing means of transport; and decentralising care through increased functional health facilities, mobile teams or community case management</p>

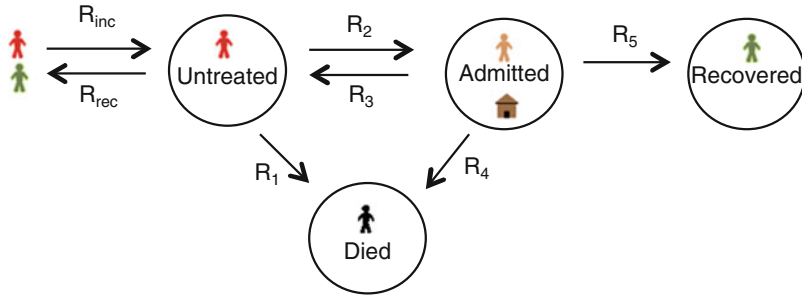


Fig. 16.5 Treatment pathways of severe acute malnutrition (SAM) in children under 5 transitioning from untreated to admitted, died or recovered states

16.2.2.3 Interactions

As shown in Fig. 16.5, children and health centres interact and define a typical treatment pathway with four health states: untreated, admitted (for treatment), recovered and died. The rules indicating the transition between each health state are referred to as case fatality when untreated (R_1), coverage (R_2), defaulting (R_3), case fatality in treatment (R_4) and recovery (R_5). In addition, R_{inc} indicates incidence and R_{rec} indicates spontaneous recovery of children with SAM living in the catchment area of a health centre. The catchment area of a health centre has been defined as the area of residence of its users. Agents are linked to one health centre. Definitions of transition rules are provided in Table 16.3 with thresholds of acceptability based on international standards.

16.2.2.4 Outcomes

Three outcomes of interest were monitored to assess the simulation evolution, together with the possible impact of health interventions: overall case fatality, coverage, and recovery (Table 16.4).

The simulation of SAM interventions also required the choice of a relevant time scale to evaluate the transition between health states. Based on the SAM expert’s opinion and supported by the definitions, 1 week appeared a suitable time unit to re-evaluate the state transition, corresponding to children’s weekly follow-up visits to the health centre for monitoring treatment progress. In the same way, each simulation ran for 100 weeks to represent approximately a 2-year time frame, which was deemed a reasonable period in which to achieve a stable level of health system performance and assess the effect of a possible health intervention strategy. For each scenario presented further, 50 simulations were run using the same set of parameters. These simulations were averaged to reveal the main trend.

Table 16.3 Definition and probability estimates of transition states of severe acute malnutrition (SAM) in children under 5 in high-burden low-income countries

Code	Transition rules (rate)	Definition	Value (range)	Justification
R ₁	Case fatality untreated	Proportion of children with untreated severe acute malnutrition (SAM) who died of SAM disease	15 %	The under 5 mortality in Niger in 2013 was 0.66 deaths per 10,000 children per day, or 2.4 % per year assuming that mortality was stable year around, which in reality is not the case (République du Niger Institut National de la Statistique 2013). Children with SAM have a nine times higher risk of mortality than other children (World Health Organization and United Nations Children's Fund 2009), thus on average 21.6 % (or $9 \times 2.4 \%$)
R ₂	Coverage	Proportion of children with SAM who are in treatment (admitted)	30 %	The acceptable coverage rate ^a in rural areas is $\geq 50 \%$ (The Sphere Project 2011). Coverage of $< 30 \%$ is most common, unless good community-based support is in place, but rarely surpasses 50 % (Becart 2014)
			(20–50 %)	
R ₃	Defaulting	Proportion of children with SAM who ended treatment against medical advice	20 %	The acceptable defaulting rate is $< 15 \%$ (The Sphere Project 2011). The defaulting rate is high ($\geq 50 \%$) in areas with low support. The national average defaulting rate in outpatient care in Niger in 2013 was very low, at 7.1 %
R ₄	Case fatality in treatment	Proportion of children with SAM in treatment who died of SAM disease	5 %	The acceptable case fatality rate in treatment is $< 10 \%$ (The Sphere Project 2011)
R ₅	Recovery	Proportion of children with SAM who restored their health after having received treatment	40 %	The acceptable and targeted recovery rate is $\geq 75 \%$ ^b (The Sphere Project 2011). Recovery rate may be low because of either high case fatality or high defaulting
			(30–75 %)	
R _{rec}	Spontaneous recovery	Proportion of children with SAM who spontaneously recovered without treatment	30 %	The spontaneous recovery rate based on low mid-upper arm circumference is 75–90 % in historical cohort studies in Senegal and the Democratic Republic of Congo (Garenne et al. 2009). When the simulation process required adjustment of parameters in the model, empirical and tacit knowledge proposed 30 % spontaneous recovery of SAM to be a reliable figure

(continued)

Table 16.3 (continued)

Code	Transition rules (rate)	Definition	Value (range)	Justification
R _{inc}	Incidence	Proportion of new SAM cases in a given time period	30 children	For a child population of 1,000 individuals, we expect 30 children with SAM at t ₀ , based on a 3 % SAM prevalence rate that is common for high-burden low-income countries

^aCoverage, in this study ‘contact coverage’, is an indicator of effectiveness of community-based interventions that measures service access and use

^bRecovery of SAM is an indicator of effectiveness of facility-based interventions. A child who has recovered, is clinically well, has no oedema and/or has improved mid-upper arm circumference (≥ 125 mm) or weight-for-height (≥ -2 z-score of the WHO Child Growth Standard population median)

Table 16.4 Definition of outcomes of severe acute malnutrition (SAM) in children under 5

Outcome (rate)	Definition
Overall case fatality	Proportion of children with severe acute malnutrition (SAM) who died of SAM disease
Coverage	Proportion of children with SAM who are in treatment (admitted)
Recovery	Proportion of children with SAM who restored their health after having completed treatment

16.2.2.5 Assumptions

The design of the model described above led to inherent assumptions that should be kept in mind when interpreting the results:

- Children do not grow old;
- Children enter the study population with a stable incidence of SAM;
- Children are admitted for treatment, and recover, default and die weekly;
- Children are treated at a unique health centre;
- Children exit the study if they die or recover; and
- Interventions target all community members regardless of their socio-economic status.

16.2.3 Markov Model: Probability Rules

In a first approach, transition rules R₁ to R₅ were described using simple probabilities. In technical terms, this reflects a simple Markov model in which children may transit through four different states: untreated, admitted (for treatment), recovered and died. Each week, new children are affected with SAM with an incidence R_{inc} and children with SAM may transit from one state to another. Untreated children have an R_{rec} probability of recovering spontaneously, an R₁ probability

of dying without treatment (case fatality untreated) and an R_2 probability of starting treatment services (admission, a proxy for coverage). Treated children have an R_3 probability of defaulting (ending treatment against medical advice), an R_4 probability of dying during treatment (case fatality during treatment), and an R_5 probability of recovering (restoring their health). During the simulation, changes in R_2 and R_5 modified the overall case fatality patterns. Search of the literature together with expert opinion referenced realistic or acceptable baseline values in high-burden low-income countries for the respective parameters used in the simulation (Table 16.3).

While realistic probabilities set up an initial baseline (if available, logistic regression could also generate probabilities based on empirical data), the simulation tools investigated several sets of probabilities leading to different scenarios. On the one hand, a demand-side intervention could impact coverage (R_2), increasing the probability of being admitted for treatment from 30 % (low) to 40 % (middle) and 50 % (high). On the other hand, a supply-side intervention could impact recovery (R_5), raising the chances of being successfully treated from 40 % (low) to 60 % (middle) and 75 % (high). Both interventions could be mixed to different extents.

Overall case fatality curves corresponding to these scenarios (Fig. 16.6) showed general convergence towards asymptotic probabilities as expected in Markov models. The upper line (thick dark grey) represents the initial set of parameters. The thin grey lines were obtained by progressively increasing the coverage (R_2 , continuous lines) and recovery parameters (R_5 , dotted lines) from low to middle and high ranges. In this simulation, the lowest line shows that increasing both coverage and recovery parameters from a low to middle range simultaneously had a greater impact on overall case fatality than increasing each parameter separately.

16.2.4 Agent-Based Model: Conditional Rules

An agent-based model was defined to integrate the role of history and conditional if-then rules that represent the impact of contextual and human factors Fig. 16.7.

16.2.4.1 The Role of History

In this model, a child needed 8 weeks of treatment to recover from SAM. Note that this procedure distorts the meaning of the probabilities accounted for in the previous Markov model. For example, the probability of defaulting now spanned 8 weeks for each child. This appears closer to the reality of the treatment process, which requires on average 6 weeks (World Health Organization 2013). In addition, defaulting within this time interval may have an effect on other factors such as recovery and awareness of care within the population, as will be discussed further.




Fig. 16.6 Outcomes of Markov model simulation of severe acute malnutrition (SAM) interventions: overall case fatality at different levels of coverage and recovery modified separately and simultaneously. The *upper line (thick dark grey)* represents a simulation that was obtained with initial values of parameters ($R_2 = 30\%$, $R_5 = 40\%$). *Thin light grey full lines* were obtained by increasing coverage (R_2) parameter values to 40% (*middle*) and 50% (*high*), and *dotted lines* were obtained by increasing recovery (R_5) parameter values to 60% (*middle*) and 75% (*high*). The *lowest line* was obtained by increasing both R_2 and R_5 parameter values from a low to middle range

16.2.4.2 If-Then Rules

Taking advantage of ABM capacities, transition rules can be described involving both demand- and supply-side factors, which impact individual factors. For example:

- R_2 (coverage): *If* an untreated child (his or her carer) has a certain awareness (i.e. knows what the problem and solution are and where to access care) and has acceptable distance (i.e. lives less than 1 h walk from a health centre), *then* the child has a chance of being admitted for treatment;

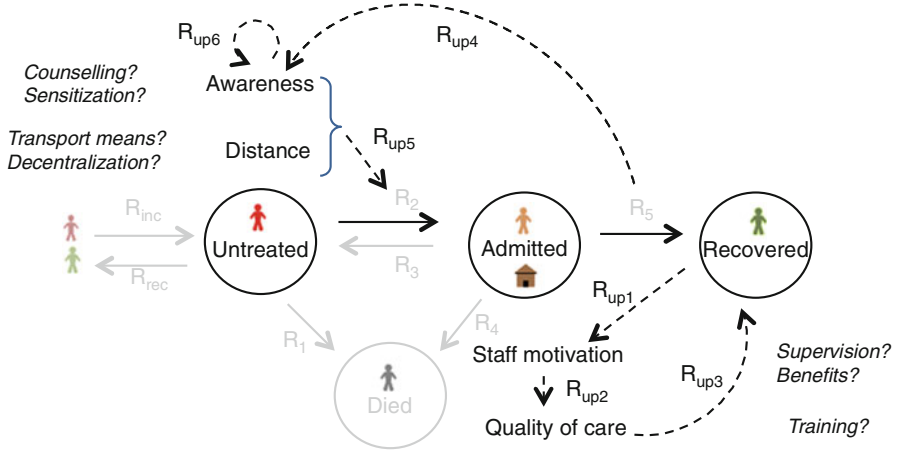


Fig. 16.7 Treatment pathways with interactions of individual factors and *intervention levers* (in *italic*). R_{up1} stands for recovery influencing staff motivation, R_{up2} for staff motivation influencing quality of care, R_{up3} for quality of care influencing recovery, R_{up4} for recovery influencing awareness, R_{up5} for awareness and distance influencing admission, and R_{up6} for individual awareness influencing community awareness

- R_3 (defaulting): *If a child defaults, then his or her (carer’s) awareness is reset to zero; and*
- R_5 (recovery): *If the quality of care is sufficient, then the admitted child has a chance of recovering.*

Further rules can be added to reflect the role of contextual and human factors.

The model introduced three feedback loops. First, the observed recovery influenced staff motivation (R_{up1}), which has an impact on the quality of care (R_{up2}), which is itself a determinant of the chance of recovery (R_{up3}). These rules influenced the supply-side factors. Second, the observed recovery influenced the carer’s awareness (R_{up4}), which is a precondition for accessing and using the services (R_{up5}). These rules made the links between supply- and demand-side attributes. They created feedback loops that led to both vicious (if one of the three elements worsens) and virtuous (if one element improves) cycles. Third, awareness generated a feedback loop on itself (R_{up6}) as individual awareness increases community awareness. This rule induced interactions among individuals that may play a critical role in diffusing behaviours in the whole population. This type of emerging pattern is referred to as a network effect (Luke and Stamatakis 2012). Figure 16.7 shows the interactions among the rules, and Fig 16.8 shows the influence of individual factors on rates used in the model. Details on the rules and interventions are provided in the pseudo code (Appendix 3).

