

Computational Risk Management

Peter Sarlin

Mapping Financial Stability

 Springer

Computational Risk Management

Series editors

Desheng Dash Wu

David L. Olson

John R. Birge

For further volumes:

<http://www.springer.com/series/8827>

Peter Sarlin

Mapping Financial Stability

 Springer

Peter Sarlin
Centre of Excellence SAFE
Goethe University Frankfurt
Frankfurt
Germany

RiskLab Finland
IAMSR Åbo Akademi University
Turku
Finland

Arcada University of Applied Sciences
Helsinki
Finland

ISSN 2191-1436 ISSN 2191-1444 (electronic)
ISBN 978-3-642-54955-7 ISBN 978-3-642-54956-4 (eBook)
DOI 10.1007/978-3-642-54956-4
Springer Heidelberg New York Dordrecht London

Library of Congress Control Number: 2014936096

© Springer-Verlag Berlin Heidelberg 2014

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. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law. 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.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Work is more fun than fun

—Noel Coward

Acknowledgments

Most parts of this book are taken from my Ph.D. thesis (TUCS Dissertation Series, No 159), which I defended on 6 June 2013. Hence, all acknowledgments in the Preface of the thesis obviously still apply. Beyond those people and institutions, I would like to thank the editors of Springer's Computational Risk Management Series, Desheng Wu, David L. Olson, and John R. Birge, for their kind help and support in turning a thesis into this book. Moreover, I once again extend my gratitude to my family, and its various extensions, for the love and patience shown with yet another project.

Frankfurt, January 31, 2014

Peter Sarlin

Contents

1	Introduction	1
1.1	Background	3
1.2	Research Objectives and Questions	7
1.3	Overview of the Book	9
1.4	Original Publications and Their Contributions	12
	References	13
2	Macroprudential Oversight	15
2.1	Financial Systems, Fragilities and Instabilities	16
2.1.1	Key Components of Financial Systems	16
2.1.2	Financial (in)stability	18
2.1.3	Fragility of Financial Systems	19
2.1.4	A Systemic Risk Cube	21
2.2	Theoretical and Empirical Underpinnings	23
2.2.1	Theoretical Models	23
2.2.2	Empirical Findings	27
2.3	Tools for Safeguarding Financial Stability	31
2.3.1	Early-Warning Indicators and Models	31
2.3.2	Macro Stress-Testing Models	33
2.3.3	Contagion and Spillover Models	35
2.3.4	Tools With Visual Capabilities	36
2.4	A Framework for Macroprudential Oversight	39
2.5	Concluding Discussion	42
	References	43
3	Macroprudential Data	51
3.1	Data: What are They?	52
3.2	Data for Macroprudential Oversight	53
3.2.1	Macroeconomic Data	54
3.2.2	Banking System Data	56
3.2.3	Market-Based Data	57
3.3	A Four-Dimensional Data Cube	58

3.4	Stylized Challenges in Macroprudential Data	60
3.4.1	Macroeconomic Data	61
3.4.2	Banking System Data	62
3.4.3	Market-Based Data	64
3.5	Concluding Discussion	64
	References	66
4	Data and Dimension Reduction	69
4.1	Some Key Concepts and Definitions	70
4.1.1	Knowledge Discovery and Data Mining	70
4.1.2	Information Visualization	75
4.1.3	Visual Analytics	82
4.2	Dimension Reduction	83
4.2.1	Aims of Dimension Reduction	84
4.2.2	An Overview of Methods	86
4.3	Data Reduction	88
4.3.1	Aims of Data Reduction	88
4.3.2	An Overview of Methods	89
4.4	The Self-Organizing Map (SOM)	92
4.4.1	Parametrizing the SOM	93
4.4.2	Supervision of the SOM	95
4.4.3	Qualities of the SOM	96
4.5	Concluding Summary	97
	References	97
5	Data-Dimension Reductions: A Comparison	101
5.1	The Optimal Method: A Literature Review	102
5.1.1	A Comparison of Dimension Reductions	102
5.1.2	A Comparison of Data Reductions	103
5.1.3	Why is the Literature so Divided?	103
5.2	DDR Combinations for the Task at Hand	104
5.2.1	Aims and Needs for the Task	104
5.2.2	Aims and Needs of DDR Combinations	105
5.3	A Qualitative Comparison	107
5.4	Illustrative Experiments	109
5.5	The SOM and Its Visualization Aids	112
5.6	Discussion	116
5.7	Concluding Summary	118
	References	119
6	Extending the SOM	123
6.1	Time in SOMs: A Brief Review	124

6.2	Extensions for Exploiting the SOM	126
6.2.1	Fuzzification of the SOM	126
6.2.2	Transition Probabilities on the SOM.	130
6.2.3	Shock Propagation on the SOM	135
6.3	The Self-Organizing Time Map	137
6.3.1	SOTM Properties	140
6.3.2	Qualities and Properties of the SOTM.	141
6.3.3	Visualizations of the SOTM.	143
6.3.4	Some Illustrative Examples	144
6.4	Concluding Summary	154
	References	155
7	Self-Organizing Financial Stability Map	159
7.1	Data	160
7.1.1	Identifying Systemic Financial Crises	161
7.1.2	Macro-financial Indicators of Vulnerabilities and Risks.	163
7.2	Model Evaluation Framework	164
7.3	Model Training Framework	170
7.4	Training and Evaluation of the SOFSM	172
7.5	Performance and Robustness of the SOFSM	174
7.6	Concluding Summary	180
	References	181
8	Exploiting the SOFSM	183
8.1	The SOFSM: Its Output and Interpretation	185
8.2	Visualizing the State of Financial Stability on the SOFSM	188
8.3	Fuzzification of the SOFSM	193
8.4	Transitions on the SOFSM.	194
8.5	Scenario Analysis on the SOFSM.	196
8.6	Shock propagation on the SOFSM	199
8.7	Outlier Analysis with the SOFSM	201
8.8	Combining the SOFSM with Predictive Methods	203
8.9	Concluding summary	208
	References	209
9	Decomposing Financial Crises with SOTMs.	211
9.1	A Decomposition of the Global Financial Crisis.	212
9.1.1	Parametrizing the SOTM.	212
9.1.2	A Univariate View of the Crisis.	212
9.1.3	A Multivariate View of the Crisis	214

- 9.2 A Decomposition of Modern Financial Crises 217
 - 9.2.1 Parametrizing the Time-to-Event SOTM 218
 - 9.2.2 A Univariate View of Crises 219
 - 9.2.3 A Multivariate View of Crises 220
- 9.3 Concluding Summary 221
- References 221

- 10 Conclusions, Limitations and the Future 223**
 - 10.1 Conclusions, Findings and Implications. 223
 - 10.1.1 Implications for Dimension Reduction 227
 - 10.1.2 Implications for Policy Use 228
 - 10.2 Limitations 229
 - 10.2.1 Limitations of Methods 229
 - 10.2.2 Limitations of Applications 230
 - 10.3 Future Research 231
 - 10.3.1 Future Methods 231
 - 10.3.2 Future Applications. 232
 - References 233

Acronyms

AANN	Auto-Associative Neural Network
AE	Advanced Economy
AI	Artificial Intelligence
ANN	Artificial Neural Network
AUC	Area Under the Curve
BIS	Bank for International Settlements
BMU	Best-Matching Unit
BSC	Banking Supervision Committee
CCA	Curvilinear Component Analysis
CCE	Coordinated Compilation Exercise
CCF	Correct Classification Frontier
CDA	Curvilinear Distance Analysis
CDS	Credit Default Swap
CHAPS	Clearing House Automated Payment System
CISS	Composite Indicator of Systemic Stress
CRISP-DM	Cross Industry Standard Process for Data Mining
DA	Discriminant Analysis
DDR	Data-Dimension Reduction
DQRS	Data Quality Reference Site
ECB	European Central Bank
EDA	Exploratory Data Analysis
EME	Emerging Market Economy
EU	European Union
FCM	Fuzzy C-Means
FDI	Financial Distress Index
FSI	Financial Soundness Indicator
FSOM	Feedback SOM
GA	Genetic Algorithm
GDDS	General Data Dissemination System
GDP	Gross Domestic Product
GFSM	Global Financial Stability Map
GTM	Generative Topographic Mapping
IA	Intelligence Amplification
IMF	International Monetary Fund

IS	Information Systems
IT	Information Technology
KD	Knowledge Discovery
KDD	Knowledge Discovery in Databases
LE	Laplacian Eigenmaps
LLE	Locally Linear Embedding
LMDS	Local MDS
MDS	Multidimensional Scaling
MPI	Macro-Prudential Indicator
MSOM	Merge SOM
MVU	Maximum Variance Unfolding
NG	Neuro-Genetic
OECD	Organisation for Economic Co-operation and Development
PCA	Principal Component Analysis
QE	Quantization Error
RecSOM	Recursive SOM
RO	Research Objective
ROC	Receiver Operating Characteristics
RQ	Research Question
RSOM	Recurrent SOM
RT	Research Theme
SDDS	Special Data Dissemination Standard
SIFI	Systemically Important Financial Institutions
SNA	System of National Accounts
SOFSM	Self-Organizing Financial Stability Map
SOM	Self-Organizing Map
SOMSD	SOM for Structured Data
SOTM	Self-Organizing Time Map
TARGET	Trans-European Automated Real-time Gross Settlement Express Transfer System
TPM	Transition Probability Matrix
t-SNE	t-Distributed Stochastic Neighbor Embedding
TSOM	Temporal SOM
UK	United Kingdom
US	United States
VQ	Vector Quantization
XOM	Exploration Observation Machine

Summary

The ongoing global financial crisis has demonstrated the importance of a systemwide, or macroprudential, approach to safeguarding financial stability. An essential part of macroprudential oversight concerns the tasks of early identification and assessment of risks and vulnerabilities that eventually may lead to a systemic financial crisis. Thriving tools are crucial as they allow early policy actions to decrease or prevent further build-up of risks or to otherwise enhance the shock absorption capacity of the financial system. In the literature, three types of systemic risk can be identified: (i) build-up of widespread imbalances, (ii) exogenous aggregate shocks, and (iii) contagion. Accordingly, the systemic risks are matched by three categories of analytical methods for decision support: (i) early warning, (ii) macro stress-testing, and (iii) contagion models. Stimulated by the prolonged global financial crisis, today's toolbox of analytical methods includes a wide range of innovative solutions to the two tasks of risk identification and risk assessment. Yet, the literature lacks focus on the task of risk communication.

This book concerns macroprudential oversight from the viewpoint of all three tasks: Within analytical tools for *risk identification* and *risk assessment*, the focus concerns a tight integration of means for *risk communication*. Data and dimension reduction methods, and their combinations, hold promise for representing multivariate data structures in easily understandable formats. The overall task of the work in this book is to represent high-dimensional data concerning financial entities on low-dimensional displays. The low-dimensional representations have two subtasks: (i) to function as a display for individual data concerning entities and their time series, and (ii) to use the display as a basis to which additional information can be linked. The final nuance of the task is, however, set by the needs of the domain, data, and methods. The following five questions comprise subsequent steps addressed in this book:

1. What are the needs for macroprudential oversight?
2. What form do macroprudential data take?
3. Which data and dimension reduction methods hold most promise for the task?
4. How should the methods be extended and enhanced for the task?
5. How should the methods and their extensions be applied to the task?

Based upon the Self-Organizing Map (SOM), the work in this book not only creates the Self-Organizing Financial Stability Map (SOFSM), but also lays out a general framework for mapping the state of financial stability. The work in this book also introduces three extensions to the standard SOM for enhancing the visualization and extraction of information: (i) fuzzification, (ii) transition probabilities, and (iii) network analysis. Thus, the SOFSM functions as a display for risk identification, on top of which risk assessments can be illustrated. In addition, this book puts forward the Self-Organizing Time Map (SOTM) to provide means for visual dynamic clustering, which in the context of macroprudential oversight concerns the identification of cross-sectional changes in risks and vulnerabilities over time. Rather than automated analysis, the aim of visual means for identifying and assessing risks is to support disciplined and structured judgmental analysis based upon policymakers' experience and domain intelligence, as well as external risk communication.

Chapter 1

Introduction

When the crisis came, the serious limitations of existing economic and financial models immediately became apparent. [...] As a policy-maker during the crisis, I found the available models of limited help. In fact, I would go further: in the face of the crisis, we felt abandoned by conventional tools.

– Jean-Claude Trichet, President of the ECB,
Frankfurt am Main, 18 November 2010

The narrative of the still ongoing global financial crisis—which undeniably has become an economic crisis, not to say a crisis of economics—has no unambiguous description. While the many factors directly and indirectly linked to the causes of the crisis are divisive, the effects of the crisis are less so. The period since the outbreak of the financial crisis in mid-2007 has been characterized by a number of multifaceted problems in financial systems and society in general: liquidity issues in large financial institutions, sovereign debt problems, government interventions in banks, the collapse of housing and stock markets, overall losses in welfare and growth, etc. While being divisive, today’s hindsight discussions illustrate a wide range of so-called systemic risks, vulnerabilities and imbalances that depicted financial systems prior to the collapse of Lehman Brothers, and the subsequent worldwide financial meltdown. Yet, as wisely put by Bezemer (2011), “*no one saw this coming*”. The aim herein is to provide tools to better *see* it coming.

So, what is financial instability and systemic risk? Paraphrasing Justice Potter Stewart’s definition of explicit content, as noted by Bisias et al. (2012), describes how vaguely financial instability and systemic risk is commonly viewed: we struggle in defining it, but we think we know it when we see it. Yet, this is no basis for measurement and analysis of threats to financial stability, however that is defined. There is obviously no undisputed definition. Herein, from the viewpoint of the antonym, financial instability is defined as an event that has adverse effects on a number of important financial institutions or markets (ECB 2009). Systemic risk, as defined by the same source, is a risk of widespread financial instability that impairs the

functioning of the financial system, with severe implications on economic growth and welfare.

Then, how costly is financial instability? Even though few foresaw financial instabilities, the above definition indicates them being of high impact. Patterns of the past and today are alike, in that financial crises have recurred throughout monetary history. Research, not to only rely on perceived occurrences, has revealed a doubling of the frequency of financial stress episodes since the end of the Bretton Woods system in 1973 (Bordo et al. 2001). Numerous sources of evidence suggest that historical financial costs of crises have been enormous. Cardarelli et al. (2011) show that out of 113 episodes of financial crisis for key advanced economies, 29 were followed by an economic slowdown and an equal number by recessions. Eichengreen (2004) reveals that the average output loss from a financial crisis is around 9 % of gross domestic product (GDP), whereas the most severe crises caused a GDP loss of over 20 %. Likewise, Hoggarth et al. (2002) find cumulative output losses from a crisis to be up to 30 % of GDP. Dell'ariccia et al. (2008) and Laeven and Valencia (2008) highlight the importance of banking crises by revealing that their median losses have been at around 20–25 % of GDP. Hence, it is a trivial fact that early identification of financial instability would be useful, in particular as it would enable policymakers to make corrective actions prior to the event.

A key concept is, however, an early enough identification of financial instabilities. The events of last years have illustrated that policy actions introduced at a late stage may be highly costly for tax payers. The global financial crisis has brought a large number of European banks to the brink of collapse, leading to bailout costs beyond anything previously experienced. Data from the European Commission shows that government assistance to stabilize the European Union (EU) banking sector exceeded €1.6 trillion at the end of 2010. Though accounting only for a moderate share of the total cost of a systemic banking crisis, this amounts to more than 13 % of EU-level GDP, not to mention the fact that the sovereign-bank nexus in Europe still remains to be resolved. Yet, defining financial instability, and its costs, provide no good means for measuring it. This accentuates a need for tools for early identification and assessment of systemic risks that might possibly lead to financial instability. These tools would allow policymakers to introduce policy actions to decrease or prevent further build-up of risks and vulnerabilities and otherwise enhance the shock absorption capacity of the financial system.

How should threats to financial stability be measured and analyzed? The current financial crisis has highlighted the importance of a system-wide, or macroprudential, approach to safeguarding financial stability, rather than one being only concerned with the stability of individual financial institutions (i.e., microprudential). This accentuates the need for a thorough understanding of not only financial entities, be they economies, markets or institutions, but also their interconnections, interlinkages and system-wide importance. Analytical tools and models provide means for two types of tasks: (i) early identification of vulnerabilities and risks, as well as their triggers, across financial instruments, markets and institutions, and (ii) early assessment of transmission channels of and a system's resilience to shocks, and potential severity of the risk materialization. As above noted by Mr. Trichet, the toolbox of

models for macroprudential oversight is, however, still in its infancy [see also, e.g., Hartmann (2009) and Schou-Zibell et al. (2010)]. In addition, Schou-Zibell et al. (2010) note that the global financial crisis hit advanced economies that lie in the very forefront of financial stability reporting. Yet, such reporting to a large extent takes the form of overall qualitative assessments and policy discussions. To this end, a systematic and data-driven approach to monitoring financial stability is as likely to support the use of the rich information provided by policymakers' judgment and experience as the latter is to support the former.