	Parameter	Rule
Supply-side	R_{up1}	If recovery in(de)creases, then staff motivation in(de)creases.
	R_{up2}	If staff motivation in(de)creases, then quality of care in(de)creases.
	R_{up3}	If quality of care in(de)creases, then recovery in(de)creases.
Demand-side	R_{up4}	If recovery in(de)creases, then awareness in(de)creases.
	R_{up5}	If a child has awareness and an acceptable distance, then (s)he has a chance to be admitted.
	R_{up6}	If individual awareness in(de)creases, then community awareness in(de)creases.

Fig. 16.8 Influence of individual factors on rates (R) used in the agent-based models of severe acute malnutrition (SAM) interventions. R_{up1} stands for recovery influencing staff motivation, R_{up2} for staff motivation influencing quality of care, R_{up3} for quality of care influencing recovery, R_{up4} for recovery influencing awareness, R_{up5} for awareness and distance influencing admission, and R_{up6} for individual awareness influencing community awareness. Because of the lack of data and formal equations to model changes, the specification of thresholds for R_{up1-3} relied on expert opinion. Several thresholds were tested and those giving the most realistic evolution were retained

Table 16.5 Plausible values for community awareness and distance and health centre quality of care, and staff motivation of severe acute malnutrition (SAM) interventions

Factor (range)	Value
Community awareness (0–100):	20
The proportion of children whose carers are aware of the severe acute malnutrition (SAM) problem and solution and treatment site	
Community distance (access) (0–100):	50
The proportion of children living at an acceptable ‘distance’	
Quality of care (0–100):	50
Adherence to the national treatment protocol	
Staff motivation (0–100):	50
Level of effort and desire to perform well	

16.2.4.3 Initial Scenario

Initial plausible values for community awareness and distance, and health centre quality of care and staff motivation were set to 20 %, 50 %, 50 % and 50 %, respectively, using realistic probabilities (Table 16.5).

The resulting simulation is presented as Scenario I in Fig. 16.9. The corresponding graph plots curves of the three outcomes: coverage, recovery and overall case fatality. As the simulation evolved, coverage and recovery outcomes remained low, and overall case fatality remained high.

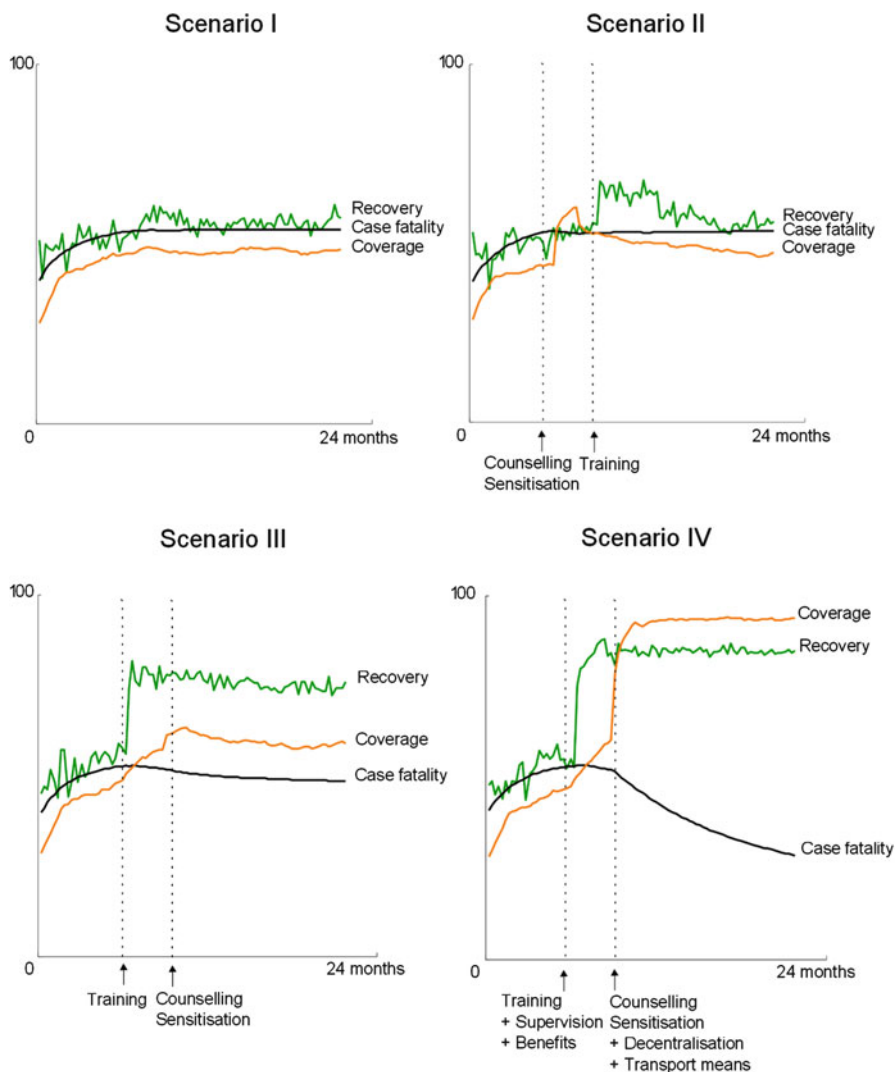


Fig. 16.9 Results of simulation model based on ABM for the four intervention scenarios showing coverage, recovery and overall case fatality rates of severe acute malnutrition (SAM) interventions. In this model, four scenarios are tested: I without intervention, II to IV with interventions applied at the sixth and ninth months. For each scenario, 50 simulations were run; the curves represent the mean values of the 50 simulations of the three outcomes, coverage, recovery and overall case fatality, for each week

16.2.4.4 Intervention Scenarios

Based on the initial Scenario I (Fig. 16.9) that used realistic probabilities, new intervention scenarios can be investigated and triggered over time. On the demand-side, interventions can address:

- Awareness through sensitisation and counselling; and/or
- Distance through decentralisation of care and availability of transport.

On the supply-side, interventions can address:

- Quality of care through staff training; and/or
- Staff motivation through supportive supervision and benefits, such as career development opportunity.

Because the Markov model showed that combining interventions decreased case fatality, we investigated a combination of demand- and supply-side interventions targeting contextual and human factors. Two subsequent interventions were introduced at the 6th month (estimated time to analyse the initial situation and prepare an intervention) and the 9th month (estimated time to verify progress and prepare a second intervention). The delay in introducing interventions allowed a first batch of children with SAM to end treatment and thus stabilised the virtual population. In scenario II (Fig. 16.9), demand-side interventions (for example, sensitisation and counselling to target awareness) were applied at the 6th month. Supply-side interventions (for example, training to target quality of care) were applied at the 9th month. Outcomes for this scenario did not improve by the end of the study period. In scenario III (Fig. 16.9), the same interventions were applied, but in reverse order. Here, outcomes slightly improved but were still not satisfactory by the end of the study period. In scenario IV (Fig. 16.9), supervision and benefits were applied together with training at the 6th month, and decentralisation and improved means of transport were applied together with sensitisation and counselling at the 9th month. Outcomes improved by the end of the study period: recovery and coverage rates increased and remained high and the overall case fatality rate decreased with time. Details on the rules and interventions are provided in the pseudo code (Appendix 3).

16.3 Discussion

In this paper, we designed a Markov model and an agent-based model to improve understanding of how health interventions address demand- and supply-side barriers and influence the effectiveness of the management of SAM. The simplicity of Markov modelling explains why it is widely used in diverse areas of public health (Sonnenberg and Beck 1993; Claxton et al. 2002) and was useful for investigating different scenarios of SAM interventions. The result of the simulation indicated possible synergy when improving both coverage and recovery. However, important limitations were inherent in this modelling approach. First, transition probabilities

were fixed in time, leaving no role for history in the dynamic process. In the same way, the impact of contextual effects may not be represented. Second, critical heterogeneous factors of human behaviour, such as awareness and staff motivation, could not be modelled in the Markov framework. The ABM approach on the other hand successfully addressed these shortcomings. While the Markov model showed that a combination of both demand- and supply-side interventions was needed, results from ABM discovered the best order and magnitude of these interventions. The benefits and limitations, and perspectives of working with a higher degree of complexity are further discussed.

16.3.1 Benefits

The specific abilities of ABM to compensate for paucity of data, provide a systems perspective, and model expert knowledge to inform health intervention strategy designs are assessed in this section. This section discusses the benefits and operational insights resulting from this.

The first benefit is that ABM allowed working in a context where critical data may not be available, solving the paucity of data in different ways. First, ABM required an ontological step of defining agents, variables and interactions, which allowed a clear articulation of a complex topic. While a traditional approach may concentrate on an exhaustive inventory of influencing factors, ABM enforced selecting a set of key variables. This selection process allowed focussing on key variables and their precise definition. Second, ABM made it possible to handle factors that were not necessarily quantifiable, such as determinants of human behaviour (awareness, motivation) and contextual factors (distance, history of treatment). While such elements may be mentioned but not included in conventional models, their role and interactions may have an important impact. Defining these factors and depicting their functioning can bring critical insight into efficient intervention strategies that would otherwise remain covered. Third, if field experts and researchers were unaware of existing data or theory, ABM allowed systematic searches for realistic values and justified validating rules of functioning within the published literature. Finally, if data were not available or precise, the use of expert opinion provided knowledgeable estimates or purposeful interpretations.

The second benefit is that ABM showed its value for studying the management of SAM from a systems perspective by examining the dynamic interactions among influencing factors. This approach may encourage researchers to develop and fine-tune a conceptual frame and avoid the trap of exhaustively listing factors without defining and understanding their interactions. A clearly defined and limited number of variables allowed concentrating on their mutual interactions. The variable labels could articulate the interactions well (Fig. 16.7). ABM also allowed simulating an agent's adaptation to the environment (i.e. diffusion or network effect). For example, in this study, the higher the proportion of children's carers that were aware of the SAM disease and its treatment, the higher the individual child carer's chance of

awareness and uptake of treatment. ABM is increasingly used for studying health interventions. In addition to our study, a variety of other studies showed its ability to study the interaction among human factors and the intervention (Brookmeyer et al. 2014), such as the level of acceptance of a antiretroviral prophylaxis, the likelihood of contact with a partner, selection of water source (Demarest et al. 2013), or perceived benefits of and barriers to influenza vaccination (Karimi et al. 2015).

The third benefit is that ABM involved expert knowledge in developing a conceptual framework to interpret health systems functioning and related intervention strategies, which is supposedly intuitive for health actors. Hence, ABM can be used as a communication tool to explain the characteristics and effects of interventions in a complex reality (Chiêm et al. 2012). This added value of ABM was clear throughout this study. For example, the specific frame of the agent-based model allowed retrieval and formulation of possible mechanisms behind observed probabilities (e.g. transition rates) in the Markov model. While these frames were critical to anchor the model in an empirical reality, they did not reproduce the entire causal context that drove this transition. Nonetheless, this reality could be obtained by using a logistic regression if data were available. While expert input helped to open the 'black box' to explain transitions (Figs. 16.7 and 16.8), ABM forced the modellers and expert to reconsider each parameter based on its effects. For example, in the first steps of the simulation exercise, no impact of the intervention on child mortality was detectable. This observation raised questions and led to readjustments (e.g. a decrease of the spontaneous recovery rate). Another example was that ABM allowed modelling the effect of spontaneous recovery. While experts know about spontaneous recovery, they ignore its effect in intervention studies because ethical standards require treatment of every sick child detected. Yet, this phenomenon is important to consider when studying the impact of health interventions and defining health intervention strategies with a universal health coverage perspective. Ignoring spontaneous recovery may lead to misinterpreting global indicators of declining mortality, and cost-effectiveness and equity of interventions.

Finally, ABM in this study allowed formulating insights that were not only observational but also operational. The Markov model showed that an intervention on different levers had a synergistic effect when a system was complex. ABM, however, allowed checking the stability of the specified probabilities of transition and showed how modifying one probability of transition affected the flow of children in different health states. Moreover, ABM showed that interventions should target individual factors of service demand- and supply-side barriers simultaneously to improve overall case fatality. Improving community factors alone was inadequate when a health centre did not perform well because trust in the health system may have been lost. It was therefore important to first ensure that the health centre had adequate quality of care and motivated staff before implementing community actions to improve service access and use. On the other hand, the agent-based model was not designed to predict outcome, but rather to support strategic planning of interventions to improve health outcome. Health interventions to improve quality and scale-up of the management of SAM are implemented in various contexts. By testing 'what-if' scenarios, the model may help identify intervention levers needed

to address barriers in implementation or prepare for contingencies. It can also ask questions and solve problems before they arise in real-time. In fact, the cognitive process experienced by the modellers may enhance their skills in representing mechanisms of underlying interventions and ultimately may prepare them to find creative solutions for unpredictable situations (Buzan and Buzan 1993).

16.3.2 Limitations

The study faced limitations that were related to the validation of the model, inherent shortcomings of the model itself, and low level of complexity.

First, the validation process of ABM is subject to much debate (Moss 2008; Moss and Edmonds 2005), but many reasons underpin a need to deviate from typical quantitative statistical validation processes. For example, not all data may be available to assess replication. In this case study, data on the long-term effects of health interventions and individual factors of behaviour were not available. Therefore, it was not possible to test whether the model replicated actual observed trends. Instead, the validation of the model relied on expert knowledge and reproduction of known data, such as case fatality and recovery rates.

These data should be handled with caution as their exact interpretation may conceal critical effects such as the role of spontaneous recovery and defaulting discussed earlier. Moreover, scales with thresholds were preferred to continuous variables, and the interpretation of results was purposively qualitative (trends). In addition, each element of the model (variables, rules) was supported by scientific evidence and expert opinion. Such steps reportedly ensure the validity of the methodology (Lee and Lee 2003). Further validation steps could be envisaged to enhance the scope of the model. For example, the pseudo code in Appendix 3 could be used as a template or tool to adapt the model to a changed context. On the one hand, results from new data may be compared with results from the present model. On the other hand, the model may help target future data collection by providing a predefined conceptual framework that has already yielded dynamic insights into sensitive factors. Another validation step could be to discuss the resulting model with other experts, possibly with other experiences and evaluate the consensus on the model. This process could lead to what is referred to as stakeholder validation (Purnomo et al. 2004).

Second, choices and assumptions used in the modelling process may have limited the scope of model interpretations in terms of SAM interventions. Three important elements of the health interventions were ignored. First, the acute child condition studied used two service delivery platforms, community-based and facility-based primary health care for uncomplicated SAM, and ignored secondary health care for complicated SAM. Second, feedback loops took into consideration interaction among health system functions for the management of SAM but not interactions in the broader health system. Third, children who defaulted and spontaneously recovered or died thereafter were not included in the calculation of overall recovery or case fatality.

The simulation showed that it was only possible to obtain satisfactory results by setting low rates of defaulting (R_3) and case fatality during treatment (R_4). While it was not discussed in the study, this meant that activities to improve compliance with treatment and case fatality during treatment should have been included in the package of interventions (e.g. interventions that improve clinical care, referral, recognition of danger signs and treatment of underlying infections). For example, the simulation did not model active tracing of defaulters, referral to a hospital or management of case overload, which could limit the effect of defaulting and dying in reality. As another example, a strong increase in coverage (as shown in scenario IV, Fig. 16.9) could negatively affect staff motivation (e.g. increasing the risk of work overload) and quality of care (e.g. increasing the risk of treatment interruption because of supply stock outs). The averaging process (of 50 simulations) did not show the multiple possible instances, which may have varied significantly. In most cases the overall situation improved following an intervention, but in a few cases the situation worsened. This may be explained by the use of probabilities and feedback loops. For example, if recovery was especially low (e.g. because of turnover of untrained staff or stockouts of supplies), staff motivation and thus quality of care decreased, and so did community awareness and coverage. As such, the situation could not be remedied by an intervention. These interesting cases could be investigated in future research.

Finally, the present model may appear too simple, relying on two agents defined by three and two variables respectively. However, the corresponding number of interactions generated much speculation. The model showed how *complexity grows exponentially* with the number of variables. Moreover, ABM relies on variables and interactions that may appear abstract. Therefore, insights resulting from the analyses of simulation models require caution in their interpretation. More abstract variables and interactions at this stage could bring even further confusion.

16.3.3 Perspectives

In its present form, ABM would benefit from the involvement of more experts to include other contexts and perceptions. Comparison of parameters elicited from various experts could reveal differences in the conception of transition mechanisms and magnitude of transition rates. Moreover, scenarios of intervention strategies can be imagined and modelled in a virtual environment within seconds. They can be tested extensively, allowing the immediate design of creative and refined solutions. Subsequent analyses and documentation of models could thus lead to intervention guidelines that are anchored in a consensual field reality of these combined expert opinions. The inventory of tacit knowledge complemented by knowledge obtained from this simulation exercises could then be used to inform decision-making for improved intervention strategy designs. Showing possible scenarios using ABM also has the potential to strengthen advocacy with policy makers, planners, and service providers and users.

The model was intended to increase understanding of the successes and failures of interventions to scale up the management of SAM. Adding more health centres with different attributes in an enhanced model would serve this purpose. It would also lead to other questions, e.g., about the effect of competition between health centres or the effect of development partner support on sustainable quality of care and overall health outcomes. As a first step, the model was informed by one health centre with real-time factors with high-burden low-income country values of parameters and scenarios. As a next step, this model could be replicated to simulate scale-up to cover an entire health district, province, or national health system, in which multiple health facilities, health services and health actors interact. In enhanced models, random effects could be introduced and parametric analyses performed with a view to testing the validity and robustness of the model. Moreover, multiplying simulations might allow identification of emergence (the process by which behaviour at a larger scale arises from the detailed structure) and self-organisation, which may be desirable behaviours for improving performance of a health system (Epstein 2008). The scale-up model should be developed through an iterative process eliciting local expert advice. This participatory process could simultaneously help health actors learn to use simulation modelling for decision-making.

16.4 Conclusion

ABM usefully complements the methodological toolkit of health intervention decision makers and planners. In contexts where data are scarce, an agent-based conceptual frame may help list existing scientific evidence, encompass realistic estimates and account for the definition of possible transition mechanisms to simulate data. These simulations may also help target data collection and use new empirical data when available.

Agent-based models representing the health systems may inform intervention strategies by accounting for system effects that could lead to unintended consequences. They may also help design multiple scenarios and test their relevance. Ultimately, they may be used to prevent health systems from collapsing and favour emergent systems behaviour (Begun et al. 2003).

While the health sector is reportedly slow in uptake of complexity (Martin and Sturmberg 2009), involving health actors could help overcome this reluctance. The participatory approach should render the insights of ABM more useful and meaningful (Kamiński and Koloch 2014). Not least, it should trigger better adherence to recommendations resulting from ABM and enhance the skills of health actors to adapt SAM interventions to field realities.

While policy design and strategic intervention planning should be grounded in empirical data, these data do not always provide an interpretation of contextual factors leading to the success of an intervention. Simulation modelling allows consideration of factors that influence the effectiveness of health interventions

for improved SAM service delivery. It also allows identification of barriers and fine-tuning of interventions to overcome them before implementation, which may prevent failure and maximise the chance of success.

References

- Atun, R. (2012). Health systems, systems thinking and innovation. *Health Policy and Planning*, 27(suppl 4), iv4–iv8.
- Becart, E. (2014). *Meta-analysis of barriers and boosters from 78 coverage assessments supported by the CMN*. London: Coverage Monitoring Network.
- Begun, J. W., Zimmerman, B., & Dooley, K. (2003). Health care organizations as complex adaptive systems. *Advances in Health Care Organization Theory*, 253, 288.
- Black, R. E., Victora, C. G., Walker, S. P., Bhutta, Z. A., Christian, P., De Onis, M., et al. (2013). Maternal and child undernutrition and overweight in low-income and middle-income countries. *The Lancet*, 382(9890), 427–451. doi:10.1016/S0140-6736(13)60937-X.
- Bousquet, F., Castella, J.-C., Trébuil, G., Barnaud, C., Boissau, S., & Kam, S. P. (2007). Using multi-agent systems in a companion modelling approach for agroecosystem management in South-East Asia. *Outlook on Agriculture*, 36(1), 57–62.
- Brailsford, S. (2005). Overcoming the barriers to implementation of operations research simulation models in healthcare. *Clinical and Investigative Medicine*, 28(6), 312.
- Brookmeyer, R., Boren, D., Baral, S. D., Bekker, L.-G., Phaswana-Mafuya, N., Beyrer, C., et al. (2014). Combination HIV prevention among MSM in South Africa: Results from agent-based modeling. *PLoS One*, 14, 9(11). doi:10.1371/journal.pone.0112668. eCollection 2014.
- Buzan, T., & Buzan, B. (1993). *The mind map book. How to use radiant thinking to maximise your brain's untapped potential*. New York: Plume.
- Chiêm, J. C., Macq, J., & Speybroeck, N. (2012). Rule-based modeling of chronic disease epidemiology: Elderly depression as an illustration. *PLoS One*, 7(8), e41452.
- Claxton, K., Sculpher, M., & Drummond, M. (2002). A rational framework for decision making by the National Institute for Clinical Excellence (NICE). *The Lancet*, 360(9334), 711–715.
- Collins, S., Dent, N., Binns, P., Bahwere, P., Sadler, K., & Hallam, A. (2006). Management of severe acute malnutrition in children. *The Lancet*, 368(9551), 1992–2000.
- Collopy, F., & Armstrong, J. S. (1992). Rule-based forecasting: Development and validation of an expert systems approach to combining time series extrapolations. *Management Science*, 38(10), 1394–1414.
- Deconinck, H., De Man, J., & Macq, J. (2015). *Health systems strengthening for integrating and scaling up management of severe acute malnutrition* (Report of a literature review). Brussels: UCL.
- Demarest, J., Pagsuyoin, S., Learmonth, G., Mellor, J., & Dillingham, R. (2013). Development of a spatial and temporal agent-based model for studying water and health relationships: The case study of two villages in Limpopo, South Africa. *Journal of Artificial Societies and Social Simulation*, 16(4), 3.
- Edmonds, B. (2010). Bootstrapping knowledge about social phenomena using simulation models. *Journal of Artificial Societies and Social Simulation*, 13(1), 8.
- Epstein, J. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), 12.
- Forrester, J. W. (1971). Counterintuitive behavior of social systems. *Theory and Decision*, 2(2), 109–140.
- Garenne, M., Willie, D., Maire, B., Fontaine, O., Eeckels, R., Briand, A., et al. (2009). Incidence and duration of severe wasting in two African populations. *Public Health Nutrition*, 12(11), 1974–1982.

- Gilbert, N., & Terna, P. (2000). How to build and use agent-based models in social science. *Mind & Society*, 1(1), 57–72.
- Joint Learning Initiative. (2004). *Human resources for health: Overcoming the crisis* (Global Equity Initiative, Harvard University). Washington, DC: Communications Development Incorporated.
- Jones, K., & Berkley, J. (2013). Severe acute malnutrition and infection. *Technical brief: CMAM Forum*.
- Kamiński, B., & Koloch, G. (2014). *Advances in social simulation*. Heidelberg: Springer.
- Karimi, E., Schmitt, K., & Akgunduz, A. (2015). Effect of individual protective behaviors on influenza transmission: An agent-based model. *Health Care Management Science*, 18(3), 1–16.
- Lee, K. C., & Lee, S. (2003). A cognitive map simulation approach to adjusting the design factors of the electronic commerce web sites. *Expert Systems with Applications*, 24(1), 1–11.
- Lempert, R. (2002). Agent-based modeling as organizational and public policy simulators. *Proceedings of the National Academy of Sciences*, 99(suppl 3), 7195–7196.
- Luke, D. A., & Stamatakis, K. A. (2012). Systems science methods in public health: Dynamics, networks, and agents. *Annual Review of Public Health*, 33, 357–376.
- Macq, J., & Chiem, J. C. (2011). Looking at the effects of performance-based financing through a complex adaptive systems lens. *Bulletin of the World Health Organization*, 89(1), 699–700. doi:10.2471.
- Mainz, J. (2003). Defining and classifying clinical indicators for quality improvement. *International Journal for Quality in Health Care*, 15(6), 523–530.
- Marshall, D. A., Burgos-Liz, L., IJzerman, M. J., Osgood, N. D., Padula, W. V., Higashi, M. K., et al. (2015). Applying dynamic simulation modeling methods in health care delivery research—the SIMULATE checklist: Report of the ISPOR Simulation Modeling Emerging Good Practices Task Force. *Value in Health*, 18(1), 5–16.
- Martin, C. M., & Sturmberg, J. P. (2009). Perturbing ongoing conversations about systems and complexity in health services and systems. *Journal of Evaluation in Clinical Practice*, 15(3), 549–552.
- Moss, S. (2008). Alternative approaches to the empirical validation of agent-based models. *Journal of Artificial Societies and Social Simulation*, 11(1), 5.
- Moss, S., & Edmonds, B. (2005). Sociology and simulation: Statistical and qualitative cross-validation. *American Journal of Sociology*, 110(4), 1095–1131.
- ONeil, M., & Reimann, S. (2013). Strengthening human resource management to improve health outcomes. *eManager*, 1, 1–17.
- Purnomo, H., Mendoza, G., & Prabhu, R. (2004). Model for collaborative planning of community-managed resources based on qualitative soft systems approach. *Journal of Tropical Forest Science*, 16(1), 106–131.
- Railsback, S. F., & Grimm, V. (2011). *Agent-based and individual-based modeling: A practical introduction*. Princeton: Princeton University Press.
- République du Niger Institut National de la Statistique. (2013). Rapport de l'enquete nutrition et survue des enfants de 6 à 59 mois. Niamey: INS.
- Rogers, E., Myatt, M., Woodhead, S., Guerrero, S., & Alvarez, J. L. (2015). Coverage of community-based management of severe acute malnutrition programmes in twenty-one countries, 2012–2013. *PLoS One*, 10(6), e0128666.
- Salerno, F., Cuccillato, E., Caroli, P., Bajracharya, B., Manfredi, E. C., Viviano, G., et al. (2010). Experience with a hard and soft participatory modeling framework for social-ecological system management in Mount Everest (Nepal) and K2 (Pakistan) protected areas. *Mountain Research and Development*, 30(2), 80–93.
- Scaling Up Nutrition Movement. (2010). *Scaling up nutrition: A framework for action*. Geneva: United Nations SUN Movement.
- Silverman, B. G., Hanrahan, N., Bharathy, G., Gordon, K., & Johnson, D. (2015). A systems approach to healthcare: Agent-based modeling, community mental health, and population well-being. *Artificial Intelligence in Medicine*, 63(2), 61–71.