How does this book serve the task? While quantitative methods have been applied for these purposes, they seldom focus on providing policymakers with representations of data in easily understandable formats. This points to the two tasks of risk identification and risk assessment lacking the component of risk communication. A visualization or abstraction of high-dimensional data can be seen as an artifact supporting the knowledge crystallization process. The work in this book puts forward a set of tools for visual identification and assessment of systemic risks, in order to support the task of risk communication. The tools, while providing means for a wide range of analytics, should rather be treated as a starting point than an ending point for the overall aims of macroprudential oversight, to which a central supporting ingredient is policymakers' judgment and experience.

The sequel of the introduction is structured as follows. First, Sect. 1.1 presents the background of macroprudential oversight and briefly positions the topic of this book. Then, Sect. 1.2 discusses the two key objectives of this book and untangles them into five research questions, whereas Sect. 1.3 provides a chapter-specific overview of the book.

1.1 Background

A comprehensive macroprudential approach to safeguarding financial stability obviously starts from a thorough understanding of the inner (dys)functioning of the financial system. In addition to the literature on financial systems, fragilities, risks and instabilities being broad, the tidal wave of research that the global financial crisis stimulated is also transforming it at a fast pace. Thus, a wide range of topics in the literature remain to be disputed. Yet, one notion that few oppose is that a key aim is to have a resilient and well-functioning financial system. One characterization of such a financial system is through the following three pillars (Fell and Schinasi 2005): well-managed *financial institutions*, efficiently functioning *financial markets* and a strong and robust *financial infrastructure*. That said, the frequent incidences of costly financial crises do, however, indicate that the three pillars of well-functioning financial systems have defects. While each recurrence of financial instability may have sources of its own kind, market imperfections like asymmetric and incomplete information, externalities and public-good characteristics and incomplete markets are a central group of defects. These imperfections, when being related to a financial sector, may lead to significant fragility of not only individual entities or firms, but

also the entire system [see, e.g., Carletti (2008)]. de Bandt and Hartmann (2002) relate fragilities in financial systems to three causes: (i) the structure of banks, (ii) the interconnection of financial intermediaries, and (iii) the information intensity of financial contracts. The material risks of these fragilities support the role of governments and other supervisory authorities in addressing and monitoring financial instability.

To concretize the notion of systemic risk, we herein follow the definition of three forms of systemic risk by de Bandt et al. (2009): (i) endogenous build-up and unraveling of widespread imbalances; (ii) exogenous aggregate shocks; and (iii) contagion and spillover. The first form of systemic risk focuses on the *unraveling of widespread imbalances* and is illustrated by a thorough literature on the presence of risks, vulnerabilities and imbalances in banking systems and the overall macro-financial environment prior to historical financial crises. Early and later literature alike have identified common patterns in underlying vulnerabilities preceding financial crises [see, e.g., Minsky (1982) and Reinhart and Rogoff (2008)]. The second type of systemic risk, *exogenous aggregate shocks*, have been shown to co-occur with financial instabilities [see, e.g., Gorton (1988) and Demirgüç-Kunt and Detragiache (1998)]. Here, an example is the collapse of banks during recessions due to the vulnerability to economic downturns. The *contagion* literature provides evidence on the final, third form of systemic risk, that is, the cross-sectional transmission of financial instability [see, e.g., Upper and Worms (2004) and van Lelyveld and Liedorp (2006)]. Here, episodes of financial instabilities have been shown to relate to the failure of one financial intermediary causing the failure of another.

For macroprudential oversight, policymakers and supervisors need to have access to a broad toolbox of models to measure and analyze system-wide threats to financial stability. Broadly speaking, tools and models can be divided into those for early identification and assessment of systemic risks. ECB (2010) provides a mapping of tools to the above listed three forms of systemic risk: (i) early-warning models, (ii) macro stress-testing models, and (iii) contagion models. First, by focusing on the presence of vulnerabilities and imbalances in an economy, *early-warning models* can be used to derive probabilities of the occurrence of systemic financial crises in the future [see, e.g., Alessi and Detken (2011) and Lo Duca and Peltonen (2013)]. Second, *macro stress-testing models* provide means to assess the resilience of the financial system to a wide variety of aggregate shocks, such as economic downturns [see, e.g., Castrén et al. (2009) and Hirtle et al. (2009)]. Third, *contagion and spillover models* can be employed to assess how resilient the financial system is to cross-sectional transmission of financial instability [see, e.g., IMF (2009)]. In addition, the literature has also provided a large set of coincident indicators to measure the contemporaneous level of systemic risk [see, e.g., Holló et al. (2012)]. While coincident measures may be used to identify, signal and report on heightened stress, they are not designed for early identification and assessment of risk.

This brief review of tools for safeguarding financial stability illustrates the approaches for identification and assessment of potential risks, vulnerabilities and imbalances. Yet, the improvement of *ex ante* prediction results has at the very least been modest, as we have clearly not been able to avert major financial crises. To the

defense of such models, their results have been neglected in the past, such as the signals of the current crisis by Borio and Lowe (2002, 2004). Supported by the fact that build-up phases prior to crises share common characteristics [see, e.g., Reinhart and Rogoff (2009)], early-warning models still have merit for early risk identification. However, one key challenge for risk identification in general and early-warning models in particular is the changing nature of crises, not the least due to financial innovation. Thus, there remains two questions: (i) *How do we better communicate results to policymakers and persuade them to take actions?* and (ii) *How should models be adapted to the changing nature of events that potentially even surpass historical experience?* Stand-alone numerical predictions are unlikely to be the answer to these questions.

In this vein, another conclusion from reviewing tools for safeguarding financial stability is the lack of visual means for identifying and assessing risks and vulnerabilities, particularly in the case of early-warning models. The literature on macroprudential oversight clearly illustrates the lack of integration of a third component, risk communication, with risk identification and assessment tools, an approach that would particularly support external communication. The soar in the availability and precision of data—both in terms of the number of reporting economies and the reporting frequency of the economies—further motivates the development of tools that provide easily interpretable views of complex, high-volume and high-dimensional data, not the least for internal use. In the case of contagion models, as well as macro stress-testing to some extent, visualizations based upon network and graph theory have been and are still gaining further interest within the policymaking community. Yet, the task of representing high-dimensional early-warning indicators on a low-dimensional display has not been addressed in an advanced manner. As a complement to numerical predictions, these visualization tools are a starting point for assessing threats to financial stability. The tools move from artificial intelligence (AI) to intelligence amplification (IA) through the effective use of information technology (IT) in augmenting human intelligence rather than only relying on computational human-like intelligence. IA refers to the notion introduced by Ashby (1957), where he stresses the abilities of human intelligence and the pattern recognition capabilities of the human brain, in particular when augmented with the effective use of visual representations. In this context, Flood and Mendelowitz (2013) note that visualization tools can make a major contribution in assessing systemic risks by pointing to the fact that certain tasks of classification and monitoring can be automated, whereas many require a human analyst, such as the difficulty to train a well-performing machine to analyze anomalous financial market activity.

We are obviously fortunate in that rapid advances in IT have enabled access to massive databases for macroprudential oversight. Alas, analyzing these data is not completely unproblematic. Except for incompleteness of data due to missing values and comparability issues due to cross-country differences in national [e.g., Hartwig (2007)] and firm-level [e.g., Nobes (2006)] accounting practices, as well as outliers and skewed distributions [e.g., Deakin (1976)], the dimensionality of the problem is a central challenge for comprehension. In the case of country-level financial stability, the large variety of sources of financial stress can be measured along several

subdimensions. Generally, the data for macroprudential oversight can be said to be of three types: macroeconomic data, banking system data and market-based data. For instance, the state of a country may be described by data that proxy asset price developments and valuations, credit developments and leverage, as well as more traditional macroeconomic and banking system measures, defined both on a domestic and a global level. Factors that further complicate the assessment of these high-volume and high-dimensional data are temporal and cross-sectional dependencies and relations, invoking also assessments of how risk is distributed in the cross section, e.g., through linkages and exposures among entities.

This particular complexity of the data may be one reason why the interpretability of the monitoring systems has not previously been adequately addressed. As with raw statistical tables, standard two- and three-dimensional plots have, of course, their limitations for high dimensions, not to mention the challenge of including a temporal or cross-sectional dimension or assessing cross sections over time. Although composite indices of leading indicators and predicted probabilities of early-warning models enable comparison across countries and over time, these indices fall short in disentangling the individual sources of vulnerability. More importantly, they lack the ability of preserving similarity relations in data. The recent work by International Monetary Fund (IMF) staff on the Global Financial Stability Map (GFSM) (Dattels et al. 2010) has sought to overcome the challenge of disentangling the sources of distress. The GFSM is a radar chart visualization of six composite indices and has appeared quarterly in the Global Financial Stability Report since April 2007, as well as in a number of other financial stability reports of national central banks. Even here, however, by plotting these types of raw indices, rather than individual indicators, the GFSM leaves a large share of the task of similarity assessment and pattern recognition for the human to solve, where even overall comparability may be questioned (e.g., areas of radar charts scale non-linearly with increases in dimensions and depend on their order).

Methods from the fields of data mining and Knowledge discovery in databases (KDD) may help in overcoming these shortcomings, not the least those for exploratory data analysis (EDA). The notion of EDA was coined by Tukey (1977) and aims at representing data in easily understandable formats through numerical, counting and graphical detective work. Data and dimension reduction methods, and their combination, are common EDA approaches that hold promise for illustrating multivariate data structures in formats easy to comprehend. Data reductions provide overviews of data by compressing information into fewer mean profiles, whereas dimension reductions provide low-dimensional overviews (or mappings) of similarity relations in data. Along these lines, a key focus of this book is to show how data and dimension reduction methods can be applied and extended to support macroprudential oversight. A particular focus of the extensions is related to two tasks in need of future research. First, Chen (2005) and Wong et al. (2012) highlight a paradigm shift from only visualizing structures to visualizing dynamics, not to say dynamics of structures. Second, to be aware of the quality and potential distortions of dimension reductions, Wismüller et al. (2010) and Wong et al. (2012) stress that they are not an

end, but provide only a means to display useful information on top of them, such as evidence, uncertainty and individual data. This leads us to the objectives of the work in this book.

1.2 Research Objectives and Questions

The background of macroprudential oversight illustrates challenges with predictive methods and a lack of visual means for the tasks. This motivates building visual tools for identifying and assessing vulnerabilities and risks. Rather than only relying on conventional early-warning models, a visual approach enables the use of judgment and domain intelligence in combination with the abstractions of data. For data to lend for analysis through summarizations and visualizations, one appealing, yet obvious, approach is the use of data and dimension reduction methods. This section summarizes the two key **research objectives (ROs)** of this book, from which concrete **research questions (RQs)** can be derived.

Yet, before turning to the concrete RQs, the commonly used term *the task at hand* needs to be defined. The overall task is to represent high-dimensional data concerning financial entities, be they countries, markets or institutions, on low-dimensional displays to facilitate the identification, assessment and communication of vulnerabilities and risks. The low-dimensional representations have two subtasks: (i) to function as a display for individual data and their time series (i.e., observation-level data concerning financial entities), and (ii) to use the display as a basis for a wide range of additional visualizations, such as qualities of models and structural properties of data. The final nuance of the task is, however, set by the needs of the domain, data and methods.

In this vein, the work in this book touches upon data and dimension reduction in macroprudential oversight and can hence be divided into two non-mutually exclusive ROs:

- (i) RO1: to choose and extend data and dimension reduction methods such that they meet the needs set by macroprudential oversight and data, and
- (ii) RO2: to apply data and dimension reduction methods in macroprudential oversight to be used by and introduced to the policymaking community.

To deliver on RO1, a large number of steps have to be explored, involving an understanding of the domain, data and methods, in order to compare and extend methods according to the demands set by macroprudential oversight. The applications of methods in RO2 can also be supported by a comprehensive understanding of the needs set by macroprudential oversight, the underlying data and used methods. An indirect implication of RO2, while being a somewhat hazy concept to be measured, is to increase awareness and acceptance of the methods in policy use and to introduce them to the policymaking community, in particular macroprudential oversight.

The two ROs are broad and may hence be refined to more precise RQs that resemble the main steps needed for successful fulfillment of the above stated objectives.




Fig. 1.1 Relations between the RQs. *Notes* The figure shows in the *upper part* the RQs as a process. The interdependence between the RQs is illustrated by separately relating each question to the previous ones. The coloring of the blocks divides the RQs into RTs. The *red blocks* relate to an understanding of the macroprudential domain and data, the *blue blocks* relate to deriving optimal methods and their extensions, and the *green block* relates to applications of the methods to the task at hand. The *color coding* of an RQ in the *lower part* of the figure refers to where its core contribution lies

In particular, whereas the ROs focus on methods and subsequent applicability in policymaking, one key ingredient of successful applied research is that it lies on a strong basis with respect to the domain and the underlying data. This provides three research themes (RTs), where two RTs derived from the above objectives (RT2 and RT3) are preceded by a thorough discussion of the general needs for macroprudential oversight (RT1).

The following five RQs comprise subsequent steps to be addressed in this book:

- (i) RQ1: What are the needs for macroprudential oversight?
- (ii) RQ2: What form do macroprudential data take?
- (iii) RQ3: Which data and dimension reduction methods hold most promise for the task?
- (iv) RQ4: How should the methods be extended and enhanced for the task?
- (v) RQ5: How should the methods and their extensions be applied to the task?

The RQs are interdependent in the sense that they define the process of this book. The process is illustrated as blocks in the upper part of Fig. 1.1, where the color coding shows in which RT each block is a member of. The red blocks relate to an understanding of the macroprudential domain and data, the blue blocks relate to optimal methods and their extensions, and the green block relates to applications of the methods to the task at hand. The lower part of the figure illustrates the interdependence between the RQs by separately relating each question to all other RQs. In the following, this section provides a more detailed discussion of the RQs.

RQ1: What are the needs for macroprudential oversight? The first RQ sets the basis for this book. It focuses on untangling the key tasks of a macroprudential supervisory body that aims at safeguarding financial stability. The aim is to shed light on the functioning of the financial system, and its inherent instabilities and fragilities, as well as the empirical and theoretical literature describing the concepts. Further, a central theme is also a review of related works on tools and models for

macroprudential oversight, and the entire process, to identify shortcomings and needs in the literature.

RQ2: What form do macroprudential data take? An issue of all types of data analysis is the form of the underlying data. The second RQ takes a broad view on macroprudential data and attempts to map them to the needs of macroprudential oversight. In addition to viewing data through the lens of a policymaker, a particular focus is obviously on stylized facts about the data. The key aim of the question, in combination with RQ1, is to provide a solid basis for the rest of the questions and the main objectives of the work in this book.

RQ3: Which data and dimension reduction methods hold most promise for the task? The main aim of the third RQ is to capture the most suitable methods for macroprudential oversight in general and the task at hand in particular. Hence, the basis for the answer to this question lies also in the answers to RQ1 and RQ2, i.e., what are the needs and demands of the domain and data. While an important task for a comparison of data and dimension reduction methods is to review and categorize existing methods, the most central problem is still to identify the methods that hold most promise for the current task.

RQ4: How should the methods be extended and enhanced for the task? With the aim of extending previous methods, the fourth RQ draws upon not only the identified methods, but also the needs and data for the task, involving the answers to RQ1, RQ2 and RQ3. In particular, the identified needs for macroprudential oversight set the needs in terms of data, which both on the other hand impact the chosen method. When then deciding to what direction the methods are to be extended, one needs to consider the limitations of the used methods, in addition to the needs for the task. While not having a substantial focus on human-computer interaction, perception and cognition, a central question is still to bridge the approaches of the so-called machine learning and information visualization communities. The former addresses mainly mathematical and algorithmic aspects of data and dimension reductions, whereas the latter focuses on visual representations of abstract data, and the interaction of humans, to reinforce cognition.

RQ5: How should the methods and their extensions be applied to the task? In the fifth RQ, the focus is on applications of the provided methods as tools for macroprudential oversight in general and the task at hand in particular. Again, the answer to this question relies upon the answers to all the previous questions: RQ1, RQ2, RQ3 and RQ4. The key use of the applications derive from the needs for macroprudential oversight and data, as discussed in the first two questions. The approach, on the other hand, comprises the methods provided in RQ3 and RQ4.