- Sonnenberg, F. A., & Beck, J. R. (1993). Markov models in medical decision making a practical guide. *Medical Decision Making*, 13(4), 322–338.
- Swanson, R. C., Cattaneo, A., Bradley, E., Chunharas, S., Atun, R., Abbas, K. M., et al. (2012). Rethinking health systems strengthening: Key systems thinking tools and strategies for transformational change. *Health Policy and Planning*, 27(suppl 4), iv54–iv61.
- The Sphere Project. (2011). *Humanitarian charter and minimum standards in humanitarian response*. Rugby: Practical Action Publishing.
- United Nations Children’s Fund. (2013). *Nutrition dashboard (NutriDash): The 2013 global report on the pilot year*. New York: UNICEF.
- Witter, S., Toonen, J., Meessen, B., Kagubare, J., Fritsche, G., & Vaughan, K. (2013). Performance-based financing as a health system reform: Mapping the key dimensions for monitoring and evaluation. *BMC Health Services Research*, 13, 367. doi:10.1186/1472-6963-13-367.
- World Health Organization. (2006a). *Quality of care: A process for making strategic choices in health systems*. Geneva: WHO.
- World Health Organization. (2006b). *The world health report: 2006: Working together for health*. Geneva: WHO.
- World Health Organization. (2008). *The world health report (Primary health care: now more than ever)*. Geneva: WHO.
- World Health Organization. (2013). *Guideline: Updates on the management of severe acute malnutrition in infants and children 2013*. Geneva: WHO.
- World Health Organization, & United Nations Children’s Fund. (2009). *WHO child growth standards and the identification of severe acute malnutrition in infants and children (A joint statement)*. Geneva: WHO.
- World Health Organization, World Food Programme, UNSCN, & United Nations Children’s Fund. (2007). *Community-based management of severe acute malnutrition (A joint statement)*. Geneva: WHO.
- World Health Organization, United Nations Children’s Fund, & The World Bank Project. (2012). *Global database on child growth and malnutrition (Joint child malnutrition estimates: levels and trends)*. Geneva: WHO.
- World Health Organization, United Nations Children’s Fund, World Food Programme, & 1000 Days. (2014). *Global nutrition targets 2025 (Wasting policy brief)*. Geneva: WHO.
- Zinszer, K., Charland, K., Kigozi, R., Dorsey, G., Kanya, M. R., & Buckeridge, D. L. (2014). Determining health-care facility catchment areas in Uganda using data on malaria-related visits. *Bulletin of the World Health Organization*, 92(3), 178–186.

Chapter 17

Exploring Contingent Inequalities: Building the Theoretical Health Inequality Model

Michael Wolfson, Steve Gribble, and Reed Beall

“an aggregate relation between income inequality and health is not necessary — associations are contingent”

(Lynch et al. 2000)

17.1 Introduction

Income inequality is pervasive and generally increasing over recent decades in most countries of the world. At the same time, and clearly in modern developed societies, there is a pervasive individual-level gradient in health that is characterized by large variations in health status and longevity across members of a population which are strongly (negatively) correlated with individuals' incomes and other measures of their socio-economic status.

Figure 17.1 juxtaposes two sets of data from the US in 1991. The downward sloping curve shows the estimated individual-level relative mortality risk on longitudinal follow-up as a function of income. (The closely adjacent dashed lines show 95 % confidence intervals.) The humped curve is the distribution of individuals by income. The reasons for this downward sloping income-mortality association are contested, with some suggesting “reverse causality”, where the main causal pathway runs from poor health to lower income. However, Wolfson et al. (1993) provide strong evidence that (at least for Canadian males) the majority of this association in Canada is causal, from income to health rather than the reverse.

This research was developed as part of the NIH-funded Network on Inequality, Complexity and Health (NICH) led by George Kaplan (NIH 2010). We are deeply indebted to NICH colleagues for many valuable discussions, and to Nate Osgood and Chris Dutchyn for invaluable collaboration in refining the algorithm specifications throughout THIM. Much of the NICH work is published in Kaplan et al. (2016).

M. Wolfson (✉) • S. Gribble • R. Beall

University of Ottawa, Ottawa, ON, Canada

e-mail: mwolfson@uottawa.ca; steve.gribble@hotmail.com; reedbeall@gmail.com

© Springer International Publishing Switzerland 2017

A. Grow, J. Van Bavel (eds.), *Agent-Based Modelling in Population Studies*,
The Springer Series on Demographic Methods and Population Analysis 41,
DOI 10.1007/978-3-319-32283-4_17

487




Fig. 17.1 Income distribution and mortality risk, US 1991 (This and all other figures except 17.2 from Wolfson 2016, used with permission)

At the same time, Fig. 17.1 provides one basis for observing not only an individual-level association between income and health (whether indicated by mortality rates or some other measure of individuals' health), but also at a group level between income inequality and health. The idea is that because the individual-level income gradient is non-linear, as individuals become more spread out along the horizontal axis – i.e. as income inequality increases – average health (however measured) necessarily must decline. Based on this well-known mathematical property, Gravelle (1998) argued that the fact that higher income inequality is observed to be associated with higher mortality is nothing more than a statistical artefact.

While Gravelle's argument is logically correct, Wolfson et al. (1999) showed that something further must be at work; the observed patterns in the curvilinear relationship shown in Fig. 17.1 by itself was nowhere near sufficient to account for the observed correlation between income inequality and mortality among the 50 US states. Further, in a multi-level analysis, Backlund et al. (2007) showed that inequality itself was significantly correlated with mortality among the 50 US states, even after controlling for individual-level income as well as a range of covariates, including race at both the individual and state level. Thus, statistical artefact alone is not enough to explain the relationship; something further must be at work.

Provocatively, Fig. 17.2 (Ross et al. 2000) with data for the US and Canada, suggests that the inequality – average health correlation is contingent. There is a strong correlation between a measure of city-level¹ income inequality (measured by

¹To be precise, the US data are referring to Standardized Metropolitan Statistical Areas (SMSAs), with the data for the other countries specifically constructed by the authors to be as close as possible in concept and definition.




Fig. 17.2 Working age mortality and median shares of income for US and Canadian metro areas (Used with permission, Ross et al. 2000)

the “median share” = share of income accruing to the bottom half) and working age mortality rates in the US and the UK. However, in Canada, Australia and Sweden, there was no observed correlation.

These clear differences in the strength of association suggest that there are other factors influencing the relationship – in other words, this evidence suggests that the income inequality-health correlation is contingent on a range of country-specific factors.²

There are a number of candidate factors which might account for the different patterns of association across the five countries. In the absence of much of the relevant data, THIM (Theoretical Health Inequality Model) was constructed to explore this contingent correlation, with a focus on Canada and the US. THIM can be considered a quasi-theoretical model – it is highly abstract and simplified, hence theoretical. But it is also constructed both to build on observed empirical

²No analysis has been done to compare the strength of the Gravelle artefact hypothesis across countries. However, given its rather small effect among US states as estimated by Wolfson et al. (1999), any such differences are unlikely to be material.

patterns and relationships, and to generate quantitative results whose patterns, while not exact, could be empirically verified were the requisite data available.

One hypothesis regarding the observations in Fig. 17.2 is that the main factor accounting for the existence of a correlation between income inequality and mortality at the city level in the US but not in Canada is the nature of economic segregation within metropolitan areas. For example, the impression is that US cities more often have gated communities where the very rich live, and poor ghettos, such that disparities in average income between these different neighbourhoods are greater than in Canada, while income heterogeneity within US neighbourhoods is less than in Canada. In the US this segregation is most visibly associated with race. However, it is also strongly associated with income. Further, when data disaggregated by race in the US are further broken down by income or other measures of socio-economic status (SES), much of the observed racial disparity is seen to be associated with SES disparities (Braveman et al. 2010). In Canada, income is often a factor in neighbourhood segregation, for example related to housing prices, while race is not. Thus, in the design of THIM, we felt it sufficient to leave out race and focus on income segregation.

Of course, there are many other candidate factors that could account for the differences in the observed health-inequality correlation for Canadian and US cities shown in Fig. 17.2. These include differences in the roles played by education and parental position in transmitting social advantage and disadvantage from one generation to the next. The development of THIM therefore reflects a process of constructing a formal theory to inform our intuition and to determine whether a set of plausible factors could possibly account for the observations in Fig. 17.2, where this formal theory takes the form of algorithms in a computer simulation model, plus a suite of parameters which appeal to the “stylized facts”.

A fundamental challenge at the outset of constructing such a model is the degree of abstraction. Generally, simpler models are easier to build, understand, and analyse. But simplicity comes at a price – too much and the theory fails to capture obviously important realities, and risks becoming little more than an entertaining story. A founder of cybernetics, Ross Ashby (1958) coined the term “requisite variety”, and in the context of ABM development, we like to adapt this to “requisite complexity”: we want sufficient complexity, richness, and detail in our theory to encompass the key attributes and processes that judgement, experience, and evidence suggest are central to the topic of interest, but no more. As in the quote attributed to Einstein, “keep it simple, but not too simple.”

To this end, THIM uses agent-based modeling (ABM). With ABM it is possible to include agent heterogeneity, multiple levels of aggregation, dynamics without any a priori assumptions of equilibrium or stability, lagged endogenous effects, agent interactions, all kinds of feedbacks and reciprocal causation, non-parametric specifications for the functional forms of various distributions, and richly textured stochasticity to reflect uncertainty inclusive of flexible notions of the “shapes” of the distributions of the random variables. Other methodologies, including compartment or differential equation models, spreadsheets, life tables, Markov models and System Dynamics models are unable to include all of these features.

THIM incorporates simplified and stylized, yet plausible empirically-based individual-level relationships, among health status, education, income, mortality rates, and neighbourhood mobility. Mobility then in turn affects the extent of neighbourhood income segregation. THIM further has multiple levels – individuals, parent-child dyads, neighbourhoods, and cities – since these are essential to theorizing about the roles of parental transmission of SES and health advantage to their children, the impacts of average neighbourhood incomes on schooling, and overall city-wide patterns of inequality and mortality – among others.

17.2 THIM Overall Structure

We begin describing this network of hierarchical relationships in THIM with its population, consisting of individual unisex agents (called “sims”). They are born to parent sims, live in neighbourhoods, receive an education, and then start earning income. They start life in perfect health, but their health gradually declines as they age, with the pace of these declines affected by their income. They then face mortality rates (death) that depend on their age, health, and income.

Each sim has a number of characteristics, attributes, or “state variables”. The relationships among these variables is multi-level, as indicated in the following three figures and Table 17.1 following. At the first individual level, these relationships are shown in Fig. 17.3.

The arrows indicate that each sim starts life with a level of education (E). “Education” is actually a marker or proxy for a confluence of factors including formal schooling, the informal “home curriculum”, innate ability, and neighbourhood barriers (e.g. presence of youth gangs) and facilitators (e.g. easily accessible public libraries). As shown by the arrow in Fig. 17.3 from E to Y, education is assumed to influence an individual sim’s income (Y) directly, but then to influence health status (H) and mortality (D = death) indirectly via income. Income also influences residential location (L) amongst neighbourhoods (see below).

Beyond these individual-level relationships, individual sims in THIM have relationships both with their parents, and with other sims in their neighbourhoods (abbreviated “nbhd” in the diagram) and cities. These multi-level relationships give rise to additional causal influences as shown in Figs. 17.4 and 17.5. In particular, there are parent-to-child effects via Y (income) of the parent affecting their child’s education (E) and income (Y), both directly and indirectly as the child sim’s education affects its income.