1.3 Overview of the Book

This work in this book aims at indirectly meeting the ROs by following the process set by the RQs. Figure 1.2 uses the same representation as in Fig. 1.1, but relates the chapters to the process of RQs. In the following, this section discusses chapter by chapter the structure of this book.

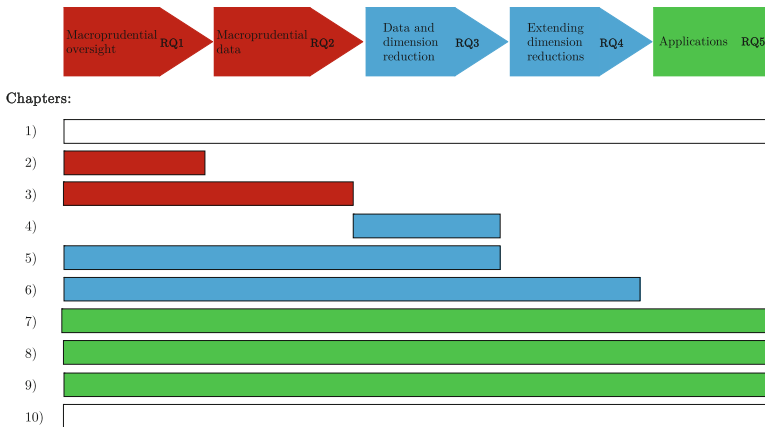


Fig. 1.2 The RQs and the chapters. *Notes* The figure shows in the *upper part* the RQs as a process. These are related to the chapters of the book. The *red blocks* relate to an understanding of the macroprudential domain and data, the *blue blocks* relate to deriving optimal methods and their extensions, and the *green block* relates to applications of the methods to the task at hand. The *color coding* of a chapter refers to where its core contribution lies. The introduction in this chapter and conclusions in Chap. 10 are intentionally left blank

Chapter 2 focuses first on the definition of financial systems, financial instability and systemic risks, as well as on the reasons for financial systems being fragile. Next, it briefly summarizes some theoretical and empirical underpinnings of system-wide risks, whereafter the chapter focuses on giving an overview of the state of the art of risk assessment and identification tools used by macroprudential policymakers, especially the use of visualization tools. The chapter concludes by relating the fragilities, risks and tools to an overall macroprudential oversight process. The process clearly illustrates the lack of integration of a third component, risk communication, with risk identification and assessment tools. A key notion for this book is the three forms of systemic risk and the respective risk identification and assessment tools, whereas a key implication is the illustrated scarcity of visualization tools, in particular for the task of identifying the build-up of widespread imbalances.

Chapter 3, while heavily relying on the previous chapter, discusses data needs and demands for macroprudential oversight, with a particular focus on early-warning models. The broad notion of macroprudential data is untangled into a four-dimensional cube representation. Finally, this chapter discusses stylized challenges related to macroprudential data. A key implication of the chapter is that the shown characteristics of macroprudential data need to be acknowledged when attempting their use to support macroprudential oversight. More importantly, rather than aggregating data into composite indices, the chapter further motivates visualizing these complex data in easily understandable formats to support disciplined and structured judgmental analysis based upon policymakers’ experience.

Chapter 4 provides an overview of data and dimension reduction methods. First, it discusses the relation of data and dimension reduction to knowledge discovery, data mining, information visualization and visual analytics. Then, the chapter sets a basis for a comparison of data and dimension reduction methods by reviewing the basics of classical methods and relating a comprehensive set of methods in a taxonomy.

Chapter 5 relates the needs for macroprudential oversight and properties of macroprudential data to the characteristics of data and dimension reductions, and their combinations. The suitability of three classical, or so-called first-generation, dimension reduction methods for the task at hand is illustrated with qualitative comparisons and illustrative experiments. A key implication of the chapter is that the Self-Organizing Map (SOM)s holds most promise for the task at hand.

Chapter 6 presents a number of extensions of the SOM to meet the needs and demands for both macroprudential oversight and macroprudential data. The enhancements not only aid in analyzing and visualizing individual cross-sectional and/or time-series data on the SOM, but also contribute to the assessment of overall properties and qualities of the SOM. Extensions to be used with a standard SOM comprise approaches for fuzzification, transition probabilities and assessing shock propagation. The chapter also presents the stand-alone Self-Organizing Time Map (SOTM) for assessing how cluster structures evolve over time (i.e., visual dynamic clustering).

Chapter 7 describes the construction of the Self-Organizing Financial Stability Map (SOFSM). First, the chapter presents the used data, including macro-financial indicators and a database of financial crises, a model evaluation framework and a model training framework. Then, the training and evaluation frameworks are applied for constructing the SOFSM based upon the standard SOM. Finally, the chapter presents a number of robustness tests on the final SOFSM. The SOFSM can be used to monitor macro-financial vulnerabilities by locating a country in the financial stability cycle on a two-dimensional display. Besides of its visualization capabilities, the SOFSM is evaluated as an early-warning model and calibrated according to policymakers' preferences between missing a crisis and issuing a false alarm (i.e., type I and II errors). The SOFSM performs on par with a statistical benchmark model and correctly calls the crises that started in 2007 in the United States (US) and the euro area.

Chapter 8 applies the SOFSM for risk identification, assessment and communication by a mapping of financial stability. Thus, the extensions in Chap. 6, except for the SOTM, are applied to macroprudential oversight in this chapter, including a fuzzification, transition probabilities and shock-propagation analysis. The chapter also shows how the SOFSM can be used for illustrating results of stress tests and detecting outliers (i.e., imbalances in macro-financial conditions). The SOFSM is also paired with a stand-alone predictive model to illustrate the complementary role of such approaches. Hence, the SOFSM not only provides means for visual early-warning exercises, but also enable superimposed visualizations of stress tests and shock-propagation assessments.

Chapter 9 applies the SOTM to macroprudential oversight in general and risk identification in particular by providing two decompositions of global financial crises. The SOTM performs temporal data and dimension reduction for visual dynamic

clustering. The first decomposition applies a standard SOTM to describe the global financial crisis that started in 2007. The second section uses a SOTM on time-to-event data to generalize patterns before, during and after financial crises.

Chapter 10 concludes with a discussion of the key contributions of this book, its limitations and suggestions for future research. The contributions are discussed both from the viewpoint of dimension reduction and policy (i.e., the research objectives), whereas limitations set a basis for future research, in addition to the numerous other questions that remain to be explored.

1.4 Original Publications and Their Contributions

As this book is based upon a number of publications, interested readers are kindly referred to the original publications for further information. A large share of the material is based upon the contribution, text, figures and tables found in the following 11 papers:

- Paper 1 Sarlin P. Data and Dimension Reduction for Visual Financial Performance Analysis. *Information Visualization*, forthcoming, Sage pub.
- Paper 2 Sarlin P, Eklund T, 2011. Fuzzy Clustering of the Self-Organizing Map: Some Applications on Financial Time Series. *Proceedings of the 8th International Workshop on Self-Organizing Maps (WSOM'11)*, Helsinki, Finland, June 13–15, pp. 40–50, Springer.
- Paper 3 Sarlin P, Yao Z, Eklund T, 2012. Probabilistic Modeling of State Transitions on the Self-Organizing Map: Some Temporal Financial Applications. *Intelligent Systems in Accounting, Finance and Management* 19(1), pp. 189–203, Wiley-Blackwell.
- Paper 4 Sarlin P, 2013. Self-Organizing Time Map: An Abstraction of Temporal Multivariate Patterns. *Neurocomputing* 99(1), pp. 496–508, Elsevier.
- Paper 5 Sarlin P, Yao Z, 2013. Clustering of the Self-Organizing Time Map. *Neurocomputing* 121, pp. 317–327, Elsevier.
- Paper 6 Sarlin P, 2013. On policymakers' loss functions and the evaluation of early warning systems. *Economics Letters* 119(1), pp. 1–7, Elsevier.
- Paper 7 Sarlin P, Peltonen TA, 2013. Mapping the State of Financial Stability. *Journal of International Financial Markets, Institutions & Money* 26, pp. 46–76, Elsevier.
- Paper 8 Sarlin P, 2013. Exploiting the Self-Organizing Financial Stability Map. *Engineering Applications of Artificial Intelligence* 26(5–6), pp. 1532–1539, Elsevier.
- Paper 9 Sarlin P, 2013. Decomposing the Global Financial Crisis: A Self-Organizing Time Map. *Pattern Recognition Letters* 34, pp. 1701–1709, Elsevier.

- Paper 10 Sarlin P, 2013. A Self-Organizing Time Map for Time-to-Event Data. *Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining (CIDM'13)*, Singapore, April 16–19, 2013, IEEE .
- Paper 11 Sarlin P. On biologically inspired predictions of the global financial crisis. *Neural Computing & Applications*, forthcoming, Springer.

References

- Alessi, L., & Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity. *European Journal of Political Economy*, 27(3), 520–533.
- Ashby, R. (1957). *An introduction to cybernetics*. London: Chapman & Hall.
- Bezemer, D. (2011). The credit crisis and recession as a paradigm test. *Journal of Economic Issues*, 45(1), 1–18.
- Bisias, D., Flood, M., Lo, A., & Valavanis, S. (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics*, 4, 255–296.
- Bordo, M., Eichengreen, B., Klingebiel, D., & Martinez-Peria, M. (2001). Is the crisis problem growing more severe? *Economic Policy*, 32, 51–82.
- Borio, C., & Lowe, P. (2002). Asset prices, financial and monetary stability: Exploring the nexus. BIS working papers, No. 114.
- Borio, C., & Lowe, P. (2004). Securing sustainable price stability: Should credit come back from the wilderness? BIS working papers, No. 157.
- Cardarelli, R., Elekdag, S., & Lall, S. (2011). Financial stress and economic contractions. *Journal of Financial Stability*, 7(2), 78–97.
- Carletti, E. (2008). Competition and regulation in banking. In A. Boot & A. Thakor (Eds.), *Handbook in Financial Intermediation* (pp. 449–482). North Holland: Elsevier.
- Castrén, O., Fitzpatrick, T., Sydow, M. (2009). Assessing portfolio credit risk changes in a sample of EU large and complex banking groups in reaction to macroeconomic shocks, ECB working paper, No. 1002.
- Chen, C. (2005). Top 10 unsolved information visualization problems. *IEEE Computer Graphics and Applications*, 25(4), 12–16.
- Dattels, P., McCaughrin, R., Miyajima, K., Puig, J. (2010). Can you map global financial stability? IMF working paper, No. 10/145.
- de Bandt, O., & Hartmann, P. (2002). Systemic risk in banking: A survey. In C. Goodhart & G. Illing (Eds.), *Financial crisis, contagion and the lender of last resort: A book of readings*. Oxford: Oxford University Press.
- de Bandt, O., Hartmann, P., & Peydro, J. (2009). Systemic risk in banking: An update. In A. Berger & M. P. Wilson, J. (Eds.), *Oxford Handbook of Banking*. Oxford: Oxford University Press.
- Deakin, E. (1976). Distributions of financial accounting ratios: Some empirical evidence. *The Accounting Review*, 51, 90–96.
- Dell'ariccia, G., Detragiache, E., & Rajan, R. (2008). The real effects of banking crises. *Journal of Financial Intermediation*, 17, 89–112.
- Demirgüç-Kunt, A., & Detragiache, E. (1998). The determinants of banking crises in developing and developed countries. *IMF Staff Papers*, 45(1), 81–109.
- ECB. (2009). The concept of systemic risk. In: *Financial stability review* (December 2009). Frankfurt, Germany: European Central Bank.
- ECB. (2010). Analytical models and tools for the identification and assessment of systemic risks. In: *Financial stability review*, (June 2010). Frankfurt. Germany: European Central Bank.

- Eichengreen, B. (2004). Financial instability. Paper written on behalf of the Copenhagen Consensus and presented in Copenhagen on May 25–28, 2004.
- Fell, J., & Schinasi, G. (2005). Assessing financial stability: Exploring the boundaries of analysis. *National Institute Economic Review*, 192, 102–117.
- Flood, M., & Mendelowitz, A. (2013). Monitoring financial stability in a complex world. In V. Lemieux (Ed.), *Financial analysis and risk management data governance* (pp. 15–45). Springer-Verlag, Heidelberg: Analytics and Life Cycle Management.
- Gorton, G. (1988). Banking panics and business cycles. *Oxford Economic Papers*, 40(4), 751–781.
- Hartmann, P. (2009). The financial crisis, systemic risk and macro-prudential supervision, keynote address at the Conference on contemporary issues in financial markets and institutions (December, 9). City University London, London.
- Hartwig, J. (2007). Is the transatlantic gap in economic growth really widening? In J. McCombie & C. Rodriguez (Eds.), *The European Union: Current problems and prospects* (pp. 68–83). Basingstoke: Palgrave-Macmillan.
- Hirtle, B., Schuermann, T., & Stiroh, K. (2009). Macroprudential supervision of financial institutions: Lessons from the scap, Staff Report No. 409. New York: Federal Reserve Bank of New York.
- Hoggarth, G., Reis, R., & Saporta, V. (2002). Costs of banking system instability: Some empirical evidence. *Journal of Banking & Finance*, 26(5), 825–855.
- Holló, D., Kremer, M., & Lo Duca, M. (2012). CISS - a composite indicator of systemic stress in the financial system. ECB working paper, No. 1426.
- IMF. (2009). *Assessing the systemic implications of financial linkages, global financial stability report April*. Washington: International Monetary Fund.
- Laeven, L., & Valencia, F. (2008). Systemic banking crises: A new database. IMF working paper, 08/224.
- Lo Duca, M., & Peltonen, T. (2013). Assessing systemic risks and predicting systemic events. *Journal of Banking & Finance*, 37(7), 2183–2195.
- Minsky, H. (1982). *Can "it" Happen Again?: Essays on Instability and finance*. Armonk, New York: M.E. Sharpe.
- Nobes, C. (2006). The survival of international differences under IFRS: Towards a research agenda. *Accounting and Business Research*, 36(3), 233–245.
- Reinhart, C., & Rogoff, K. (2008). Is the 2007 us sub-prime financial crisis so different? an international historical comparison. *American Economic Review*, 98(2), 339–344.
- Reinhart, C. M., & Rogoff, K. (2009). *This time is different: Eight centuries of financial folly*. Princeton: Princeton University Press.
- Schou-Zibell, L., Albert, J., Song, L. (2010, March). A macroprudential framework for monitoring and examining financial soundness. ADB working paper series on regional economic intergration, Asia: Asian Development Bank.
- Tukey, J. (1977). *Exploratory data analysis*. Reading, PA: Addison-Wesley.
- Upper, C., & Worms, A. (2004). Estimating bilateral exposures in the german interbank market: Is there a danger of contagion? *European Economic Review*, 48(4), 827–849.
- van Lelyveld, I., & Liedorp, F. (2006). Interbank contagion in the Dutch banking sector: A sensitivity analysis. *International Journal of Central Banking*, 2(2),
- Wismüller, A., Verleysen, M., Aupetit, M., & Lee, J. (2010). Recent advances in nonlinear dimensionality reduction, manifold and topological learning. Proceedings of the European Symposium on artificial neural networks (ESANN 10) (pp. 71–80). Bruges, Belgium.
- Wong, P., Shen, H., Johnson, C., Chen, C., & Ross, R. (2012). The top 10 challenges in extreme-scale visual analytics. *IEEE Computer Graphics and Applications*, 32(4), 63–67.

Chapter 2

Macroprudential Oversight

In the absence of clear guidance from existing analytical frameworks, policy-makers had to place particular reliance on our experience. Judgement and experience inevitably played a key role. [...] But relying on judgement inevitably involves risks. We need macroeconomic and financial models to discipline and structure our judgemental analysis. How should such models evolve?

–Jean-Claude Trichet, President of the ECB, Frankfurt am Main,
18 November 2010

Paraphrasing Milton Friedman’s statement about Keynesians, Borio (2011) stated “*We are all macroprudentialists now*”. Since the date when the still ongoing global financial crisis broke out, the notion of a macroprudential approach to safeguarding financial stability has grown consensus among the academic and policymaking communities alike. Yet, it is by no means a new concept. The central bank of central banks, the Bank for International Settlements (BIS), applied the term to describe a system-wide orientation of regulatory frameworks already in the 1970s, and the term appeared in publicly available material in the mid-1980s [see, e.g., BIS (1986) and Borio (2011)], but the use of the concept remains somewhat ambiguous.

So, what is a macroprudential vis-à-vis a microprudential approach? With the help of a comparison to the microprudential approach, Borio (2011) summarizes the macroprudential orientation as follows. First, while the aim of the macroprudential approach is to limit system-wide stress and possible costs for the macroeconomy, a microprudential orientation attempts to limit an individual institution’s risk of failure with the aim of minimizing costs for depositors and investors. Second, the macroprudential approach explicitly accounts for the fact that risk is dependent on the collective behavior of financial institutions (i.e., endogenous), rather than being something outside their influence (i.e., exogenous) as is in the microprudential case.

This chapter is partly based upon previous research. Please see the following work for further information: Sarlin (2014a).

Third, the macroprudential approach has a system-wide perspective, where a top-down approach works out a desirable safety standard for the system as a whole, rather than the stand-alone soundness of individual institutions approached from the bottom-up. Thus, a macroprudential approach takes a holistic view on the financial system with the aim and mandate to ensure system-wide stability, rather than only being concerned with the failure of individual entities. Yet, the two approaches are difficult to compartmentalize because they most often co-exist.

The comprehensive macroprudential approach thus obviously also involves an understanding of a large number of other concepts. It is crucial for regulatory decision-makers to have a broad and deep understanding of financial systems, fragilities and instabilities, as well as risks and vulnerabilities, in the economy. Hence, to carry out macroprudential oversight aiming at ensuring system-wide stability, policymakers need a thorough information basis and a large variety of risk identification and assessment models and tools for data to become actionable information. Macroprudential oversight, while also requiring a large share of domain intelligence and plain analysis of statistical data, has its core in analytical models and tools for analyzing, summarizing and interpreting the widely available masses of data.

This chapter focuses first in Sect. 2.1 on the definition of financial systems and financial stability—or rather its antithesis, financial instability—as well as fragilities in financial systems and the concept of systemic risk. Section 2.2 briefly summarizes some theoretical and empirical underpinnings of three identified forms of systemic risk. Then, Sect. 2.3, and the main focus of this chapter, attempts to give an overview of the state of the art of risk assessment and identification tools used by macroprudential policymakers, especially the use of visualization tools. Finally, Sect. 2.4 relates the fragilities, risks and tools to the macroprudential oversight process, to be followed by a summary of key implications of this chapter for the rest of the book in Sect. 2.5.

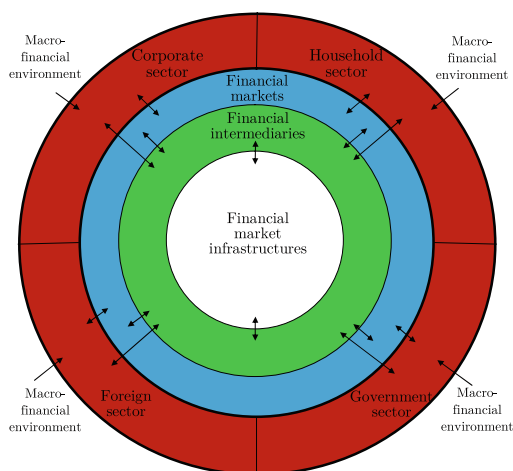
2.1 Financial Systems, Fragilities and Instabilities

Understanding the key concepts related to financial systems, and their (in)stability and fragility, is essential for a broader understanding of macroprudential oversight. This section presents some key principles of financial systems and defines the notions of stability and instability, as well as discusses why they are so fragile and what are the main risks to stability. Hence, this section provides a basis for the rest of the chapter, not the least by untangling systemic risks into three forms, to which we oftentimes refer in the sequel.

2.1.1 Key Components of Financial Systems

The basis for any discussion of financial fragilities, instabilities or risks ought to be an understanding of the notion of a financial system. Hence, the main questions are: *What is a financial system, which components does it comprise and how do they interact?*

Fig. 2.1 The financial system and its components. *Notes* The figure shows interrelations of the three components of the financial system: financial markets, intermediaries and market infrastructures. The figure is an adapted version of that in ECB (2012a) and follows the description in Fell and Schinasi (2005)



Broadly speaking, financial markets may be thought of as a mechanism for people to trade various financial securities, commodities and other fungible items at prices that reflect the markets. In a larger context, Schinasi (2004) summarizes the key functions of a financial system in fostering and supporting the real economy by matching investors with savers, allocating and pricing financial risks and resources and supporting various intertemporal economic processes like wealth accumulation, economic growth, and social prosperity. However, the functioning of financial systems is a multifaceted concept with multiple inter and intra relationships. Key components of financial systems, as well as their relationships, are illustrated in Fig. 2.1. As is pointed out in the figure, the financial system comprises three interrelated, yet separable, components [see, e.g., ECB (2005) and Fell and Schinasi (2005)]:

- (i) financial intermediaries (green layer);
- (ii) financial markets (blue layer); and
- (iii) financial market infrastructures (white layer).

Following the description in Fell and Schinasi (2005), entities of the household, corporate, foreign and government sectors (red layer) invest their savings and obtain funding for their activities through these three components. First, *financial intermediaries* comprise mainly financial institutions and have as their main task to pool risks and funds of one counterparty and allocate them to another. Financial institutions provide a wide range of services, in addition to those traditionally provided by banks. Depending on their profile, e.g., insurers, banks, pension funds, hedge funds and hybrids of financial and non-financial companies (e.g., General Electric) provide multiple different types of financial services. Second, *financial markets* mainly aim at matching those who need capital with those who have it (i.e., spenders with savers). The trading of financial securities (e.g., stocks and bonds), commodities (e.g., precious metals and agricultural goods) and fungible items in general occurs between people and firms, be they financial or not. For financial markets to support

the provision of credit, transfer of risk and risk management in general, it is crucial that they function smoothly and are resilient under various circumstances. Third, the *financial infrastructure* of the financial system is comprised of privately and publicly owned and operated institutions through which financial market operations are concretely carried out. The infrastructure may be provided by institutions like payment, clearing and settlement systems for financial transactions and other types of monetary, legal, accounting, regulatory, supervisory and surveillance infrastructures. Payment systems commonly transfer funds electronically from one institution to another, clearing systems commonly transfer credit risk in the derivatives market to a clearinghouse from each counterparty of a trade, and settlement systems complete transactions like securities trades. Thus, we herein follow Schinasi (2004) by defining the financial system as a term that encompasses “*both the monetary system with its official understandings, agreements, conventions, and institutions as well as the processes, institutions, and conventions of private financial activities*”.

While financial intermediaries connect to the financial architecture, the household, corporate, government and foreign sectors are connected both directly and indirectly to financial intermediaries, where financial markets may function as a middleman. Like private market participants, governments may borrow in markets and hedge risks. The working principles of the financial system, and the general performance of its key tasks, is based upon these components and their interrelations. Further, the external macro-financial environment will not only have a direct impact on private and public participants, but will also indirectly affect the functioning of financial markets and intermediaries, and in some cases even affect the design of infrastructures.

It is hence obvious to conclude that a resilient and well-functioning financial system is characterized by well-managed financial institutions and efficiently functioning financial markets, as well as by a strong and robust financial infrastructure. This might also be associated with less frequent and costly incidences of financial crisis. The fact that we have experienced frequent incidences of financial crisis does, however, indicate that financial institutions are not always well-managed, the functioning of financial markets may be inefficient, and the financial infrastructures may have cracks and weaknesses. Before discussing the reasons to this, we need a working definition of financial stability, particularly its antithesis.

2.1.2 Financial (in)stability

The term financial stability, not to paraphrase Justice Potter Stewart once again, belongs to the group of concepts that are broad and vague, yet implicitly understood. Still, we need to agree upon the definition of stable and unstable financial systems before delving into the causes of fragilities and risks. Coining financial stability with a commonly accepted and used definition has indeed been an elusive goal ever since it has shifted towards a common policy objective. In spite of numerous proposals, there is, as yet, no single, widely accepted definition for the concept.

Some define financial stability broadly, such as “*a condition where the financial system is able to withstand shocks*” (Padoa-Schioppa 2003), while others focus on situations when the financial system supports, rather than impedes, the functioning of the real economy [e.g., Schinasi (2004)]. However, guided by macroprudential thinking, with an aim to ensure system-wide stability, the definition of financial stability ought to be narrowed down along those lines. ECB (2009) provides a somewhat long, but descriptive definition of system-wide stability: “*a condition in which the financial system—comprising of financial intermediaries, markets and market infrastructures—is capable of withstanding shocks and the unravelling of financial imbalances, thereby mitigating the likelihood of disruptions in the financial intermediation process which are severe enough to significantly impair the allocation of savings to profitable investment opportunities*”.

Via its antithesis, Allen and Wood (2006) favor to define a financially stable system as simply one: “*which is not prone to episodes of financial instability*”. This leads to the question: What is financial instability? While being somewhat easier to define, also a broad variety of definitions of financial instability exist. We may want to call it “*a situation in which normal-sized shocks to the financial system are sufficient to produce financial distress*” (Borio and Drehmann 2009b) or “*any deviation from the optimal saving—investment plan of the economy that is due to imperfections in the financial sector*” (Haldane et al. 2004). From the sample definitions, it is easy to see the lack of unanimity with regards to these concepts.

Again, to meet the demands of a macroprudential approach, we narrow down from the broad concept of financial instability to systemic financial crises or strong systemic events. Such a crisis may be defined as an event that “*adversely affects a number of systemically important intermediaries or markets*” (ECB 2009). Rather than only being interested in the systemic events *per se*, an obvious central theme is to have an understanding of the underlying risk of experiencing a systemic financial crisis, i.e., systemic risks. In broad terms, systemic risk is defined as “*the risk that financial instability becomes so widespread that it impairs the functioning of a financial system to the point where economic growth and welfare suffer materially*” (ECB 2009).

Now, when we have defined the concepts of a stable and unstable financial system, we can move forward in discussing what makes financial systems particularly fragile and what are the underlying risks to stability.

2.1.3 Fragility of Financial Systems

Any form of systemic risk, while having sources of its own kind, is most often preceded at an early stage by various market imperfections. Imperfections in markets may take the form of asymmetric and incomplete information, externalities and public-good characteristics, incomplete markets, etc., and are to some extent present in most economic sectors. However, the imperfections, when being related to a financial sector, may lead to significant fragility of not only individual entities,

but also the entire system (Carletti 2008; ECB 2009). Carletti (2008) illustrates the need for regulation of the banking sector with a large sample of examples of market imperfections, such as banks being exposed to deposit runs due to the maturity transformation by investing short-term deposits in long-term assets and informational asymmetries between depositors and borrowers, as well as debtholders and firm managers having so-called misaligned principal-agency problems, leading to agents not acting in the best interest of the principal. Another example is the parallel of financial stability to a public good and its absence to externalities like pollution, as each entity manages its own risks with no need to consider its impact on the system-wide risk as a whole. Bandt and Hartmann (2002) relate fragilities in financial systems to three causes:

- (i) the strong information intensity and intertemporal nature of financial contracts and transactions;
- (ii) the balance-sheet structures of financial intermediaries with a high reliance on debts or leverage, and maturity mismatches between assets and liabilities; and
- (iii) the high degree of interconnectedness between financial intermediaries and markets.

In the following, this subsection focuses on the above mentioned three main features behind the fragility of financial systems as identified by Bandt and Hartmann (2002).

First, the *information intensity and control intensity* relates to the fact that financial decisions concern intertemporal allocation of purchasing power [see, e.g., Stiglitz (1993)]. This relates to the issue of asymmetric information, in which lenders do not have full information about the intentions of the borrower, such as whether or not they are capable and/or willing to repay their debt. Likewise, the intertemporal nature leads to an inherent need for a lender to trust either the borrower to repay her debt or a third party to enforce the contract, not the least as the intertemporality leaves room for renegotiations of contracts. Thus, the decisions have their basis in whether the outcome of future asset values and future cash flows promised in contracts will meet expectations, such as is the case with deposit contracts. Another obstacle is changes in uncertainty affecting investment and disinvestment decisions [see, e.g., Shiller (1989)]. This leads, for instance, to substantial changes in asset prices not being explained by their fundamentals (e.g., companies' earnings and inflation rates are fundamentals to shares and exchange rates, respectively).

Second, the *maturity-mismatch structure of banks* is described by taking fixed-value deposits and enabling them to be withdrawn at a short notice, as well as by lending long term to the industry [see, e.g., Bryant (1980)]. When exceptionally high withdrawals occur and long term loans cannot be liquidated, the small fraction of held reserves may lead to insolvency. Hence, the strength of a bank depends on both the capability of lending to profitable investment projects and the confidence of depositors on the bank's loan book, as well as the confidence that other depositors will not run the bank. Yet, the better the deposit insurance scheme the less likely are confidence crises. While many fragilities relate to financial intermediaries in general, Goodhart et al. (1998) note that these types of confidence problems do most often

only apply to banks, except for cases when the non-bank intermediary is a part of the same entity as a bank.

Third, the complex *interconnectedness and network structure of banks* in particular and financial intermediaries in general implies that the failure of one bank may affect others [see, e.g., Humphrey (1986) and Folkerts-Landau (1991)]. Bandt and Hartmann (2002) relate the networks of real exposures among banks to consist partly of interbank lending and partly of those in wholesale and retail payment and settlement systems. While the aim of the interbank lending market is to provide a channel for short-term lending and borrowing to banks, a sudden low transaction volume in this market due to various reasons may lead to liquidity problems, such as during the financial crisis of 2007. Likewise, the exposures in payment and settlement systems may be large enough for a failure to meet payment obligations of one bank to impact the capability of other banks fulfilling their payment obligations. This could subsequently lead to failures spreading through amplified domino effects. However, the better the risk management measures, margin requirements and portfolio insurance, the more robust are payment and settlement systems (Bandt and Hartmann 2002).

These particularities support the role of governments and other supervisory authorities in addressing and monitoring systemic risks.

2.1.4 A Systemic Risk Cube

Above, we discussed systemic risk in broad terms, whereas the inherently complex issue can reasonably not be covered by such a simple definition. Hence, there is a need for a more precise and structured definition. To give some structure to the concept, the definition used herein is untangled with the help of the systemic risk cube (henceforth the risk cube) shown in Fig. 2.2. The risk cube presented here is an adapted version of that in ECB (2010). It represents the European Central Bank (ECB)'s conceptual framework for systemic risk and has its origin in the works by Bandt et al. (2009), ECB (2009), Trichet (2009) and ECB (2010). The sequel of this chapter is to a large extent guided by, and often paired with, the systemic risks identified through the risk cube.

Due to the great complexity of systemic risk, a virtue of the risk cube is that it not only helps untangling the forms of systemic risks, but also enables a subsequent mapping of them to the theoretical and empirical literature, as well as to analytical tools for identification and assessment of risks. The three dimensions of the risk cube are the *triggers*, *origins* and *impacts*. The nature of triggers unleashing the crisis could take the form of an *exogenous* shock, which stems from the outside of the financial system (e.g., a macro-economic shock and events like natural disasters or political turmoil), or could emerge *endogenously* from within the financial system or some other part of the economy (e.g., from financial intermediaries, markets and infrastructures). The origins of the events may be distinguished to limited *idiosyncratic* shocks and widespread *systematic* shocks. While idiosyncratic shocks




Fig. 2.2 Systemic risk cube with three forms of risks. *Notes* The figure represents the systemic risk cube with three dimensions and systemic risks, as well as possible market imperfections underlying systemic risk. It is an adapted version of that in ECB (2010)

are those that initially affect only the health of a single financial market, financial intermediary or asset, systematic shocks are those that, in the extreme, affect the financial system as a whole, such as the entire banking sector. Here, it is important to pay regard to the differentiation of the terms *systemic* and *systematic*. Bandt and Hartmann (2002) note that a systematic shock may cause a systemic event, but a systemic event does not need to have its origin in a wide systematic shock. Further, the impact of the events may be divided into those causing problems for a range of financial intermediaries and markets in a *sequential* and *simultaneous* fashion.

To reduce the complexity of the risk cube, combinations of its elements (the triggers, origins and impacts) may be limited to the materialization of three broad and interrelated forms of systemic risk (see Fig. 2.2):

- (i) endogenous build-up and unraveling of widespread imbalances (red boxes);
- (ii) exogenous aggregate shocks (blue boxes); and
- (iii) contagion and spillover (green boxes).