THIM also posits neighbourhood level effects where average neighbourhood income affects sims’ education, and income, and then indirectly health, and mortality. These posited neighbourhood effects reflect, for example, the influences of a neighbourhood’s affluence or poverty on the quality of the school system, and the peers and social milieux to which the children are exposed. Further, while not easily shown in Fig. 17.4, all of these state variables and their interacting effects co-evolve through time, indicated by the arrow in the upper-right.

Table 17.1 Summary characteristics of main variables and events in THIM

Variable or event	Units	Timing	Explanatory variables by level of aggregation				Randomness	Other parameters
			Self	Parent	Neighbourhood	City		
Education (E)	Years	At birth		Income relative to neighbourhood	Income relative to city	Overall average income	Normal	
Potential income (YBase)	Positive real	At birth	E	Income relative to neighbourhood	Income relative to city	Overall average income	Positively skewed	
Income (Y)	\$	Birthday	YBase, E				Lognormal	Age-income profile
Health (H)	Real in [0,1]	Birthday	Y			Average incomes near age	Negatively skewed	
Giving birth	Time	Continuous					Constant hazard	Fertile age range
Death (D)	Time	Continuous	H, Y			Average health, average incomes near age	Age-specific hazard	Survival curve
Neighbourhood mobility	Neighbourhood index	Birthday	Y		Average Income		Uniform	

Source: Wolfson (2016), with permission

Fig. 17.3 A sim's attributes and their relationships

E = education
 Y = income
 H = health
 D = death
 L = location

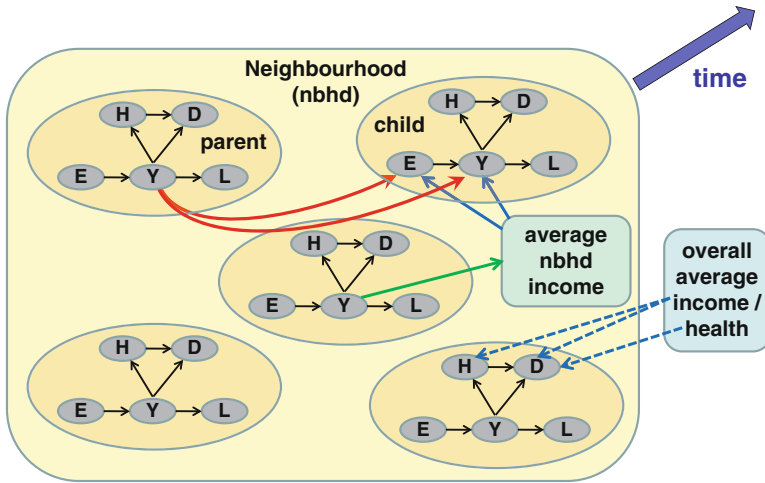



Fig. 17.4 Sims in a neighbourhood

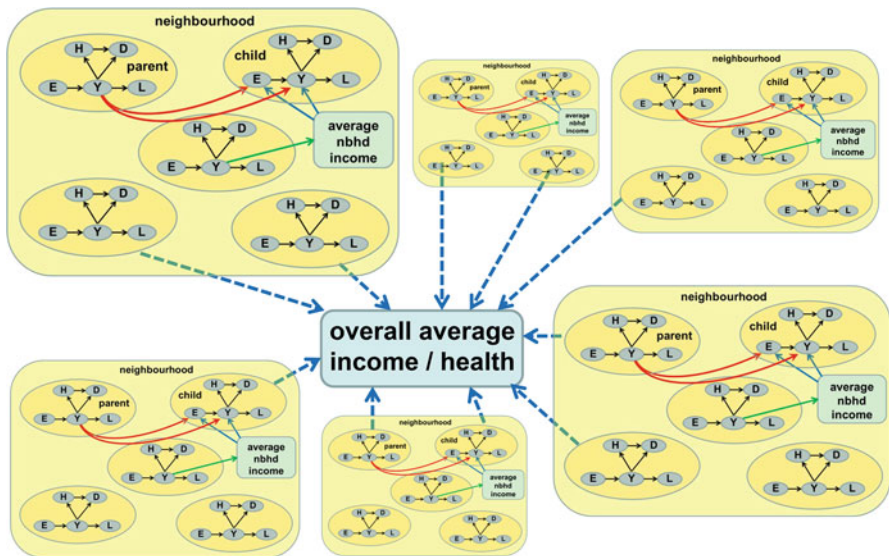


Fig. 17.5 A THIM City with multiple neighbourhoods, and multiple sims within each neighbourhood

Finally, there is a third level, the society or city, comprised of multiple neighbourhoods (Fig. 17.5). One key factor at this broader level is geographic mobility. Analogous to observed behaviour in the US and Canada, adult sims are posited to have a general propensity (of strength which can be varied with parameters) to move to a neighbourhood where their own income is closer to that neighbourhood's average (in a phrase, income homophily). Depending on this propensity, neighbourhoods can either be quite heterogeneous in terms of residents' incomes and levels of education (weak homophily), or more segregated with neighbourhoods more internally homogenous, but more polarized comparing one to another (strong homophily, e.g. more "gates (for gated communities) and ghettos" as in the US compared to Canada).

Table 17.1 provides another perspective on the relationships shown in Figs. 17.3, 17.4, and 17.5 above. The first column lists the five main variables that are attributes of each and every sim, as shown in Fig. 17.3 above (E, Y, H, D, and L), plus one intermediate variable (YBase), as well as the process of giving birth. The next two columns indicate the units of measurement for the variable or event, and its timing. As will be discussed in the following section, timing is different for the different variables and events.

In a simulation model like THIM, everything must be self-contained, so every variable and event must be endogenous. The four columns under "explanatory variables by level of aggregation" recapitulate the arrows shown in Figs. 17.3, 17.4, and 17.5 by showing the variables posited in THIM to explain or determine the given variable or event. Since a core objective of THIM is to theorize about the inter-relationships of real world processes occurring at different levels of aggregation, these explanatory variables are in turn defined at different levels of aggregation, from the sim's own attributes to those of its city.

As widely observed in the real world, and found ubiquitously in statistical analysis of real world phenomena, especially in social science and epidemiology, statistical relationships always include substantial "unexplained" variance. THIM reflects this reality by incorporating a substantial variety of randomness, indicated in the penultimate column of the table. While this randomness is realistic and unavoidable, it is possible to know something about its form or "shape", and this is indicated by the entries in this column.

Finally, in some cases it is useful to appeal to some broadly observed relationship in the form of a "stylized fact", such as the typical age profile of income ("humped") or the age profile of mortality rates ("accelerating"). These are indicated in the last column.

Given this multi-level structure of posited causal relationships, the dynamics of each of a given sim's attributes, and those of the population sub-groups and groups of which it is a member, co-evolve endogenously. This theoretical structure is designed to strike a balance: to be as simple as possible so that THIM's dynamics will be reasonably understandable, but at the same time to be rich enough that the phenomena of interest can be represented at least somewhat realistically, albeit in simplified and stylized form – i.e. to have requisite complexity, but no more.

The following sections provide more detail and intuition regarding each of these processes.

17.3 Timing

THIM has been designed as a discrete *event* model, with continuous time. Continuous time formulations of agent-based models are fundamentally more realistic than discrete time for relevant processes.

For THIM, we have used a mixture of discrete and continuous time formulations for the various dynamic processes. The two major (discrete) demographic events – birth and death – are characterized by hazard rates and can occur at any moment in (continuous) time. For simplicity, however, other events are posited to occur at fixed times. In general, the unit of time is the “year”, and THIM defines two discrete events that occur at regular annual intervals: (1) the start of a new calendar year, and (2) for any given sim, its birthday, though birthdays themselves occur at real valued times throughout each year in model time, since birth itself is a continuous time process.

The remaining events that characterize sims’ dynamics and the dynamics of their neighbourhoods and city are then “hooked” to these two discrete “trigger” events – a new year and a birthday. Specifically, a sim’s annual flow rate of income (Y), its level of health (H), and its neighbourhood location (L) are fixed for 1 year at a time, and can only change at the sim’s birthday. As is characteristic of discrete time models, the ordering of these events is arbitrary, but it must be clearly specified because changing the ordering will change the simulation results. THIM posits that this order is income, health, and then location.³

At the neighbourhood and city level, various forms of average income and average health enter into the equations for each sim’s attributes (E , Y , H , L , and D per Figs. 17.3, 17.4, and 17.5 above). While these averages are changing continually over time (whenever any sim in the relevant population has a birthday), THIM posits a “statistical office” that takes a snapshot of the relevant income flow rates $\{Y_i\}$ and levels of health $\{H_i\}$ (where i indexes sims) an instant before the new year begins. These snapshot values are then the basis for computing the required averages which in turn are used as inputs to various algorithms for individual sims’ dynamics over the ensuing calendar year.

³If A affects B, and B affects A, then a different value for A will result in a different effect on B, and a different value for B will result in a different effect on A. Thus, if transitions in A are simulated in discrete time before the probabilities of transitions in B are computed, the possible outcomes for A and B can be different than if B’s transitions are simulated first. For example, in THIM, if a change in location were simulated before changes in income, there could be a different move than if income change were simulated first.

17.4 Demography and Birth

THIM demography involves birth of a sim at age $a = 0$ (in continuous time), a period of “schooling”, followed by a period during which the sim can give birth to another sim, and ultimately death. There is also an initial population, and a fertility rate.

The demographic parameters for the THIM population are

- Starting number of individuals (e.g. 50,000)
- Minimum age to reproduce (e.g. 20)
- Maximum age to reproduce (e.g. 40)

While reference is made to “adults” and “children” in THIM, these categories are not defined explicitly in the model; rather an adult is implicitly any sim with age $a \geq$ minimum age to reproduce; correspondingly, a “child” is any sim with $a <$ minimum age to reproduce.

The maximum (continuous) age of a sim is assumed to be <100 ; i.e. sims die with 100 % probability by the instant of their 100th birthday. The algorithm for time of death is described in the Mortality section below.

New “baby” sims are born at a constant fertility rate to “eligible” (unisex) parent sims in their “fertile” age range (e.g. exact ages 20–40). (A sim can have multiple children during this reproductive period.) A constant fertility rate is determined endogenously just before the simulation begins, and is defined as the one that will result in a constant population for the given mortality rate schedule and fertile age range.

This fertility rate calculation assumes there are no other influences on mortality rates. However, THIM is constructed such that other factors may indeed be at play, specifically the impacts of H and Y on chances of dying (see Mortality below). As a result, the population in THIM may increase or decrease over the course of a simulation, depending on the parameters used to set the strengths of the H and Y effects on mortality. Hence, there is no a priori assumption in THIM of a steady-state or even asymptotic population size and distribution by age. These demographic characteristics emerge from any given simulation.

17.5 Education

Education is measured by the last year of age when the sim is still “in school” (an integer). This “amount” of education, E, is in a specified range, such as 1–20 (both parameters). Note that “education” E can start as early as at birth (age = 0) because it is intended as a marker for an entire suite of early developmental influences, including parenting, formal schooling, home curriculum, and neighbourhood socialization.

Even though this abstraction for the number of years of “education” is conceptualized as being broader than the more conventional measure, years of schooling, THIM posits determinants of educational attainment that are similar to those in the empirical literature. Specifically, the amount of education E attained by a child sim is assumed to depend log-linearly on the relative income of the parent and log-linearly on its average neighbourhood income relative to that of the city overall. Thus, sims born in neighbourhoods with higher (lower) relative average income will tend to have higher (lower) ages of school completion. And children of parents with higher (lower) incomes relative to the average of their neighbourhood have higher (lower) ages of school completion.

Further, education is assumed to be “pre-ordained” at birth, and is not acting on/affecting/affected by anything during the educational period. This is clearly a simplification, but is designed to allow a focus on the impacts of early and adolescent life events on the rest of the life course. In other words, we judged that the critical feature was how education was affected by the parent and neighbourhood socio-economic context at birth only; i.e. there would be no important loss in the theory if we ignored how education evolved between the times of birth and school-leaving.

Given this assumption that education is “pre-ordained”, E to be attained is set for each newborn sim at age $a = 0$, and then fixed for the rest of life. The inputs used to determine E are thus computed just once at the time of birth. These inputs are as follows:

- The parent’s income in dollars at the moment of the sim’s birth
- Average income for all “adult” sims in the given sim’s parent’s neighbourhood for the most recent complete calendar year (recall the section on Timing above)
- Average income for all “adult” sims across all neighbourhoods, also for the most recent calendar year

17.6 Income

Given demography and education, the next main process in THIM (conceptually speaking) is modeling income. In principle, income dynamics can be partitioned into several components, specifically (a) some notion of unobserved but highly important innate ability, (b) “permanent income” (i.e. longer term average income,), (c) a characteristic age-income profile, and (d) transitory income. THIM’s formulation, while simplified and abstract, builds on these ideas.