The first form of systemic risk refers to the risk that *widespread imbalances*, that have built up over time, *unravel abruptly*. The underlying problems are caused by an endogenous build-up of imbalances in one or several parts of a financial system, such as high concentrations of lending in certain parts of the economy or credit booms in general. While these imbalances, some may even say bubbles, may in the short term last with mainly profitable implications, a shock leading to a repricing of risk may be triggered by even a small event or change in expectations. This resembles Kindleberger’s (1978) and Minsky’s (1982) financial fragility view of a boom-bust credit or asset cycle. Hence, the subsequent abrupt unraveling of the imbalances may be endogenously or exogenously caused by idiosyncratic or systematic shocks, and may have adverse effects on a wide range of financial intermediaries and markets in a simultaneous fashion. Second, systemic risk may also refer to a widespread *exogenous aggregate shock* that has negative systematic effects on one or many financial intermediaries and markets at the same time. For instance, if banks go bad during recessions, they can be said to be vulnerable to economic downturns. The

third form of systemic risk is *contagion and spillover*, which usually refers to an idiosyncratic problem, be it endogenous or exogenous, that spreads in a sequential fashion in the cross section. For instance, a failure of one financial intermediary causing the failure of another financial intermediary, which initially seemed solvent, was not vulnerable to the same risks and was not subject to the same original shock as the former. It is worth noting that contagion refers to a situation when the initial failure is entirely responsible for subsequent ones, whereas the term spillover is commonly used when the causal relationship is not found or cannot be tested [see, e.g., ECB (2010)].

A categorization of systemic risks into the three forms provides means for a further discussion on the empirical and theoretical literature.

2.2 Theoretical and Empirical Underpinnings

This section draws upon the above defined terms and concepts. In light of the above discussion, the section reviews and discusses theoretical and empirical works on systemic risk. In both subsections, three parts match the identified forms of systemic risk. This chapter draws upon literature reviews in Bandt and Hartmann (2002), Bandt et al. (2009) and ECB (2009), in addition to a wide range of other sources to which in-text references are provided.

2.2.1 Theoretical Models

This subsection discusses the theoretical literature related to the three forms of systemic risk. While the literature is currently developing at a tremendous pace, many important older works continue to be relevant. We start by discussing the literature on lending booms and build-ups of imbalances that goes half a century back in time, then we focus on theoretical works on macroeconomic aggregate shocks to the economy, and finally on the literature on interbank contagion.

Endogenous Build-up of Widespread Imbalances

The notion of financial fragility and lending booms relates back to early work by Minsky (1977, 1982) and Kindleberger (1978), who pinpointed common historical reasons for financial crises to be the endogenous build-up and abrupt unraveling of widespread imbalances. The early authors explain the boom and bust cycle as follows. The imbalances oftentimes derive from the pro-cyclicality of financial behavior; in good times consumption and investment increases, which generates income, and further fuels consumption and investment. During this time of “euphoria” and “gregarious behavior”, the financial activities become more speculative, or even so-called

Ponzi finance, in which a lack of expected income flows causes a reliance on the rise of market value of assets or income to pay off interest or principal. In this “virtuous” circle, risks are often neglected with mainly profitable implications in the short term. Then, even a small trigger, shock, change in expectations, or other type of event, be it exogenous or endogenous, may lead to a repricing of risk, an end of the boom, unraveling of imbalances and possibly simultaneous adverse effects to intermediaries and markets. This event may even be called a Minsky moment—a term coined by the managing director of PIMCO, Paul McCulley, in 1998 when describing the Asian financial crisis. The early literature has its core in uncertainty rather than only risk, such as the discussion on the relation between Knightian uncertainty and investment returns and risk premiums in Guttentag and Herring (1984). The same authors also explain disaster myopia by subjective probabilities of disastrous events diminishing when time elapses after the previous realization of such an event. The here described characteristics of a financial stability cycle emerges, according to the early authors, endogenously in economies with particularly unregulated financial markets.

There are a number of reasons to the build-up of imbalances, of which four key notions are summarized, as is categorized in ECB (2009). *First*, financial markets are inherently featured by herd behavior, leading to entities sharing similar risks. Banerjee (1992) and Bikhchandani et al. (1992) describe these as rational herding waves, if relative returns of investments are highly uncertain. Likewise, the herding by Scharfstein and Stein (1990) involves investment or fund managers and loan officers that mimic each other when they are evaluated, which steers pay or reputation, in relation to the rest of the market. *Second*, the so-called curse of low interest rates may diminish incentives to screen borrowers when interest rates are low [see, e.g., Dell’ariccia and Marquez (2006)]. Low interest rates over a wide maturity spectrum have more often than not been quoted as an element of the imbalances prior to the current crisis. Another obvious channel is an increase in collateral values, such as real estate prices, when interest rates are low. For further discussions on the effect of low rates on crises, see Allen and Gale (2007). *Third*, positive shocks to collateral, while enhancing the borrowing capacity in an economy, may also contribute to leverage cycles (Kiyotaki and Moore 2002). When an industry, or another industry with similar collateral, benefits from an increase in collateral value, it also allows more borrowing and investment, and thus further amplifies leverage. Likewise, Geanakoplos (2010) asserts that variation in leverage impacts volatility in asset prices and thus contributes to financial booms and busts. He explains it by there being high-leverage buyers for whom an asset is more valuable than it is for others, for instance, due to them being more sophisticated investors, better in hedging exposures to the assets or less risk averse. This drives prices up, whereas losses in wealth will, due to leverage, move the assets into more pessimistic hands, which again amplifies the decrease in value. *Fourth*, risk-taking and moral hazard may also be amplified by better safety net provisions. One example is a decrease in depositors’ incentives to screen bank risks through deposit insurance [see, e.g., Boot and Greenbaum (1993)]. Similar effects can be derived from public bailouts or lenders of last resort, that is, an institution providing credit in the lack of other sources and with the aim of preventing failures of important institutions.

Exogenous Aggregate Shocks

It is no new notion that macroeconomic shocks or economic downturns have been a trigger of many historical financial crises [see, e.g., Gorton (1988)]. Yet, the theoretical literature directly addressing the topic is somewhat scarce. Even though direct interbank connections and contagion is missing, banking crises have still occurred simultaneously with aggregated shocks. Banks may be seen as vulnerable to aggregated shocks as credit risks occur on the asset side while liabilities are most often unaffected. A key point by Hellwig (1994) is that the effect of macroeconomic shocks would be decreased by letting liabilities be dependent on the macroeconomic state and depositors share the burden of asset losses. Still, banks expand credit, relating to the above discussed lending booms, while knowing that the risks may lead to problems as banks cannot pass on the risk to depositors. In individual bank models, any information on the the macroeconomic state provides a signal about the quality of banks' loans to depositors. Accordingly, Allen and Gale (1998) show that macroeconomic shocks may lead to a banking crisis if depositors make their withdrawal decisions based upon leading indicators of business cycle fluctuations. Likewise, Chen (1999) illustrates in his model that adverse macroeconomic events also increase the probability of bank contagion. One may also assert the reverse when the business cycle is affected by restrictions in bank lending caused by financial fragility (Mishkin 1991).

Contagion, Spillover and Shock Propagation

A common feature of financial instabilities, in particular banking crises, is the notion of contagion. There is a rich and broad literature on the phenomenon. The theoretical literature may be distinguished into three types of contagion: (i) bank runs, and (ii) contagion through interbank lending and (iii) payment systems. This relates to two types of transmission channels. The first type of contagion can be defined to occur through the *information channel*, such as deposit withdrawals of creditors to whom the health and exposures of banks are imperfect, whereas the two latter types occur through *real channels*, such as domino effects through common exposures in inter-bank markets and payment systems.

The first type of contagion is related to *bank runs*. These events are mostly characterized by two features. First, the most prone banks and banking systems to runs to retail depositors are those not covered by deposit insurance schemes. Second, imperfectly informed investors judge the health of their own bank based upon the health of other banks. There is a wealth of literature on single banks' health based upon the balance-sheet structure and the intertemporal nature of financial contracts (as previously noted in Sect. 2.1.3), such as Bryant (1980), Diamond and Dybvig (1983) and Jagannathan (1988). The classical Diamond-Dybvig model illustrates how depositors' expectations of a bank run increase their incentives to withdraw their deposits, as late withdrawers lose all or some of their deposits. However, today's thorough deposit insurance schemes function as safety nets for this type of contagion, which might be one of the main reasons why recent waves of crisis have not, as yet,

experienced this transmission channel. Hence, it is important to distinguish between the notions of a bank run affecting one entity and a banking panic affecting multiple entities, i.e., the systemic nature of runs.

Bank run models have, accordingly, been extended to multiple banks. Chen (1999) presents an extension of the Diamond-Dybvig model, where the difference is that Chen includes two kinds of depositors: those who are informed and uninformed about the value of a bank's assets. As informed depositors are able to withdraw earlier when they comprehend that the bank cannot repay all depositors, the uninformed depositors may have an incentive to disregard their own information and respond to other sources of more noisy information (e.g., the failure of other banks). These misinterpretations may cause bank runs to become contagious. One might also reason that bank runs based upon noisier information incur higher societal costs as it might lead to defaults of healthier banks than those caused by expectations based upon correct information.

The second type of contagion focuses on the *interbank market*. This has also been a key focus of many contagion studies since the 1990s. While differences in liquidity shocks may be solved through interbank lending, the physical exposures among banks provide a channel for contagion. For instance, Rochet and Tirole (1996) show that peer monitoring, while resolving problems with moral hazard among bank shareholder managers and bank debt holders, also causes contagion risk. Along the same lines, Allen and Gale's (2000) model of interbank market exposures shows that even a small aggregate liquidity shock in a particular region can lead to systemic risk. A bankruptcy of one bank may cause other banks, which have deposits in it, to also go bankrupt. The key implication of many studies, yet not all, is that the more complete, or diversified, the markets in terms of lending relationships, the more resilient to contagion is the system.

More recent research has applied network theory to model connections between banks in the asset and liability side of the balance sheet. For instance, the findings of Babus (2006) corroborate those of Allen and Gale (2000) by considering optimal interbank network formations to reduce the risk of contagion. Leitner (2005), on the other hand, finds that the more interbank linkages a network exhibits, the better the risk sharing among banks, while the higher the potential for contagious multiple-bank failures. Conversely, emergency liquidity assistance by central banks may be motivated by surplus banks in the interbank market under-providing banks with a cash shortage, as suggested in Acharya et al. (2012). Further, already early literature has pointed out potential effects of information problems on interbank contagion. For instance, Flannery (1996) relates asymmetric information on the quality of rivals' borrowers. A shock to the financial system may hence lead to a stop in interbank lending and hoarding of liquidity, something related to the recent crisis by Cassola et al. (2008).

The third type of contagion relates to *payment systems*. The interbank lending between financial intermediaries is determined by large-value payment systems. The lending through payment systems, while not being as explicit as interbank lending, is a more detailed view of interbank exposures that may influence the propagation of shocks. From the larger family of payment systems, the main source of systemic risk derives from pure net settlement systems as netting of payments and infrequent

settlements may continue for a longer time, such that they accumulate to significant exposures [see, e.g., Freixas and Parigi (1998)]. Kahn et al. (2003) relate vulnerabilities of gross settlement systems to gridlocks and payment delays. The problems in pay-ins may be driven by high opportunity costs in foregone interest rate and doubts about other banks' solvency.

2.2.2 Empirical Findings

Next, we survey empirical works with a focus on explaining the three forms of systemic risks. The main focus lies on comparing the scope of the theoretical studies to the evidence provided by empirical studies.

Endogenous Build-up of Widespread Imbalances

The build-up phase of widespread imbalances and the relation between a financial system's pro-cyclicality and fragility is, due to numerous reasons, not an entirely straightforward question. This is illustrated by a multifaceted literature. Gourinchas et al. (2001) point to the importance of lending by showing in a large cross-country study that the likelihood of a banking crisis is higher directly after a lending boom than during tranquil periods. Likewise, findings by Dell'ariccia et al. (2012) and Mian and Sufi (2009) suggest that lending standards related to the mortgage market in the US declined prior to the ongoing financial crisis, in particular in areas with larger mortgage credit booms, house price booms and mortgage securitization rates. However, a key monetary policy tool that obviously plays a vital role in pro-cyclicality is the interest rate. Jiménez et al. (2007) and Ioannidou et al. (2009) find that reductions in interest rates often first affect positively the net present value of loans, but then with low loan rates banks attempt to re-establish profitability by moving into riskier loans. These risks, while often having somewhat long build-up episodes, may materialize suddenly and strongly either to rises in interest rates or some other unexpected trigger. Another factor leading to pro-cyclical effects is financial regulation. Repullo et al. (2010) illustrate the pro-cyclicality through capital requirements that are increasing functions of various regulatory measures of default likelihood, which often affect the the supply of credit by decreasing in good times and rising in bad times.

Another line of research has focused on the determinants of banking crises through the analysis of univariate indicators (i.e., the so-called signaling approach) and multivariate regression. In general, periods prior to systemic banking crises have been shown to be explained by traditional vulnerabilities and risks that represent imbalances like lending booms. By an analysis of univariate indicators, Alessi and Detken (2011) show that best-performing indications of boom/bust cycles are given by liquidity in general and the global private credit gap in particular. Borio and Drehmann (2009a) show that banking crises tend to be preceded by strong deviations of credit

and asset prices from their trend. Likewise, in a multivariate regression setting, vulnerabilities and risks have, overall, been shown to precede country-level crises on a large sample of developed and developing countries in Demirgüç-Kunt and Detragiache (1998) and for the US, Colombia and Mexico in Gonzalez-Hermosillo (1999), as well as on a bank level in Eastern European transition economies in Männasoo and Mayes (2009). Borio and Lowe (2002) and Borio and Lowe (2004) show that already several years prior to the current financial crisis a lending boom was awaiting behind the corner if not already visible. Lo Duca and Peltonen (2013) show that modern financial crises have been preceded by a range of macro-financial vulnerabilities and risks, particularly credit growth, equity valuations and global measures like GDP growth, real credit growth and leverage. This only provides a snapshot of the broad literature, but clearly illustrates the unanimity of imbalances preceding modern financial crises.

Exogenous Aggregate Shocks

Aggregate shocks in terms of economic downturns have commonly been shown to precede systemic banking crises. Gorton (1988) shows that a large share of banking crises in the US in the latter part of the 19th and the early part of the 20th century occurred as reactions of depositors to cyclical downturns and could hence have been correctly called with a standard model for forecasting the business cycle. While partly being related to the literature on the build-up of imbalances, systemic crises may be explained with traditional macroeconomic fundamentals (e.g., current account imbalances, gross domestic product (GDP) growth, real interest rates and inflation). Macroeconomic fundamentals have been shown to be statistically significant explanatory variables on a sample of the United States (US), Colombia and Mexico (Gonzalez-Hermosillo 1999), the Eastern European transition economies (Männasoo and Mayes 2009) and European banks during the ongoing crisis (Betz et al. 2014). These studies have, however, long forecast horizons, which relates them to imbalances and vulnerabilities prior to the crises. Yet, a number of authors show that also the timing of banking crises is related to macroeconomic fluctuations, rather than other competing factors, such as contagion. Gorton (1988) illustrates evidence for the US, Gonzalez-Hermosillo et al. (1997) for the Mexican crisis of the mid-1990s and Demirgüç-Kunt and Detragiache (1998) for a sample of developed and developing countries. Further, whereas Alfaro and Drehmann (2009) show that a large number of banking crises were preceded by decreases in GDP growth, the share that do not experience weakened GDP points at other driving factors, e.g., macroeconomic feedback effects due to the fact that GDP generally drops during post-crisis episodes.

While extreme value theory is mostly used to understand the third category of systemic risk, the study of interbank contagion, it may also be used to compute so-called tail-betas for banks. Given an extreme crash in the market, the tail-betas illustrate how the probability of crashes in individual bank stocks would be influenced. The significance of aggregate shocks in stock markets in the US (Straetmans et al. 2008)

and Europe (de Jonghe 2010) relates this to the concept of systemic risk. In European context, de Jonghe (2010) finds that banks with a large share of non-interest generating activities are more vulnerable to these aggregate shocks. Further, a comparative analysis of the shocks in the two continents is put forward by Hartmann et al. (2005). They find the effects of macro shocks on banking systems to be relevant, but similar, in the euro area and the US. Interestingly, they also show that the introduction of the euro had close to no effect on banking system risk, and relate it to the possibility that the better risk sharing and ability to absorb shocks would be offset by increases of cross-border crisis transmission channels.

Contagion, Spillover and Shock Propagation

In the early contagion literature, the main attempts were related to measuring contagious effects of bank failures on stock prices of other entities. In addition to those studies, the empirical literature on measuring interbank contagion can be matched to the three types of theoretical works: *bank runs*, *interbank lending* and *payment systems*.

Early studies have attempted to capture contagion through variation in stock prices, e.g., by measuring effects of bank failures on stock prices of other entities using event studies. Aharony and Swary (1983) and Peavy and Hempel (1988) focused on (US) banks, and their resilience to a number of failures. However, many pieces of work along this line [see, e.g., Slovin et al. (1993) and Dockinga et al. (1997)] found mixed results on contagion effects depending on the considered banks. The concept of contagion in terms of adverse stock market reactions also has been asserted as being intertwined with flight-to-quality effects, where losses of someone are benefits of others [see, e.g., Caballero and Kurlat (2008)], and to similar exposures rather than pure interbank contagion [see, e.g., Smirlock and Kaufold (1987) and Wall and Peterson (1990)]. One explanation to the mixed results might be typically observed differences in patterns during tranquil and crisis periods, where crises include non-linear and extreme stock-price movements. Hence, the more recent literature has turned the focus from regular stock price reactions to substantial ones. One potential line of research is the use of extreme value theory to estimate the spillover risk among large and complex banks [see, e.g., Hartmann et al. (2005)]. The findings of Gropp et al. (2009) illustrate that cross-border contagion risk among key European countries was significant and increased between the early 1990s and early 2000s. Yet, the focus herein is on matching the empirical works to the three groups of theoretical studies.

The first group of models based upon theoretical research aiming at capturing contagion through *bank runs* focuses on analyzing deposit flows. When there is no deposit insurance, such as during the Great Depression in the US, Saunders and Wilson (1996) have identified episodes when “bad news” about one bank caused on some occasions withdrawals from other banks (i.e., herding behavior), and on other occasions depositions in other banks (i.e., flight-to-quality effects). The results of Calomiris and Mason (1997, 2003) show equally divisive results, as they observe

contagious behavior of uninformed investors on some occasions and not on other. Allen and Gale's (2000) assertion of interbank lending explaining contagious deposit withdrawals is corroborated in a case study on an Indian bank failure in 2001 by Iyer and Peydró (2011). They show that interbank exposures to a failing bank drive retail deposit withdrawals from the exposed banks. Likewise, Van Rijckeghem and Weder (2003) test in an international context the directions of bank flows during three major financial crises. After the Mexican crisis in the mid-1990s and the Asian crisis in the end of the 1990s, the authors show that spillovers from one country to another was caused by creditor banks' exposures, whereas not during the Russian crisis in 1998.

The second group of contagion models focuses on using counterfactual simulations on balance-sheet data to assess contagion risk through the channel of *interbank lending*. The network exposures are most commonly balance-sheet linkages and the simulations often test the effects of a failure of one or several banks on the rest of the network. The simulations are, however, somewhat sensitive to underlying assumptions like the share of recovered assets from failed banks. Accordingly, the literature has presented far from unanimous results, as simulated contagion risk is negligible in Austria, Belgium, Italy and US (Elsinger et al. 2006; Furfine 2003; Mistrulli 2011), whereas the risks are larger in Germany and the Netherlands (van Lelyveld and Liedorp 2006; Upper and Worms 2004).

The third group of contagion models focuses on using simulations in large-value *payment systems* to assess interbank contagion risk. Contagion in payment systems has been explored through similar simulations. Using payment data and Monte Carlo simulations, the early literature has identified significant contagion risks in net settlement systems [see, e.g., Humphrey (1986)]. However, given appropriate risk management in payment systems (e.g., legal certainty for multilateral netting, limits on exposures, collateralization and loss sharing), some later studies have shown that interbank contagion risk may be contained. Soramäki et al. (2007) explore the network topology of the interbank payments over the Fedwire Funds Service, the payment system operated by the 12 Federal Reserve Banks of the US. Whereas they show a low average path length and connectivity for the network, as well as a tightly connected core of banks and a close to scale free degree distribution, they still point out that it is not clear how the degree distribution and other topological measures relate to contagion. Wetherilt et al. (2010) make use of a dataset of individual trades in the United Kingdom (UK) Clearing House Automated Payment System (CHAPS) to construct a network of overnight market lending. They illustrate a diversification of lending relationships that decreases their dependence on the core during the crisis, in order to attempt reducing funding liquidity risk, but make no direct conclusions about overall resilience of money market liquidity. Further, using data from the pan-European large-value payment system (i.e., the Trans-European Automated Real-time Gross Settlement Express Transfer System (TARGET)), Galos and Soramäki (2005) illustrate low systemic consequences of one bank's failure on the solvency of other banks. This indicates that today's payment systems exhibit a low risk of having systemic consequences.

2.3 Tools for Safeguarding Financial Stability

The literature, while in many aspects being in its infancy, has provided a variety of tools for safeguarding financial stability. This section focuses particularly on tools for early identification and assessment of risks. Following ECB (2010), models can be distinguished into three broad analytical approaches that match the identified forms of systemic risks:

- (i) early-warning models,
- (ii) macro stress-testing models and
- (iii) contagion and spillover models.

While the first approach aids in *risk identification*, the second and third approaches provide means for *risk assessment*. From the viewpoint of the risk cube (see Fig. 2.2), each of these aim to detect at an early stage one of the three forms of systemic risk: (i) imbalances, (ii) aggregate shocks and (iii) contagion. First, early-warning models can be used to derive probabilities of impending systemic financial crises. Second, macro stress-testing models provide a means to assess the resilience of the financial system to a wide variety of aggregate shocks. Third, contagion and spillover models can be employed to assess how resilient the financial system is to cross-sectional transmission of financial instability. In addition to models for early identification and assessment, the literature has provided a large set of coincident indicators that measure the current state of instability in the financial system. While these serve as means to measure the contemporaneous level of systemic risk, and thus may be used to identify and signal heightened stress, they are not designed to have predictive capabilities. This is not the focus of this book, but it is worth noting that *ex post* measures may serve a function in communicating the occurrence of unusual events to resolve fear and uncertainty, e.g., after the so-called flash crash of May 6, 2010 in the US (Bisias et al. 2012). In the sequel of this section, we focus on the three analytical approaches to derive tools for early identification and assessment of risks. In line with the focus of this book, the final subsection summarizes advances in visualization approaches in both risk identification and risk assessment.

2.3.1 Early-Warning Indicators and Models

Early-warning exercises may be performed with a wide range of methods and indicators, which are also known in the literature as Early Warning Systems. The main aim of these tools is to predict vulnerable states prior to financial instabilities and crises. Hence, they oftentimes first define an index of financial instability or stress in an entity, e.g., country, bank or market. The contemporaneous level of systemic risk may, for instance, be derived from coincident stress indices, such as the Composite

Indicator of Systemic Stress (CISS) by Holló et al. (2012).¹ A threshold, or some combination with other rules, on the index value defines binary crisis/tranquil events for the entities. For the models to focus on imbalances, risks and vulnerabilities, a binary pre-crisis variable is then set to 1 during some specific horizon prior to the crisis events, and to 0 in all other periods. The other part of data used is a set of vulnerability and risk indicators. These are chosen and transformed according to their performance in explaining and predicting the binary pre-crisis variable. The outputs of such models mostly take the form of a probability of a crisis within a specific time horizon and are monitored with respect to threshold values (or cut-off values).

The early univariate signaling literature used country-specific percentile transformations of single indicators and turned them into signals by choosing an optimal threshold. The optimal threshold is commonly chosen based upon specified weights on the loss of type I and II errors (see Sect. 7.2 for an overview of evaluation frameworks and the one used in this book). Kaminsky and Reinhart (1996) and Kaminsky et al. (1998) introduced the signaling approach for predicting currency crises. Lately, it has been applied to boom/bust cycles (Alessi and Detken 2011), banking system crises (Borio and Drehmann 2009a), and to sovereign debt default (Knedlik and Schweinitz 2012). However, the key limitation of this approach is that it does not enable any interaction between or weighting of indicators, while an advantage is that it demonstrates a more direct measure of the importance and provides a ranking of each indicator.

Much of the early-warning literature deals, however, with models that rely on conventional statistical methods, such as logit/probit models. Logit or probit regressions use on the left-hand side the binary pre-crisis variable and on the right-hand side the early-warning indicators. The linear regression models make use of a cumulative probability function to force the value of the predicted variable within the interval $[0, 1]$. This estimation provides a direct aggregate measure of the intensity of the signal, i.e., the probability of an impending crisis. The two models are similar, except that the probit model uses the cumulative normal distribution and the logit the cumulative logistic function to transform variables into the $[0, 1]$ interval. Logit and probit models have frequently been applied to predicting financial crises. Eichengreen and Rose (1998), Frankel and Rose (1996) and Sachs et al. (1996) provide some early applications of probit/logit analysis to currency crisis prediction. Later, Berg and Pattillo (1999) apply a probit model to predicting currency crises; Schmidt (1984) and Fuertes and Kalotychou (2006) to predicting debt crises; Barrell et al. (2010) to predicting banking crises; and Lo Duca and Peltonen (2013) to predicting systemic crises. For an early, yet comprehensive, review, see Berg et al. (2005).

In comparison to the signals approach, binary-choice methods allow for a multivariate approach to estimating crisis probabilities, while providing means to rank

¹ There are many coincident stress indices. For instance, Illing and Liu (2006) focus on measuring financial stress in Canada and Hakkio and Keeton (2009) discuss more broadly what financial stress is, how it can be measured and why it matters. Cardarelli et al. (2011) and Balakrishnan et al. (2009) construct financial stability indices for a broad set of advanced and emerging economies, whereas the CISS aims at measuring stress in the euro area.

risks and assess most significant indicators, but still depend largely on a number of restrictive assumptions. While being non-linear in nature, the relationship between indicators and the events is still assumed to consistently follow some specific function (e.g., logistic or normal). Further, the lack of interactions between indicators may also limit performance as indicators of debt, currency, and systemic crises have been shown to be non-linearly related (Fioramanti 2008; Lo Duca and Peltonen 2013; Arciniegas Rueda and Arciniegas 2009). While interaction terms can be included in logit/probit specifications, manually specifying the complex relations between and interactions among various economic and financial factors is a demanding task. This should be accounted for when choosing a predictive method.

A new approach to early-warning modeling has been the introduction of methods commonly used in subfields of computer science, such as data mining, machine learning and pattern recognition. Since the turn of last century, the use of such intelligent, oftentimes also distribution-free and non-parametric, techniques in crisis monitoring have increased. Indeed, the flexible non-parametric techniques have slightly improved results in *ex post* crisis prediction [see Demyanyk and Hasan (2010) for a review]. The key methods in non-parametric early-warning models have so far been based upon biologically inspired computing in general and artificial neural networks (ANNs) in particular (Nag and Mitra 1999; Franck and Schmied 2003; Peltonen 2006; Fioramanti 2008). The first to publicly try predicting financial crises with the help of an ANN were Nag and Mitra (1999). Their findings on predicting the Malaysian, the Thai and the Indonesian currency crises suggested that their ANN approach performed better than the signaling approach. Similarly, Franck and Schmied (2003) also concluded that their application of an ANN for predicting the speculative attacks in Russia in 1998 and Brazil in 1999 outperformed a logit model. Peltonen (2006) used an ANN to predict the Asian currency crisis and showed that it outperforms a probit model. Fioramanti (2008) shows in his study that a non-parametric ANN-based early-warning model outperforms analyses using the signals approach and probit or logit models. Yet, when the focus is on the introduction of one specific method, it is important to note that mostly “successful” experiments are reported.

A task that remains to be unexplored is the choice of indicators in the models. A large number of studies use univariate predictive performance in terms of the signaling approach to assess the extent of discriminatory power of individual indicators [e.g., Kaminsky et al. (1998), Alessi and Detken (2011) and Lo Duca and Peltonen (2013)], of which Lo Duca and Peltonen (2013) use the best predictors as an input to a logit regression. Still, due to the possibly complex interactions, the choice of indicators should be performed in a multivariate setting.

2.3.2 Macro Stress-Testing Models

The key family of tools for assessing risks of exogenous aggregate shocks is that of macro stress-testing models. Hence, while the above discussed tools aim at

risk identification, the tools for risk assessment are literally of different nature. Stress-testing models allow policymakers to assess the consequences of assumed extreme, but plausible, shocks for different entities. As stress-testing is no new concept, there is a broad literature not only on micro stress-testing, but also on the macro level. While being macro stress-tests, they commonly follow many principles used in micro stress-testing and risk management to assess the loss potential of specific portfolios given extreme market conditions [see, e.g., McNeil et al. (2005)]. Kida (2008) pinpoints the differences between micro and macro stress-testing to three key factors. First, macro models commonly include multiple banks with different portfolios, where differences affect how resilient one bank is to shocks and how shocks to one bank affects the system. Second, macro models oftentimes include multiple time points by enabling shocks to propagate for several periods. Third, a macro stress-test focuses on how risk is propagated between banks or between sectors. The recent handbook edited by Quagliariello (2009) provides a comprehensive overview of the macro versions of such models.

The key question of macro stress-testing, or stress-testing in general, is finding the balance between plausibility and severity of the stress scenarios such that they are plausible enough to be taken seriously and severe enough to be meaningful [see, e.g., Alfaro and Drehmann (2009) and Quagliariello (2009)]. Then, the assessment of shocks most often includes also the propagation of the shock among entities. Kida (2008) pinpoints the feedback (or risk transmission and propagation) mechanisms into four key types: (i) interbank contagion (e.g., when exposures to risk spread through the interbank loans market), (ii) correlation between credit and market risks (e.g., when increases in interest rates raise the probability of default of borrowers of a bank, and causes thus also increases in interest rates), (iii) correlation between asset prices and the portfolio adjustment mechanisms of a bank (e.g., when increases in asset prices damage banks' balance sheets, leading to large-scale sales of assets, and thus further decreasing asset prices), (iv) propagation of shocks between the financial system and the real economy (e.g., when banking system shocks affect economic activity, and thus further weaken banks' credit environment). Contrary to the early-warning model literature, stress-testing does not attempt to derive the likelihood and severity of shocks, but rather takes that as given. This information could, obviously, come from an early-warning indicator or model. In a macro setting, a policymaker is more interested in the resilience of the financial system more broadly, or the banking system in particular. Policymakers may hence test various adverse scenarios and design policy actions related to individual institutions or the general architecture if the resilience of the system is judged not to be strong enough.

A macro stress-testing approach to assessing a banking system uses multiple inputs and consists of a number of different steps. First, most often a basis for the test is a scenario of an adverse aggregate macroeconomic or macro-financial shock. This shock may be defined on hypothetical grounds or estimated from data, such as a tail density forecast of a macroeconomic model. The second step uses a set of exposures and other mechanisms to link banks to the impact of the adverse scenario. The links may be banks' loan books or other credit risk exposures of a bank or a country-level banking system. Thus, the effects of the scenario are shown as changes

in the probabilities of default and losses given default, and also lead to indications of whether and how many banks fail [see, e.g., Castrén et al. (2009)]. Likewise, Castrén et al. (2010) estimate a so-called global vector autoregressive model and link it to firms' default probabilities for a model that may be used for analyzing a financial sector's probability of default given a range of macro shocks. Alfaro and Drehmann (2009) use country-specific univariate autoregressive models to forecast GDP growth, but focus more on showing that stress scenarios derived from historical data are not severe enough in comparison to actual events. Hirtle et al. (2009) describe the stress-testing model of the Supervisory Capital Assessment Program, which tests a range of macroeconomic scenarios, e.g., variation in GDP growth, housing prices and unemployment. For comprehensive reviews of the stress-testing literature, see Sorge (2004) and Drehmann (2009).

2.3.3 Contagion and Spillover Models

The main aim of contagion models is to assess the transmission of financial instabilities in the cross section. Hence, they attempt to answer the question: *With what likelihood, and to what extent, could the failure of one or multiple financial intermediaries cause the failure of other intermediaries?* Further, they may also focus on the failure of one or several financial markets and their likelihood to cause failures of other markets. Thus, contagion and spillover models attempt to grasp, show and quantify the transmission channels of instability across financial intermediaries and markets, as well as market infrastructures [(for comprehensive reviews, see Bandt et al. (2009) and Upper (2007)]. Herein, we discuss the use of three data sources to answer these questions: (i) market-based data, (ii) interbank balance-sheet data, and (iii) interbank payments data.