While being educated, sims have no income; they are assumed to start receiving income at the birthday when they complete school at age $a = E$. Once they do begin receiving income at age E and then throughout life, THIM posits that their incomes are positively correlated with their relative educational attainment, relative parental income, and relative neighbourhood income. Further, incomes follow a stereotypical age profile, and are distributed in a characteristically positively skewed manner.

More specifically, THIM generates income levels during a given year of age, and their dynamics from one birthday to the next, according to the following steps.

First, as soon as a sim's educational attainment, E , is determined (i.e. at birth), the sim is assigned (also at birth) a relative potential lifetime income index Y_{Base} – for example a value of 1.57 or 0.83 indicating incomes over the life course that will be 57% higher or 17% lower than the city- (or society-) wide (age-specific) average, respectively. This value is determined, as a first step, by a random number drawn from a user-specified positively skewed probability distribution (hence area under the curve = 1.0) with mean = 1.0 and taking only positive values.

This resulting Y_{Base} thus corresponds to an index representing a mix of innate ability and permanent income, where permanent income is here defined in terms of percentile rank in the distribution rather than in dollar terms. Y_{Base} embodies both an a priori quantification of the degree of inequality in this distribution as well as the influences of a set of factors generally accepted as empirically important – parental income, own education, and neighbourhood average income – all relative to a broader average.


To the extent that Y_{Base} is interpreted as innate ability, it may appear strange that its degree of inequality may be varied. However, the stochastic process generating Y_{Base} in THIM is intended to reflect both dispersion in innate ability and other unobserved “environmental” influences on future income potential. THIM is effectively positing that these influences can vary across societies in general, and cities specifically in THIM, and that the foundational dispersion reflected in the distribution of Y_{Base} is in substantial part a social construction.

The second main step in THIM is using Y_{Base} , as endowed at birth and fixed throughout life, in order to generate a dollar income Y , as an annual flow rate, constant between each birthday at or after exact age $a = E$. For this step, THIM embodies two further stylized facts. One is that incomes tend to follow a characteristic life course pattern, often summarized in an age-income profile. The other is that incomes are observed to vary from one year to the next in a way similar to the outcome of memoryless random shocks, where these shocks are positively skewed (Luong and Hébert 2013; Johnson and Neumark 1996; Murphy and Welch 1990).

As a result, THIM updates Y at each successive birthday event by following a stereotypical age-income profile, which is modified by random lognormal proportional shocks. Before age $a = E$, dollar income Y is kept at zero.

This stereotypical average age-income profile (ignoring the fact that income is always zero for age $a < E$) is assumed to take the general form of a piecewise linear function, as shown in Fig. 17.6. Plausible hypothetical values would be \$0 at age zero, \$35,000 at age 30, \$55,000 at age 55, \$25,000 at age 65, and \$20,000 at age 100 (Luong and Hébert 2013; Johnson and Neumark 1996; Murphy and Welch 1990). This function is fully determined by a pair of vector parameters – the points along the horizontal age axis, and the dollar levels of these incomes at each of these ages.

Fig. 17.6 Simplified stereotypical age income profile



We then derive the sim’s dollar income level Y_a for each age $a \geq E$ at each birthday using YBase combined with the posited average age-income profile, and a lognormal random term.

With this formulation, which combines a constant lifetime YBase endowed at birth with a stereotypical age-income profile and yearly random shocks, THIM (with appropriate parameter settings) should be able to reproduce the strong autocorrelations typically observed over time in actual individual-level income trajectories. It should also be able to reproduce observed patterns of income inequality (e.g. Gini coefficients), and of course patterns of income over the lifecycle, since an average age-income profile is an explicit input parameter.

17.7 Health

In popular media, health is most often discussed in terms of diseases. However, in the health economics literature, especially discussions of cost-effectiveness, health is often discussed in terms of QALYs = quality-adjusted life years. A QALY is an index number, usually in the [0,1] interval, where zero is equivalent to being dead, and 1 is considered full health. For THIM, we have built on this health economics literature and posited that health can be described like a QALY as a real number H in the [0,1] interval. THIM further assumes that all sims start life at age $a = 0$ in full health, with $H = 1$, but their health starts changing at their first birthday (i.e. $a = 1$), and changes thereafter at each birthday.

There is a considerable literature where specific QALY indices have been examined empirically, including in longitudinal surveys like the National Population Health Survey (Statistics Canada, NPHS) using the McMaster Health Utility Index (Feeney et al. 2001). Based on these kinds of analyses, health does not decline inexorably with age at an accelerating pace. Rather there are ups and downs, though generally superimposed on an increasingly downward trend (Kaplan et al. 2007). THIM therefore posits that H follows a biased random walk, more down than up.

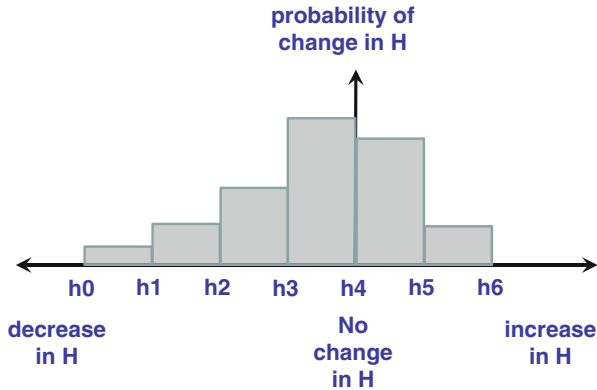


Fig. 17.7 Hypothetical distribution of health changes

In turn, changes in H are influenced by two factors – age and income: higher age predisposes to larger steps in the random walk, and higher income predisposes to less down and more up in these steps. The character of these influences can be controlled by a THIM user via several parameters.

More specifically, this process of determining a sim's change in health occurs on each birthday. It starts with a draw from a negatively skewed probability distribution, such that there is a chance that health can go up (e.g. as a result of cataract surgery or a joint replacement) as well as down (e.g. as a result of continuing cognitive decline associated with dementia, or from a heart attack) from age $a-1$ to a , but it is more likely to decrease. To provide flexibility in the specification of this random walk, THIM assumes the simple yet flexible (non-parametric) piecewise constant probability density histogram shown in Fig. 17.7. With this basic probability density for defining the change in health at each age, the actual change is based on two further steps.

First, a further change in health is added depending on the sim's income. This dependence on income is based on the ubiquitous finding, illustrated in Fig. 17.1 above, that health (as well as mortality) is positively (negatively for mortality) and significantly correlated with income, and this relationship is substantially causal (Wolfson et al. 1993). While there is some debate, there is also strong evidence that the relationship is with relative rather than absolute income. In other words, what matters is not the dollar level of an individual's income, but rather the ratio of the individual's income to the average of his or her peers.

The social epidemiology literature also shows important correlations, likely causal, of health with other key variables like educational attainment (Braveman et al. 2010). It would of course be possible to posit that in THIM, the random walk for H would also be influenced by E . However, E can already affect changes in H indirectly via Y , so it was judged that this kind of added complexity in the model

was not warranted; it would not add appreciably to the insights likely to be gained by experimenting with THIM.⁴

As a result, THIM uses the sim's income relative to the average income of all sims in the city "near" in age as the relevant variable to affect the steps in the random walk of H. The restriction to nearby ages is included in THIM due to the character of the typical age-income profile, where sims at higher ages have much lower incomes than those of middle age. If an elderly sim's income were compared to the average income of all "adult" sims, it would almost always be significantly below average. It is better therefore to use an income relative to those of similar age.

17.8 Mortality

The THIM mortality algorithm starts with a standard age-specific mortality schedule. This schedule is shown in Fig. 17.8.

Specifically, the mortality schedule is represented for the THIM user by a simple piecewise-linear function like that on the left of Fig. 17.8 with six cut points (e.g. at ages {0, 20, 40, 60, 85, 99, 100}) and corresponding mortality rates (e.g. {0, .001, .005, .025, .13, 1}). These specific mortality rate parameters correspond to a life expectancy of 77 years. From this simplified mortality schedule, THIM then creates a vector of 100 constant annual mortality rates for intervening ages by linear interpolation. In other words, internally THIM is using piecewise constant hazards by single year of age. The simplified schedule with six intervals is a convenience for users to make plausible patterns of age-specific mortality rates easier to input.

Given these mortality rates, a new tentative $D = \text{age at death} \geq a$ for the given sim is generated randomly at each birthday, based on the constant hazard h_a for




Fig. 17.8 Baseline piecewise linear age-specific mortality rates and corresponding survival curve (a) Mortality rates by age (b) Survival by age

⁴Of course, other THIM users might judge differently in this case and many others. If so, the source code is available and these kinds of modification to THIM can be made given some programming knowledge and experience.

each integer age. If $D > a + 1$, no death occurs during that year of age, and a new tentative D is drawn at the next birthday. Otherwise the sim dies at the randomly drawn (continuous) time in the $[a, a + 1]$ age interval.

17.9 Neighbourhood Mobility

Central to THIM is that sims live in neighbourhoods, and the characteristics of these neighbourhoods which together comprise a city can have an influence on them through several pathways, as illustrated in Figs. 17.3, 17.4, and 17.5 above – specifically influencing education E and “permanent income” Y_{base} . Moreover, sims can move from one neighbourhood to another, though only at the moment of a birthday. This capacity to simulate mobility between neighbourhoods in THIM is designed to enable varying degrees of neighbourhood sorting and neighbourhood segregation to be modeled explicitly, along with their effects on income inequality and population health.

In a seminal paper by Schelling (1971), using an extremely simplified ABM, it was shown that even very slight neighbourhood mobility homophily preferences can evolve toward quite dramatic differences in neighbourhood composition. THIM builds on these insights and widely held views on neighbourhood segregation to include an abstract but substantive role for neighbourhoods, and sims’ mobility between neighbourhoods.

To start, THIM has a parameter for the number of neighbourhoods in the city (e.g. 5 or 50). THIM then posits that neighbourhood mobility is based on a simple propensity such that sims (beyond school age, i.e. age $a \geq E$) “want” to live in a neighbourhood with other sims who have similar incomes. Thus, the more discrepant an individual sim’s income is with the average income of the neighbourhood where it is currently residing, the more likely it is to want to move. If the sim is going to move, it seeks the neighbourhood where its own income is least discrepant. If all neighbourhoods with less discrepant incomes are fully occupied (see below), or there are no other neighbourhoods with less discrepant incomes, the attempted move fails. A “year” must pass before another attempt at moving can be made at the sim’s next birthday.

If it is determined that the sim does want to move, THIM then checks whether there are any neighbourhoods that have “available space”. (Note that child sims do not count in determining a neighbourhood’s available space.) The maximum size of neighbourhoods is endogenous; it evolves over simulated time with the size of the population that has completed school.

The underlying intuition here is that neighbourhoods have limited housing stock, so that there are in reality practical constraints on how many people can live in a given neighbourhood. If a sim is already in a neighbourhood, and that

neighbourhood has exceeded its occupancy potential, the sim is able to stay; sims can age into a neighbourhood as it becomes overcrowded without having to move.⁵

Note also that it is possible to have orphan sims if a parent dies before its child sim reaches age $a = E$. In these cases, the young sim stays in the same neighbourhood until it reaches age $a = E$, whereupon it starts receiving income and can then check whether it “wants” to move to a less discrepant neighbourhood in terms of income.

If a sim does want to move, and there is at least one neighbourhood with “available space”, this subset of neighbourhoods is rank ordered by how close each neighbourhood’s average income is to the sim’s own. As long as there is a neighbourhood with space, the sim moves to the least discrepant in terms of income.

A sim’s “desire to move” is parameterized in terms of income thresholds. Setting these thresholds low results in more mobility, hence a higher degree of income sorting/segregation. On the other hand, setting them high will mean that even if incomes diverge substantially within a neighbourhood, there will still not be a lot of mobility. In turn, this means that the within-neighbourhood income inequality will be higher and will approach the overall city-wide level of income inequality, while between-neighbourhood income inequality will be low – and vice versa for a high degree of mobility, hence sorting and segregation.⁶

17.10 Parameters for THIM Simulation Results

In order to illustrate how THIM works, we present in a later section results focusing on one of the key challenges: whether THIM can “account for”, “explain”, or more concretely, generate results that look like those in Fig. 17.2 above. As in other areas of science, we judge THIM as a theory according to whether, given plausible inputs, it can generate realistic outputs.

But before examining any simulation results, it is first necessary to develop a set of input parameter values. This process was in fact quite lengthy. The basic idea was to find a set of plausible input parameters to define a “base case” scenario in line with the stylized facts. THIM – even with its apparently simple structure – is actually quite complex; the behaviour of the agents with the multi-level influences across their hierarchy of aggregated entities, even though described by only a handful of equations, can be highly non-linear. Thus, even though our intuition about the

⁵Indeed, even if a sim has aged into an “overcrowded” neighbourhood, and there are other candidate neighbourhoods, but none of these is less discrepant in terms of income, the sim remains in its current neighbourhood.

⁶Note that this formulation embodies a symmetry assumption – that sims with both high and low income relative to their neighbourhood average income are equally likely to want to move to another neighbourhood. This assumption may not be that realistic, but it was judged adequate for the experiments planned for THIM.

directions of effects of various changes in parameters might be correct, feelings for the magnitudes involved were not obvious – so that many sets of parameters had to be tried out.

One pair of criteria for this process of searching the input parameter space was that, on the one hand, the input parameters should reflect at least stylized versions of observed relationships, and on the other, so should the key outputs. The results of this search of the parameters space is described in Wolfson and Beall (2016). For example, a number of large empirical studies have been referenced to inform the comparative Canada-US parameters on education in relation to parental SES and to average school SES (OECD 2012, 2013) and for differences in parent to child income correlations (Corak 2013).

In this section we simply indicate the main parameters used in that analysis, and then in a later section illustrate their impacts with a few simulation results which seek to replicate the patterns observed in Fig. 17.2 above. These are described in greater detail, along with the algorithms, in Wolfson (2016).

One group of parameters characterizes the neighbourhood structures for the C and U cities, i.e., stylized Canadian and US cities respectively. In line with the greater fragmentation of governance structures in US cities – which often reflects the ability of wealthier neighbourhoods to “opt out” of collective local goods, including education, that would otherwise flow to poorer neighbourhoods, we posit a larger number of neighbourhoods for the U cities than the C cities. We further posit different “mobility” patterns between neighbourhoods.

Recall that THIM embodies a general desire to live in a neighbourhood whose average income is as close as possible to the sim’s own income. However, such mobility only occurs if the proportional difference between the sim’s income and that of the neighbourhood where it resides is above some threshold. If this threshold is high, there is less movement between neighbourhoods. In turn, such a high parameter setting reflects a willingness of higher income sims to live in a neighbourhood with more variation in incomes within the neighbourhood, a greater tolerance for income diversity. For lower income sims, it means they do not have as strong a need to move to a neighbourhood with lower average income when their own income is well below the neighbourhood average. For example, a high mobility threshold embodies the premise that the gentrification of a neighbourhood exerts less pressure on lower income sims to move to another neighbourhood.

Thus, intuitively, lower mobility thresholds plus more neighbourhoods will lead to more income homogeneity within neighbourhoods, and larger income differences between neighbourhoods. In turn, to the extent that neighbourhood average income affects sims’ educational and income prospects, there will be greater inequality in these outcomes between neighbourhoods.

Another set of parameters is for education. They essentially posit that sims living in U cities have parental and neighbourhood income influences on educational attainment that are twice those in the C cities. Further, there is twice the variability in attained years of “education” in U versus C cities. These parameter settings are generally in line with the empirical evidence (OECD 2012).

A third group of parameters characterizes incomes relative to a common age-income profile. These parameters reflect the impacts of own education, parental income and neighbourhood average income respectively on the sim's lifetime potential income, YBase (see below). There is also a parameter for the magnitude of transitory skewed lognormal variability from 1 year to the next in actual income, given age (hence the average level from the common age-income profile), and potential income (YBase). The U values are all posited to be significantly higher than the C values.

There is one more set of parameters differentiating C and U cities – for health transitions and mortality rates. The U cities' parameters embody a larger impact of income and health on mortality than those for the C cities, and income also has a larger impact on health transitions. This latter difference can be seen in part as an implicit reflection of the differences in access to health care between Canada and the US, where in particular there are significant segments of the US lower income population who do not have effective access to care – far more so than in Canada with its universal publicly funded health care.

In addition to these parameters, there is one further major factor – the degree of income inequality. Figure 17.2 (horizontal axis) shows a wide variation in median shares across cities, with US cities spanning a wider range. In THIM, the distribution of income is an output, not an input. A sim's actual income at any given age is the product of a series of factors including those just outlined. But there is an important further influence on this key output – each sim's endowment of a heterogeneous “potential income”, reflecting in part innate ability, personality, and other unobserved characteristics that remain fixed throughout life, YBase. Recall that this variable is defined before applying the influences of education, and parental and neighbourhood relative income.

We have posited eight such distributions of for YBase. In general, these eight distributions have been designed to span a range of income inequality that is very wide – wider than is observed in Canada and the US, both overall and in the cities shown in Fig. 17.2 above.

The most widely used measure of income inequality is the Gini coefficient. However, the construction of comparable income inequality data at the city level shown in Fig. 17.2 above was based on a conceptually simpler measure, the share of income accruing to the bottom half of the population ranked by income, the median share.⁷ Still, income distribution density functions embody an infinity of points, so the “shapes” of these densities can be summarized in an “inequality” index any number of ways. Indeed, there are features of income distributions which are broadly associated in the public's mind with inequality, but are mathematically inconsistent with the usual axiomatization of income inequality measures like the

⁷Strictly speaking, the median share should not be called an inequality measure because it need not be consistent with the partial ordering of income distributions induced by the criterion of Lorenz domination, i.e. that one Lorenz curve is everywhere closer to (or everywhere further from) the 45° line (Atkinson 1970).

Table 17.2 Three inequality measures for eight Ybase input distributions

Input Ybase inequality scenario	1	2	3	4	5	6	7	8
YBase Gini	0.093	0.269	0.358	0.408	0.425	0.471	0.556	0.571
YBase median share	0.450	0.314	0.238	0.223	0.123	0.177	0.127	0.098
YBase polarization	0.015	0.206	0.330	0.291	0.657	0.350	0.380	0.467

Source: Wolfson (2016), with permission

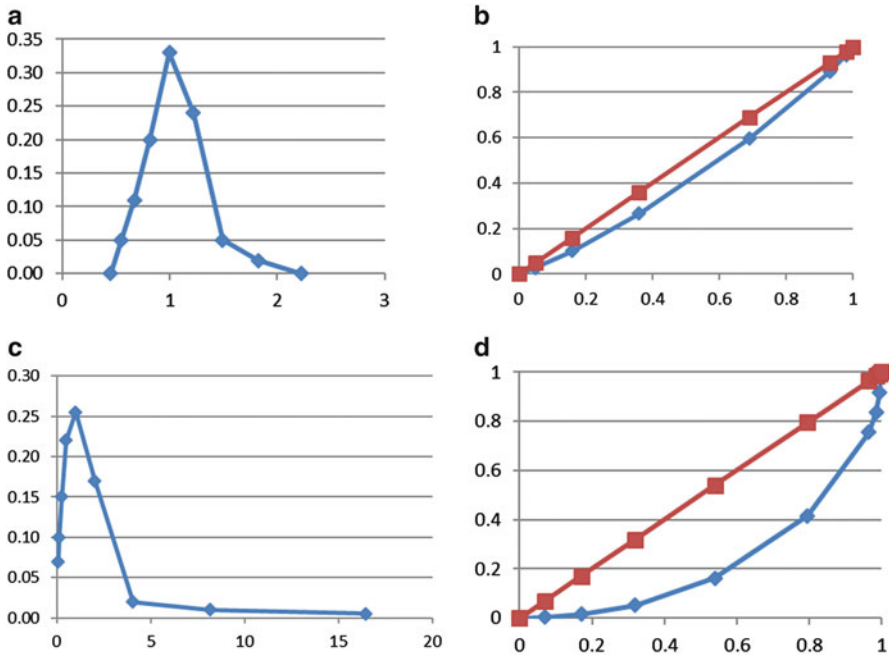


Fig. 17.9 Density functions and Lorenz curves for input Ybase distribution scenarios 1 and 8 (per Table 2) (a) density function for scenario 1 (b) Lorenz curve for scenario 1 (c) density function for scenario 8 (d) Lorenz curve for scenario 8

Gini, specifically polarization (Foster and Wolfson 2010). To assess the robustness of our results, we focus on three summary measures – the Gini, the median share, and the polarization index. Table 17.2 shows these three measures for each of the eight YBase distributions used as inputs.⁸ Figure 17.9 then shows the density functions and Lorenz curves for the YBase distributions in scenarios 1 and 8. (Note that the horizontal axes in Figs. 17.9a and c are different.)

⁸Note that these measures need not be rank order correlated, and in fact are not for the YBase distributions used as inputs. Such situations are typically associated with crossing Lorenz curves (Atkinson 1970).

17.11 Software Environment

THIM was originally developed using NetLogo. However, it quickly became apparent that NetLogo could not handle a model of this complexity. As a result, it was re-implemented and completed using Statistics Canada's ModGen software (Statistics Canada, ModGen). ModGen is a dialect of C++ (actually a C++ pre-compiler) designed originally for large scale policy oriented microsimulation modeling. Its main use is still for that purpose. But ModGen is also sufficiently flexible and convenient that it provides an excellent development environment for agent-based models like THIM.

17.12 Simulation Outputs

A critical part of any ABM like THIM is the kinds of outputs it can produce. As noted by Hegselmann (2012) in his discussion of the Schelling model, being able to visualize the dynamics of the simulation as it unfolds is often very informative. However, such real time data visualization exacts a heavy price in execution time – one of the reasons NetLogo had to be abandoned. ModGen, the software environment now used for implementing THIM, has been designed (among other things) to optimise execution time. This is accomplished by having everything execute in RAM, with essentially no disk reads or writes. As a result, it is critical to plan the outputs carefully.

Based on many years of experience, we have found it best to design outputs to cover a wide variety of attributes and processes being simulated. In the first instance, these outputs support debugging and exploration of the model. Then, depending on the questions of interest, the focus can be on a subset of these outputs.

ModGen natively produces two main kinds of outputs – multi-dimensional tables, and samples of trajectories of the various actors. In the case of THIM, these actors are sims, neighbourhoods, and the city. For the tables, the most flexible way to view them is via export to Excel in pivot table format (using a specially designed Excel macro). For agent trajectories, ModGen has a companion piece of software called Biobrowser (Statistics Canada, ModGen). It is most useful in debugging, and for explaining how the simulation model works.

17.13 THIM Simulation Results

By construction, THIM embodies among other things a theory of the potential determinants of the patterns in Fig. 17.2.⁹ We now describe a few key results of a series of simulations to determine whether the factors included in THIM are sufficient to generate both the US and the Canadian patterns.

⁹Further details on the model, including parameters, source code, and an executable version are available from the authors (mwolfson@uottawa.ca).

Of course, it may well be possible to devise other theories that can account for the patterns observed in Fig. 17.2. All we are exploring with THIM is whether there is at least one satisfactory theory.

It is possible that THIM as constructed is unable to do reproduce the patterns in Fig. 17.2. If so, it is lacking key factors and/or causal pathways, and is therefore an inadequate theory.

To start, though, we examine several simulation results from the perspective of face validity. While THIM is theoretical, it is based to a substantial degree on empirical evidence, albeit summarized in the form of stylized facts. As a result, THIM should be able to generate plausible outputs – results that appear reasonable qualitatively when compared to actual observations. Further, these outputs should be “emergent”; they should not be simple transforms of inputs, but rather the results of the full interacting richness of the sims’ posited and then simulated behaviours. This is the case.

But first some background. THIM simulations start with a population with some very strong simplifying assumptions, including that everyone is perfectly healthy ($H = 1$) at birth, and they are randomly scattered across neighbourhoods. As a result, THIM simulations need to run at least 100 “years” until the population “settles down” to something like an asymptote – though there is no a priori assumption of equilibrium in THIM. Further, since THIM includes a number of stochastic processes (e.g. individual sim’s health as a form of random walk), even asymptotes may only be approached approximately and “fuzzily”; population aggregates can be expected to continue at least to “wobble” around asymptotes indefinitely.¹⁰ In order to assess whether any given set of results is reasonably stable, the THIM simulations have all been run for 500 “years” and the results for various decades over this time span checked for stability.

One example for the purpose of establishing face validity of the THIM simulations is the univariate distributions of health status for each of several age groups. In THIM, this is not an input, but rather the outcome of many interacting factors in a simulation. Figure 17.10 shows count frequencies for levels of H (the summary index of health in the $[0,1]$ interval used in THIM) by age group as generated by a typical THIM simulation. The horizontal axis is first broken down by selected 5-year age groups, and then within age groups by levels of H . Curves for each of five different decades are shown. Since these curves lie very close to one another, in this simulation the results are quite stable over time.¹¹ Since these curves lie very close to one another, in this simulation the results are quite stable over time.

Qualitatively, the distributions become less tall and more negatively skewed (i.e. less concentrated at high levels of H) as we move from left to right up the age spectrum, indicating both a general decline in health and a decline in the total

¹⁰These stochastic elements also give rise to Monte Carlo error. The THIM results presented have been assessed to ensure that this source of error is not material.