First, one can use *market-based estimates* to measure the extreme dependence of negative asset returns, the so-called tail dependence. The first approach measures the extent of losses, after controlling for common factors, caused by a large loss of market value or a large increase in default probability. These approaches commonly identify tail-risk drivers in a tail-dependence network and enable assessing which entities are particularly vulnerable to large losses in the market. For instance, IMF (2009) presents a co-risk model for assessing interdependence among banks under extreme events and a distress dependence matrix for assessing pairs of banks' distress probabilities, both using market data. Likewise, Hautsch et al. (2011) propose the systemic risk beta as a measure for financial companies' contribution to systemic risk given network interdependence between firms' tail risk exposures measured using equity prices. While being widely available and capturing other contagion channels than those in direct linkages between banks (Acharya et al. 2010), market price data assume that asset prices correctly reflect all publicly available information on bank exposures. Yet, it has repeatedly been shown that securities markets are not always efficient in reflecting information about stocks and are thus vulnerable to mispricing distortions [see, e.g., Malkiel (2003)]. In addition, market prices are most often

contemporaneous, rather than leading indicators, and it might be difficult to separate the factors driving market prices in order to observe bilateral interdependence (Borio and Drehmann 2009b).

The second approach, on the other hand, uses counterfactual simulations on *balance-sheet data*, or some proxy of them. These studies simply simulate to what extent and whether the failure of one financial intermediary would lead to losses of other intermediaries. For instance, Castrén and Kavonius (2009) provide a tool for assessing contagion and the transmission of risk in the euro area financial system. They construct a sector-level network of bilateral balance sheet exposures of the euro area financial accounts data, as well as include sensitivity of the balance sheets to changes in leverage and asset volatility, to illustrate the propagation of local shocks in the network. Likewise, Chan-Lau (2010) evaluates, under extreme adverse scenarios, interconnectedness risk in banking systems among mature and emerging market economies, and between individual financial institutions in Chile, using balance sheet-based network analysis. Along these lines, Battiston et al. (2012) developed a network measure of centrality, the DebtRank, as one approach to capture the impact of distress in a financial institution to the cross section across the entire network. Moreover, the IMF (2009) presents a default-intensity model that uses both direct and indirect linkages in the financial sector, as well as combines them with failure probabilities of banks, to achieve a measure of the probability of failure of a large fraction of financial institutions. Yet, balance-sheet data, while measuring direct linkages between banks, are mostly not publicly disclosed. In many cases, even supervisors and other market oversight authorities have access to only partial information.

The literature on the third group of models focusing on *payments data* is somewhat scarce. A concern once again is that interbank payment data, likewise interbank lending data, are locked behind the doors of confidentiality. Yet, while not always making use of real data, there exist some tools based upon payments data. In particular, the three compilations edited by Leinonen (2005, 2007, 2009) provide a broad overview of policy-oriented research on tools for payment systems simulation. Recently, along the lines of DebtRank for balance-sheet data, Soramäki and Cook (2010) developed a network metric for payments data, the SinkRank, for identifying systemically important banks and most affected banks in the case of distress.

2.3.4 Tools With Visual Capabilities

Data visualization can serve multiple purposes in macroprudential oversight. First, visual representations can generally be classified to be used to enhance communication with two audiences: (i) internal and (ii) external. The purpose of use in *internal* communication relates to enhancing the understanding of policymakers on various levels. One task is obviously to support the analysts themselves, and within other groups of active participants in the process of deriving analytical models. Further, one may also want to communicate to the outside of the involved counterparties,

which involves making use of visuals when presenting to the management, entire divisions and even on the level of the institution or organization as such. The key task, at the lower level, is to provide means for interaction with visuals in order to amplify cognition, that is, to better understand and model the task at hand (for further discussion see Sect. 4.1), whereas the higher level focuses more on reporting and presentation of information by the means of oftentimes static visuals. While the case of low-level analysts can easily be imagined, an example at a higher level could be the dissemination of identified risks by the risk identification division for assessment at the risk assessment division. *External* communication, on the other hand, refers to conveying information to other authorities with responsibility for financial stability and overall financial-market participants, such as laymen, professional investors and financial intermediaries. Whereas this mainly relates to communication of readily processed and finalized data products, such as on the high level of internal communication, it obviously is a more challenging task due to the large heterogeneity in the audience. A direct example of such communication is quarterly or biannual Financial Stability Reports, a recent phenomenon that has quickly spread to a large number of central banks.

In the context of low-level internal communication of systemic risk modeling, Flood and Mendelowitz (2013) note that data exploration is an area where visualization tools can make a major contribution. They point to the fact that certain tasks of classification, analysis and triage can be automated, whereas many require a human analyst, such as the difficulty to train a well-performing machine to analyze anomalous financial market activity. This follows the very definition of visual analytics (see Sect. 4.1.3). Ekholm (2012)—the Deputy Governor of Sveriges Riksbank, the first central bank to publish a stability report in 1997—notes that there is a strive for not only openness and transparency, but also clear external communication, in particular during times of crisis when “a *“negative” but reliable announcement can [...] be better for confidence than a “positive” but uncertain announcement*”.

Herein, we discuss a brief overview of used visualization tools for the above categories of models: (i) early-warning models, (ii) macro stress-testing models, and (iii) contagion and spillover models.

First, the standard predictive *early-warning models* may be complemented by the use of tools amplifying cognition. Due to the complexity of financial systems, a large number of indicators are often required to accurately assess the sources of financial instability. As with statistical tables, standard two- and three-dimensional visualizations have, of course, their limitations for high dimensions, not to mention the challenge of including a temporal or cross-sectional dimension or assessing multiple countries over time. Although composite indices of leading indicators and predicted probabilities of early-warning models enable comparison across countries and over time, these indices fall short in describing the numerous sources of distress.

Some recent approaches make use of techniques for multidimensional visualization to assess sources of risk and vulnerability. Work by International Monetary Fund (IMF) staff on the Global Financial Stability Map (GFSM) (Dattels et al. 2010) has sought to disentangle the sources of risks by a mapping of six composite indices with a standard radar-chart visualization. Even here, however, the GFSM falls short

in disentangling individual sources, for which separate visualizations are needed. In addition, familiar limitations of radar charts are, for example, the facts that area does not scale one-to-one with increases in variables and that the area itself depends on the order of dimensions. This is illustrated in Fig. 2.3, where Country A and Country B have an area of significantly (i.e., infinitely) different size but the same aggregate risks (i.e., mean value). In addition, the use of adjustment based on market and domain intelligence, especially during crisis episodes, and the absence of a systematic evaluation gives neither a transparent data-driven measure of financial stress nor an objective anticipation of the GFSM's future precision. Indeed, the GFSM comes with the following caveat: “*given the degree of ambiguity and arbitrariness of this exercise the results should be viewed merely illustrative*”.²

Data and dimension reduction methods have also been used to represent these complex data. In terms of Fuzzy c-means (FCM) clustering, a combination of clustering models and the reasoning of fuzzy logic have been introduced to the early-warning literature by finding risky clusters and treating relationships in data structures as true or false to a certain degree (Marghescu et al. 2010). This type of analysis has the benefit of not only signaling a crisis in a timely manner, but also signaling the type and degree of various sorts of financial imbalances. In an exploratory study, Arciniegas Rueda and Arciniegas (2009) found, with the help of the Self-Organizing Map (SOM), strong associations between speculative attacks' real effects and 28 indicators, yet did neither focus on visualizing individual data nor on early-warning performance. Resta (2009) also has applied the SOM to a large set of indicators, but with a focus on rather general economic and financial performance of countries and with limited evaluations of classification performance.

Second, *macro stress-testing models*, to the best of my knowledge, make no use of advanced visualization techniques for representing the results of the tests, including the processing of data at the input, interim and output stage. The visualizations seldom go beyond a framework or schematic structure for the designed transmission mechanisms in the model and plots of loss distributions in various formats. Obviously, standard visualizations from graph theory may be used in representing networks, if such are used in the models. For instance, the macro stress-testing model by Boss et al. (2006), which integrates satellite models of credit and market risk with a network model for evaluating default probabilities of banks, enable one to make use of concepts from graph theory in visualizing the network structure. Network visualizations are, however, more common in contagion models.

As said, the third group of *contagion and spillover models* commonly make use of concepts from graph or network theory to visualize the structure of linkages in

² The authors state that the definitions of starting and ending dates of the assessed crisis episodes are somewhat arbitrary. Similarly, the assessed crisis episodes are arbitrary, as some episodes in between the assessed ones are disregarded, such as Russia's default in 1999 and the collapse of Long-Term Capital Management. Introduction of judgment based upon market intelligence and technical adjustments are motivated when the GFSM is “*unable to fully account for extreme events surpassing historical experience*”, which is indeed an obstacle for empirical models, but also a factor of uncertainty in terms of future performance since nothing assures manual detection of vulnerabilities, risks and triggers.




Fig. 2.3 Radar charts of two countries. *Notes* The figure provides an example of a radar chart, such as the one in Dattels et al. (2010)

the models [see, e.g., Estrada (2011)]. This provides means to represent entities as nodes (or vertices) and their links as edges (or arcs). The combination of nodes and edges provide all constituents for a network, where the edges may be directed *versus* undirected and weighted *versus* unweighted. However, rather than a visualization, a network is a data structure. The interpretability of networks has been enhanced by the means of various methods. For instance, positioning algorithms, such as force-directed layout methods, are commonly used for locating nodes with similar edges close to each other, as well as ring and chord layouts for more standardized positioning. Yet, the so-called hairball visualization, where nodes and edges are so large in number that they challenge the resolution of computer displays, not to mention interpretation, is not a rare representation of complex financial networks [see, e.g., Bech and Atalay (2010)]. Still, it is worth noting that recent advances in software for visualizing financial networks, such as Financial Network Analytics (www.fna.fi), hold promise in bringing aesthetics and the ease of use to visualizations in the financial domain. An additional essential feature, not the least to deal with hairballs, is the use of interaction techniques with visualizations.

2.4 A Framework for Macroprudential Oversight

To connect the concepts defined in this chapter, we discuss them in how they relate to safeguarding financial stability. One might thus also say that this section attempts to provide a holistic view of the macroprudential oversight process. The ECB's conceptual framework not only includes a systematic way of structuring risks through the risk cube, but also includes a process of the steps that a macroprudential super-

visory body would follow.³ The process in Fig. 2.4 is an adapted version of that in ECB (2010), where red components represent risks and vulnerabilities, the green components represent the need for risk identification and assessment, and the blue components represent the need for risk communication. Thus, the **black frames** mark the use of tools for safeguarding financial stability, where the **solid** lines represent the current approach and the **dashed** lines proposes an integration of means for risk communication into the tools. A discussion of policy assessments and implementations represented by gray components is beyond the scope of this book.

The macroprudential oversight process begins with underlying market imperfections that at a later stage propagate as possible risks. In the *first step* of the supervisory process (risk identification), the key focus is on identifying risks to stability and potential sources of vulnerability. The vulnerabilities and risks could exist in any of the three components of the financial system: financial intermediaries, financial markets and the financial infrastructure. The necessary tools to identify possible risks, vulnerabilities and triggers come from the set of early-warning models and indicators, as well as the use of market intelligence, and expert judgment and experience. This provides means for ranking risks and vulnerabilities as per intensity, as well as for assigning probabilities to specific shocks or future systemic events.

In the *second step* of the process (risk assessment), the rankings and probabilities may be used to assess the identified risks. The used tools come mainly from the set of macro stress-testing models and contagion models. In macro stress-testing, simulations of most plausible risk scenarios show the degree of impact severity on the general financial system, as well as its components. The contagion models, on the other hand, might be used through counterfactual simulations to assess the impact of specific failures on the entire financial system and individual institutions. The first and the second step of the process should not only provide a list of risks ordered according to possible severity, but also contain their materialization probabilities, losses given their materialization, and losses in macroeconomic output and welfare, as well as their possible systemic impact. Hence, these two initial steps in the process aim at early risk identification and assessment and provide means for safeguarding financial stability.

The *third step* (policy assessment) involves the assessment of policy actions as early preventive measures. Based upon the identified and assessed risks, a macroprudential supervisory body can consider giving a wide variety of risk warnings and recommendations for other parties to use policy instruments, as well as an implementation of policies given the instruments at hand. To steer their decisions, the policy assessment step can make use of the same analytical tools used for risk identification and assessment. While policy tools and their effectiveness is slightly outside macroprudential oversight and the general scope of this book, it is worth noting that actions tailored to the needs of a system-wide orientation are a key part of macroprudential

³ A macroprudential supervisory body is an institution tasked with macroprudential oversight of the financial system and the mandate of safeguarding financial stability. Examples are the European Systemic Risk Board in Europe, the Financial Policy Committee in the UK, and the Financial Stability Oversight Council in the US.

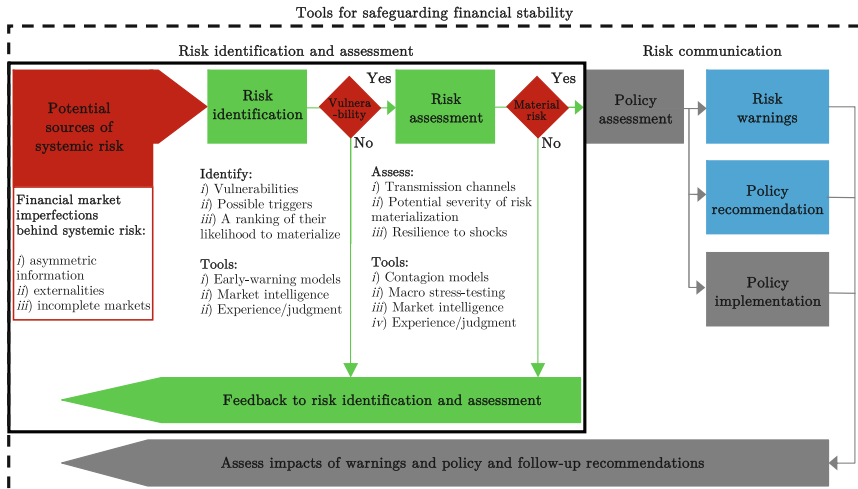


Fig. 2.4 The macroprudential oversight process. *Notes* The figure represents the role of analytical models and tools for identification and assessment of systemic risk in the macroprudential oversight process. The *red components* represent risks and vulnerabilities, the *green components* represent the need for risk identification and assessment, and the *blue components* represent the need for visual means facilitating risk communication. Thus, the *black frames* mark the need for tools for safeguarding financial stability, where the *solid lines* represent the current approach and the *dashed lines* represent means for risk communication integrated in the tools. The *gray components* are beyond the scope of this book. The figure is an adapted version of that in ECB (2010)

regulation and supervision. As interest rate policy may be a too blunt and powerful tool with material damage to other parts of the economy, the policies could take the form of tighter standards—e.g., requirements on capital adequacy, provisioning, leverage ratios, and liquidity management—for individual financial institutions with larger contributions to systemic risk and calibrated to address common exposures and joint failures. Macroprudential regulation and tools may also be used for accumulating buffers or reserves in good economic times to be used during worse times.

Performing risk identification and assessment is generally seen as the key task of tools for safeguarding financial stability (solid **black** frame in Fig. 2.4). This points to a lack of integration between the tools for safeguarding financial stability and the communication that occurs after the policy assessment step, in particular the tasks of issuing risk warnings, giving policy recommendations and publishing Financial Stability Reports, as represented by the blue components in Fig. 2.4. To answer the question, *what is the overall purpose of communication through a Financial Stability Report?*, a survey among central bankers by Oosterloo and Haan (2004) pinpoints three main reasons for publishing these reports:

- (i) to contribute to overall financial stability,
- (ii) to increase the transparency and accountability, and
- (iii) to strengthen co-operation between authorities with financial stability tasks.

Thus, following the discussion in the previous section, a major concern is how the results of these risk identification and assessment tools are communicated to a wide range of stakeholders in easily understandable formats, with the ultimate aim of achieving transparency and accountability. The broader perspective proposed by the dashed **black** frame in Fig. 2.4 argues for relating the third step to risk communication, which would be supported by visual representations of the tools used in the prior steps. Although not being illustrated in the figure, internal and external risk communication would obviously have separate feedback loops: the former to risk identification and assessment (green components), and the latter to potential sources of systemic risk, vulnerabilities, and material risks (red components). This would translate to a threefold focus of tools: risk identification, risk assessment and risk communication.

2.5 Concluding Discussion

This chapter has provided an overview of macroprudential oversight. Not only have we discussed how financial systems work and what makes them fragile, but also the specific systemic risks and tools for safeguarding financial stability. Finally, the chapter ends by summarizing all the above ingredients within a larger framework of the macroprudential oversight process.