¹¹Not all age groups are shown; of those shown, for example, “40–45” indicates all sims at least exact age 40 and less than exact age 45.




Fig. 17.10 Cumulative H distributions for selected age groups, THIM simulation

population as a result of higher (cumulative) mortality with higher age. These distributions are very similar to those observed in Statistics Canada’s National Population Health Survey (Statistics Canada) for the McMaster Health Utility Index (HUI; Feeney et al. 2001), the real world counterpart of the H variable included in THIM.

THIM also produces “reasonable” looking individual-level health-income gradients. This is evident by comparing Fig. 17.11 to Fig. 17.1 above (though Fig. 17.1 shows mortality rather than H, so slopes downward rather than upward, and all age groups are combined) and to McIntosh et al. (2009). Again, the horizontal axis is first broken down by selected 5-year age groups, but this time within age groups by intervals of income of increasing width. The vertical axis is the average level of H within each age/income interval. Curves for each of five different decades are again shown, for years 90–100, 190–200, . . . 490–500.¹²

In this case, moving from left to right to higher age groups, the gradients become steeper and more variable/noisy. The noisiness is due to the decreasing numbers of sims at higher ages. (The absence of sims with very low health status H in the 90–95 age interval is due to mortality selection, as posited, with mortality rates depending on H.)

Table 17.3 presents our first set of main results. The top three rows are identical to Table 17.2, showing measures of inequality for the input Ybase distributions. However, the resultant income (Y) inequality values in the following six rows are considerably different – since actual incomes depend not only on potential income YBase, but also on the influences of parental and neighbourhood factors working

¹²As in Fig. 17.10, the simulation spans 500 years, and only selected decades over this 500-year span that have been graphed.




Fig. 17.11 Health status (H) by income for selected age groups – THIM simulation

Table 17.3 Income inequality’s path from Ybase to final population income inequality for three inequality measures, eight Ybase scenarios, and U versus C parameters

Input Ybase inequality scenarios	1	2	3	4	5	6	7	8
YBase Gini	0.093	0.269	0.358	0.408	0.425	0.471	0.556	0.571
YBase median share	0.450	0.314	0.238	0.223	0.123	0.177	0.127	0.098
YBase polarization	0.015	0.206	0.330	0.291	0.657	0.350	0.380	0.467
Output income inequalities								
C Gini	0.281	0.383	0.441	0.477	0.403	0.496	0.511	0.530
C median share	0.298	0.233	0.187	0.174	0.202	0.165	0.160	0.152
C polarization	0.140	0.190	0.266	0.270	0.247	0.280	0.276	0.287
U Gini	0.393	0.409	0.416	0.420	0.420	0.467	0.514	0.474
U median share	0.225	0.213	0.208	0.206	0.206	0.178	0.155	0.179
U polarization	0.192	0.213	0.219	0.221	0.221	0.259	0.291	0.249

Source: Wolfson (2016), with permission

through education, and directly on income, plus variation over the life cycle based on a typical average age-income profile, plus an annual stochastic disturbance reflecting short run income volatility, plus (especially at higher ages) the possibility of mortality selection by income.

For example, the most equal potential income YBase distribution (scenario 1) has a Gini of 0.093. But when this potential income scenario is played out in a THIM simulation, the income distribution Gini’s for the C and U cities end up much more unequal, at 0.281 and 0.393 respectively. At the other end of the YBase inequality scenarios, the highly unequal input Gini of 0.571 (scenario 8) for the YBase distribution ends up generating somewhat less unequal Gini’s of 0.530 and 0.474 for C and U cities respectively.

Our key health outcome measures in THIM are life expectancy (LE) and health-adjusted life expectancy (HALE). These are both broader population health

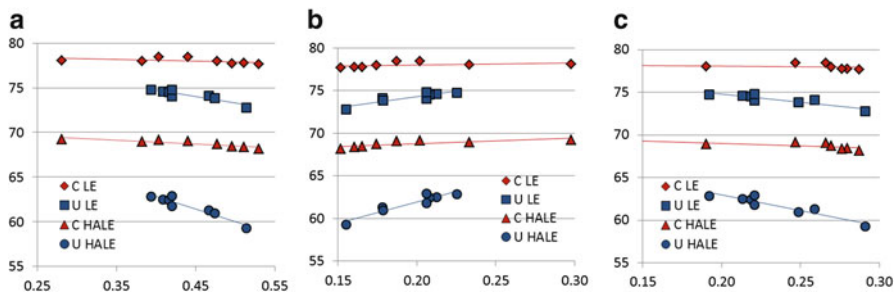


Fig. 17.12 LE and HALE for U and C Cities, by (a) Gini, (b) median share, and (c) polarization index

measures than the working age mortality used in Fig. 17.2 above, especially HALE (Sullivan 1971; McIntosh et al. 2009; Wolfson 2014).

These outcomes are plotted in Fig. 17.12a–c with the (output) income inequality measures in Table 17.3 plotted along the horizontal axis.¹³

Recall that the input baseline age-specific mortality schedule is associated with a conventional period life expectancy (LE) of 77 years. The graphs, however, show that there is a range of LEs resulting from the THIM simulations – some a bit higher and most somewhat lower.¹⁴

The first observation is that in all cases, the LEs and HALEs (years, on the vertical axes) of the C cities are everywhere higher than those for the U cities. This result corresponds to the observation in Fig. 17.2 that Canadian cities generally have lower working age mortality rates at similar levels of the median share.

Most importantly from the viewpoint of our main hypothesis, the slopes of the relationships are systematically different. In all cases, there is an almost flat relationship for both LE and HALE with all three of the inequality measures for the C cities. But for the U cities, there is in all cases an evident slope.

We have therefore established that the theory and stylized factual inputs embodied in THIM can indeed generally reproduce the observations shown in Fig. 17.2 above.

17.14 Discussion and Conclusions

We have constructed a theory to account for the observation that correlations between income inequality and health appear contingent on country-specific factors.

¹³While the vertical axes in the three graphs are identical, the horizontal axis scales are specific to each measure.

¹⁴LE and HALE are computed in THIM as in usual real world practice using cross-sectional data from overlapping birth cohorts and the Sullivan (1971) method.

This theory is embodied in an agent-based Theoretical Health Inequality Model (THIM). As with virtually all theories, major simplifications have been made. At the same time, we have appealed to real world observations to infer a set of “stylized facts” which have been incorporated into THIM – both as algorithms reflecting causal pathways, and as quantified parameters reflecting specific strengths of relationships.

As a first test, THIM has demonstrated face validity in Figs. 17.10 and 17.11, where it has generated realistic patterns of outputs.

And most importantly, THIM can account for the patterns observed in Fig. 17.2. THIM, given a plausible set of input parameters, is able to generate contingent patterns of correlation between income inequality and health similar to those observed, specifically between Canada and the US.

Further experiments with THIM (Wolfson and Beall 2016) explore these results in more detail, by assessing the relative contributions of various sub-groups of parameters to the overall result. The surprising result is that neighbourhood income segregation does not look to be the main factor accounting for the different patterns for US and Canadian cities in Fig. 17.2. Rather, the factors more likely accounting for the Canada-US differences include stronger parent to child transmission of advantage and disadvantage, less fragmentation of municipal governance structures resulting in more even distributions of local public goods, and stronger effects of individual income on health in the US compared to Canada.

References

- Ashby, W. R. (1958), Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2), 83–99. Available at <http://pcp.vub.ac.be/Books/AshbyReqVar.pdf>. Republished on the web by F. Heylighen—Principia Cybernetica Project.
- Atkinson, A. B. (1970). On the measurement of inequality. *Journal of Economic Theory*, 2(3), 244–263.
- Backlund, E., Rowe, G., Lynch, J., Wolfson, M., et al. (2007), Income inequality and mortality: a multilevel prospective study of 521 248 individuals in 50 US states. *International Journal of Epidemiology*, 36(3), 590–596. <http://www.ncbi.nlm.nih.gov/pubmed/10514157>
- Braveman, P. A., Cubbin, C., Egerter, S., Williams, D. R., & Pamuk, E. (2010). Socioeconomic disparities in health in the United States: What the patterns tell us. *American Journal of Public Health*, 100(Suppl 1), S186–S196. <http://doi.org/10.2105/AJPH.2009.166082>.
- Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, 27(3), 79–102. <http://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.27.3.79>.
- Feeny, D., Furlong, W., Torrance, G., Goldsmith, C., Zhu, Z., DePauw, S., Denton, M., & Boyle, M. (2001). Multi-attribute and single attribute utility functions for the health utility index mark 3 system. *Medical Care*, 40(2), 113–128.
- Foster, J., & Wolfson, M. (2010). Polarization and the decline of the middle class: Canada and the U.S., *Journal of Economic Inequality*, 8(2), 247–273.
- Gravelle, H. (1998). How much of the relation between population mortality and unequal distribution of income is a statistical artefact? *BMJ: British Medical Journal*, 316(7128), 382–385.

- Hegselmann, R. (2012). Thomas C. Schelling and the computer: Some notes on Schelling's Essay "On Letting a Computer Help with the Work", *Journal of Artificial Societies and Social Simulation*, 15(4), 9.
- Johnson, R. W., & Neumark, D. (1996). Wage declines among older men. *The Review of Economics and Statistics*, 78(4), 740–748.
- Kaplan, M. S., et al. (2007). The predictive validity of health-related quality of life measures: Mortality in a longitudinal population-based study. *Quality of Life Research*, 16(9), 1539–1546.
- Kaplan, G. A., Diez Roux, A., Galea, S., & Simon, C. P. (eds) (2016). *Growing Inequality: Bridging Complex Systems, Population Health, and Health Disparities*. Washington, D.C.: Westphalia Press.
- Luong, M., & Hébert, B. P. (2013). Age and earnings. *Statistics Canada*. Available: <http://www.statcan.gc.ca/pub/75-001-x/2009101/article/10779-eng.htm#Robinson>.
- Lynch, J. W., Davey Smith, G., Kaplan, G. A., & House, J. S. (2000). Income inequality and mortality: Importance to health of individual income, psychosocial environment, or material conditions. *BMJ: British Medical Journal*, 320(7243), 1200–1204.
- McIntosh, C. M., Fines, P., Wilkins, R., & Wolfson, M. (2009). *Income disparities in health-adjusted life expectancy for Canadian adults, 1991–2001*. Ottawa: Health Reports, Statistics Canada.
- Murphy, K. M., & Welch, F. (1990). Empirical age-earnings profiles. *Journal of Labor Economics*, 8(2), 202–229.
- NIH (2010) <http://nih.gov/news/health/aug2010/od-05.htm> and http://obssr.od.nih.gov/scientific_areas/social_culture_factors_in_health/health_disparities/nich.aspx.
- OECD. (2012). *PISA database* (Table II.2.10 Figure 11 and Table II.2.6 Figure 9).
- OECD. (2013). *Skills outlook* (Figure 3.8c Fig 10).
- Ross, N. A., Wolfson, M. C., Dunn, J. R., Berthelot, J.-M., Kaplan, G. A., & Lynch, J. W. (2000). Relation between income inequality and mortality in Canada and in the United States: Cross sectional assessment using census data and vital statistics. *BMJ: British Medical Journal*, 320(7239), 898–902.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2), 143–186.
- Statistics Canada (NPHS). <http://www.statcan.gc.ca/eng/survey/household/3225>.
- Statistics Canada (ModGen). <http://www.statcan.gc.ca/microsimulation/modgen/modgen-eng.htm>.
- Sullivan, D. F. (1971). A single index of mortality and morbidity. *HSMHA Health Reports*, 86(4), 347–354.
- Wolfson, M. C. (2014). Measuring the health component of quality of life. *Statistical Journal of the IAOS: Journal of the International Association for Official Statistics*, 30(3), 193–207.
- Wolfson, M. C. (2016). *Exploring contingent inequalities – Building the theoretical health inequality model, model algorithm description*. http://micromodel.beyond2020.com/demodoc/thimn_1_2014_5_6_/Workbooks/Exploring%20Contingent%20Inequalities%20-%20THIM%20Algorithm%207-16.pdf.
- Wolfson, M. C., & Beall, R. (2016). Contingent inequalities – A quasi-theoretical exploration of health inequalities in the U.S. and Canada. In G. Kaplan (Ed.), *Growing inequality: Bridging complex systems, population health, and health disparities*. Washington, DC: Westphalia Press.
- Wolfson, M. C., Rowe, G., Gentleman, J. F., et al. (1993). Career earnings and death: A longitudinal analysis of older Canadian men. *Journal of Gerontology: Social Sciences*, 48, S167–S179.
- Wolfson, M. C., Kaplan, G., Lynch, J., Ross, N., & Backlund, E. (1999). Relation between income inequality and mortality: Empirical demonstration. *BMJ: British Medical Journal*, 319(7215), 953–955.