Macroprudential oversight as such is not a new concept. Yet, supervisory bodies with the mandate of safeguarding system-wide financial stability have only recently been created, all in the aftermath of the financial instabilities of 2007–2008. The European Systemic Risk Board in Europe, the Financial Policy Committee in the UK, and the Financial Stability Oversight Council in the US were all either established or announced in 2010. While we have discussed the complexity of factors affecting financial systems, how fragilities may build up and what form systemic risks may take, as well as empirical and theoretical underpinnings, an obvious focus of this chapter is on tools and models for macroprudential oversight. Given the mandate of multiple macroprudential supervisory bodies, the central task ought to be timely and accurate measurement of systemic risks. In this chapter, we have discussed the following three categories of systemic risks (*and tools*):

- (i) endogenous build-up of widespread imbalances (*early-warning models*);
- (ii) exogenous aggregate shocks (*macro stress-testing models*); and
- (iii) contagion and spillover (*contagion and spillover models*).

This sets an inherent need for a broad basis of tools for the identification and assessment of potential risks, vulnerabilities and imbalances. One key conclusion of the review of tools and models is the lack of visual means for identifying and assessing risks and vulnerabilities, particularly macro stress-test and early-warning models. In the case of contagion models, visualizations based upon network models and graph theory have been applied and are still gaining further interest within the policymaking community. Yet, the task of representing high-dimensional early-warning indicators

on a low-dimensional display has not been addressed in a sufficient manner. Visual aids to the representation of macro stress-test models may also hold promise due to their complex nature, but to provide a sufficient abstraction of the problem seems like an inherently different, yet highly interesting, task to address. However, this is generally beyond of the scope of this book.

Another line of research is to purely focus on the forecasting capabilities of models. The early-warning literature has indicated that ANNs are suitable for the complex task. They are effective data-driven non-linear function approximators, but are alas no *panacea* for binary-choice classification. To fully benefit from capabilities of ANNs, they need to be provided with their computational demands (i.e., large samples and computing power) and specific training schemes for generalization. In addition, the literature showed that the choice of the optimal set of indicators is either performed according to economic significance or univariate predictive performance, whereas the choice has not been performed in a multivariate framework.

Yet, in all above tasks, it is worth remembering that the quality of a model is highly dependent on the quality of the underlying data. The early-warning models are generally dependent upon country-level macroeconomic, banking system and market-based indicators of risks, vulnerabilities and imbalances. This takes us to the topic of data in macroprudential oversight.

References

- Acharya, V., Pedersen, L., Philippon, T., & Richardson, M. (2010). *Measuring systemic risk*. Discussion Paper, NYU, Stern School of Business.
- Acharya, V., Gromb, D., & Yorulmazer, T. (2012). Imperfect competition in the interbank market for liquidity as a rationale for central banking. *American Economic Journal: Macroeconomics*, 4(2), 184–217.
- Aharony, J., & Swary, I. (1983). Contagion effects of bank failures: evidence from capital markets. *Journal of Business*, 56(3), 305–322.
- Alessi, L., & Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: a role for global liquidity. *European Journal of Political Economy*, 27(3), 520–533.
- Alfaro, R., & Drehmann, M., (2009). Macro stress tests and crises: What can we learn? [BIS Quarterly Review], (December).
- Allen, F., & Gale, D. (1998). Optimal financial crises. *The Journal of Finance*, 53, 1245–1284.
- Allen, F., & Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1), 1–33.
- Allen, W. A., & Wood, G. (2006). Defining and achieving financial stability. *Journal of Financial Stability*, 2(2), 152–172.
- Allen, F., & Gale, D. (2007). *Understanding financial crises*. Oxford: Oxford University Press.
- Arciniegas Rueda, I., & Arciniegas, F. (2009). Som-based data analysis of speculative attacks' real effects. *Intelligent Data Analysis*, 13(2), 261–300.
- Babus, A. (2006). *The formation of financial networks*. Tinbergen Institute Discussion Papers, 06–093/2.
- Balakrishnan, R., Danninger, S., Elekdag, S., & Tytell, I. (2009). *The transmission of financial stress from advanced to emerging economies*. IMF Working Paper, No. 09/133.
- de Bandt, O., & Hartmann, P. (2002). Systemic risk in banking: A survey. In C. Goodhart & G. Illing (Eds.), *Financial crisis, contagion and the lender of last resort: A book of readings*. Oxford: Oxford University Press.

- Banerjee, A. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797–817.
- Barrell, R., Davis, P., Karim, D., & Liadze, I. (2010). Bank regulation, property prices and early warning systems for banking crises in oecd countries. *Journal of Banking and Finance*, 34(9), 2255–2264.
- Battiston, S., Puliga, M., Kaushik, R., Tasca, P., & Caldarelli, G. (2012). Debrank: too central to fail? financial networks, the fed and systemic risk. *Nature*, 2(541), 1–6.
- Bech, M., & Atalay, E. (2010). The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, 389(22), 5223–5246.
- Berg, A., & Pattillo, C. (1999). Predicting currency crises—the indicators approach and an alternative. *Journal of International Money and Finance*, 18(4), 561–586.
- Berg, A., Borensztein, E., & Pattillo, C. (2005). Assessing early warning systems: How have they worked in practice? *IMF Staff Papers*, 52(3), 462–502.
- Betz, F., Oprica, S., Peltonen, T., & Sarlin, P. (2014). Predicting distress in European banks. *Journal of Banking and Finance* (Forthcoming).
- BIS. (1986). *Recent innovations in international banking, Report prepared by a study group established by the central banks of the group of ten countries*, Basel.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy*, 100(5), 992–1026.
- Bisias, D., Flood, M., Lo, A., & Valavanis, S. (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics*, 4, 255–296.
- Boot, A., & Greenbaum, S. (1993). Bank-regulation, reputation and rents: Theory and policy implications. In C. Mayer & X. Vives (Eds.), *Capital markets and financial intermediation*. Cambridge: Cambridge University Press.
- Borio, C. (2011). Implementing a macroprudential framework: Blending boldness and realism. *Capitalism and Society*, 6(1), 1–23.
- Borio, C., & Drehmann, M. (2009a). Assessing the risk of banking crises—revisited. [BIS Quarterly Review], (March).
- Borio, C., & Drehmann, M. (2009b). *Towards an operational framework for financial stability: 'fuzzy' measurement and its consequences*. BIS Working Papers No. 284.
- Borio, C., & Lowe, P. (2002). *Asset prices, financial and monetary stability: Exploring the nexus*. BIS Working Papers No. 114.
- Borio, C., & Lowe, P. (2004). *Securing sustainable price stability: Should credit come back from the wilderness?* BIS Working Papers No. 157.
- Boss, M., Krenn, G., Pühr, C., & Summer, M. (2006). Systemic risk monitor: A model for systemic risk analysis and stress testing of banking systems. *Oesterreichische Nationalbank Financial Stability Report*, 11(June), 83–95.
- Bryant, J. (1980). A model of reserves, bank runs, and deposit insurance. *Journal of Banking and Finance*, 4(4), 335–344.
- Caballero, R., & Kurlat, P. (2008). *Flight to quality and bailouts: Policy remarks and a literature review*. MIT Working Paper, 08–21.
- Calomiris, C., & Mason, J. (1997). Contagion and bank failures during the great depression: the June 1932 Chicago banking panic. *American Economic Review*, 87(5), 863–883.
- Calomiris, C., & Mason, J. (2003). Fundamentals, panics, and bank distress during the depression. *American Economic Review*, 93(5), 1615–1647.
- Cardarelli, R., Elekdag, S., & Lall, S. (2011). Financial stress and economic contractions. *Journal of Financial Stability*, 7(2), 78–97.
- Carletti, E. (2008). Competition and regulation in banking. In A. Boot & A. Thakor (Eds.), *Handbook in financial intermediation* (pp. 449–482). North Holland: Elsevier.
- Cassola, N., Drehmann, M., Hartmann, P., Lo Duca, M., & Scheicher, M. (2008). A research perspective on the propagation of the credit market turmoil. *European Central Bank Research Bulletin*, 7.

- Castrén, O., & Kavonius, I. (2009). *Balance sheet interlinkages and macro-financial risk analysis in the euro area*. ECB Working Paper No. 1124.
- Castrén, O., Fitzpatrick, T., & Sydow, M. (2009). *Assessing portfolio credit risk changes in a sample of EU large and complex banking groups in reaction to macroeconomic shocks*. ECB Working Paper No. 1002.
- Castrén, O., Dées, S., & Zaher, F. (2010). Stress testing euro area corporate default probabilities using a global macroeconomic model. *Journal of Financial Stability*, 6(2), 64–78.
- Chan-Lau, J. (2010). *Balance sheet network analysis of too-connected-to-fail risk in global and domestic banking systems*. IMF Working Paper 10/107.
- Chen, Y. (1999). Banking panics: the role of the first-come, first served rule and rational expectations equilibrium. *Journal of Political Economy*, 107(5), 946–968.
- Dattels, P., McCaughrin, R., Miyajima, K., & Puig, J. (2010). *Can you map global financial stability?* IMF Working Paper No. 10/145.
- de Bandt, O., Hartmann, P., & Peydro, J. (2009). Systemic risk in banking: An update. In A. Berger & M. P. Wilson, J. (Eds.), *Oxford handbook of banking*. Oxford: Oxford University Press.
- Dell'ariccia, G., & Marquez, R. (2006). Lending booms and lending standards. *The Journal of Finance*, 61(5), 2511–2546.
- Dell'ariccia, G., Igan, D., & Laeven, L. (2012). Credit booms and lending standards: evidence from the subprime mortgage market. *Journal of Money, Credit and Banking*, 44(2–3), 367–384.
- Demirgüç-Kunt, A., & Detragiache, E. (1998). The determinants of banking crises in developing and developed countries. *IMF Staff Papers*, 45(1), 81–109.
- Demyanyk, Y., & Hasan, I. (2010). Financial crises and bank failures: a review of prediction methods. *Omega*, 38(5), 315–324.
- Diamond, D., & Dybvig, P. (1983). Bank runs, deposit insurance and liquidity. *Journal of Political Economy*, 91(3), 401–419.
- Dockinga, D., Hirschey, M., & Jones, E. (1997). Information and contagion effects of bank loan-loss reserve announcements. *Journal of Financial Economics*, 43(2), 219–239.
- Drehmann, M. (2009). Macroeconomic stress testing banks: A survey of methodologies. In M. Quagliariello (Ed.), *Stress testing the banking system: Methodologies and applications*. Cambridge: Cambridge University Press.
- ECB. (2005). Assessing financial stability: Conceptual boundaries and challenges. [Financial Stability Review]. European Central Bank: Frankfurt (June, 2005).
- ECB. (2009). The concept of systemic risk. [Financial Stability Review]. Frankfurt: European Central Bank. (December, 2009).
- ECB. (2010). Analytical models and tools for the identification and assessment of systemic risks. [Financial Stability Review]. Frankfurt: European Central Bank. (June, 2010).
- ECB. (2012). *Identifying emerging risks in the euro area—the work at the ECB*, Paper Presented at the Seminar on Macprudential Tools and Frameworks, 29–31 May.
- Eichengreen, B., & Rose, A. K. (1998). *Staying afloat when the wind shifts: External factors and emerging-market banking crises*. NBER Working Paper No. 6370.
- Ekholm, K. (2012). *Macprudential policy and clear communication contribute to financial stability* Speech to the Swedish Banker's Association, 30 March.
- Elsinger, H., Lehar, A., & Summer, M. (2006). Risk assessment for banking systems. *Management Science*, 52(9), 1301–1314.
- Estrada, E. (2011). *The structure of complex networks: Theory and applications*. Oxford: Oxford University Press.
- Fell, J., & Schinasi, G. (2005). Assessing financial stability: exploring the boundaries of analysis. *National Institute Economic Review*, 192, 102–117.
- Fioramanti, M. (2008). Predicting sovereign debt crises using artificial neural networks: a comparative approach. *Journal of Financial Stability*, 4(2), 149–164.
- Flannery, M. (1996). Financial crises, payment system problems, and discount window lending. *Journal of Money, Credit and Banking*, 28(4), 804–824.

- Flood, M., & Mendelowitz, A. (2013). Monitoring financial stability in a complex world. In V. Lemieux (Ed.), *Financial analysis and risk management data governance, analytics and life cycle management*. (pp. 15–45). Heidelberg: Springer-Verlag.
- Folkerts-Landau, D. (1991). Systemic financial risk in payment systems, *Determinants and systemic consequences of international capital flows*. IMF Occasional Paper No. 77.
- Franck, R., & Schmied, A. (2003). *Predicting currency crisis contagion from East Asia to Russia and Brazil: An artificial neural network approach*. IMCB Working Paper No. 2.
- Frankel, J. A., & Rose, A. K. (1996). Currency crashes in emerging markets: an empirical treatment. *Journal of International Economics*, 41(3–3), 351–366.
- Freixas, X., & Parigi, B. (1998). Contagion and efficiency in gross and net interbank payment systems. *Journal of Financial Intermediation*, 7(1), 3–31.
- Fuertes, A.-M., & Kalotychou, E. (2006). Early warning system for sovereign debt crisis: the role of heterogeneity. *Computational Statistics and Data Analysis*, 5, 1420–1441.
- Furfine, C. (2003). Interbank exposures: quantifying the risk of contagion. *Journal of Money, Credit and Banking*, 35(1), 111–128.
- Galos, P., & Soramäki, K. (2005). *Systemic risk in alternative payment system designs*. ECB Working Paper No. 508.
- Geanakoplos, J. (2010). The leverage cycle. In: D. Acemoglu, K. R. Rogoff, & Woodford, M. (Eds.), *NBER Macroeconomics Annual 2009* (Vol. 24, pp. 1–65). Chicago: University of Chicago Press.
- Gonzalez-Hermosillo, B. (1999). *Determinants of ex-ante banking system distress: A macro-micro empirical exploration of some recent episodes*. IMF Working Papers No. 99/33.
- Gonzalez-Hermosillo, B., Pazarbasioglu, C., & Billings, R. (1997). Banking system fragility: likelihood versus timing of failure—an application to the Mexican financial crisis. *IMF Staff Papers*, 44(3), 295–314.
- Goodhart, C., Hartmann, P., Llewellyn, D., Rojas-Suarez, L., & Weisbrod, S. (1998). *Financial regulation: Why, how and where now?*. London: Routledge.
- Gorton, G. (1988). Banking panics and business cycles. *Oxford Economic Papers*, 40(4), 751.
- Gourinchas, P.-O., Landerretche, O., & Valdés, R. (2001). Lending booms: Latin America and the world. *Economía*, 1(2), 47–99.
- Gropp, R., Lo Duca, M., & Vesala, J. (2009). Cross-border bank contagion in Europe. *International Journal of Central Banking*, 5(1), 97–139.
- Guttentag, J., & Herring, R. (1984). Credit rationing and financial disorder. *The Journal of Finance*, 39(5), 1359–1382.
- Hakkio, C., & Keeton, W. (2009). Financial stress: what is it, how can it be measured and why does it matter? [Federal Reserve Bank of Kansas City Economic Review Q2], (pp. 5–50).
- Haldane, A., Hall, S., Saporta, V., & Tanaka, M. (2004). Financial stability and macroeconomic models. [Bank of England Financial Stability Review], (June).
- Hartmann, P., Straetmans, S., & de Vries, C. (2005). *Banking system stability: A cross-atlantic perspective*. NBER Working Papers No. 11698.
- Hautsch, N., Schaumburg, J., & Schienle, M. (2011). *Financial network systemic risk contributions*. SFB 649 Discussion Paper 2011–072, Humboldt University zu Berlin.
- Hellwig, M. (1994). Liquidity provision, banking, and the allocation of interest rate risk. *European Economic Review*, 38(7), 1363–1389.
- Hirtle, B., Schuermann, T., & Stiroh, K. (2009). *Macprudential supervision of financial institutions: Lessons from the scap*, Staff Report No. 409, Federal Reserve Bank of New York.
- Holló, D., Kremer, M., & Lo Duca, M. (2012). *CISS—a composite indicator of systemic stress in the financial system*. ECB Working Paper No. 1426.
- Humphrey, D. (1986). Payments finality and risk of settlement failure. In A. Saunders & L. White (Eds.), *Technology and the regulation of financial markets: Securities, futures, and banking* (pp. 97–120). Lexington: Lexington Books.
- Illing, M., & Liu, Y. (2006). Measuring financial stress in a developed country: an application to Canada. *Journal of Financial Stability*, 2(3), 243–265.

- IMF. (2009). Assessing the systemic implications of financial linkages. *Global financial stability report April, international monetary fund*. Washington, (pp. 73–110).
- Ioannidou, V., Ongena, S., & Peydro, J. (2009). *Monetary policy, risk-taking, and pricing: Evidence from a quasi-natural experiment*. Tilburg University Discussion Paper 2009–31 S.
- Iyer, R., & Peydró, J. (2011). Interbank contagion at work: evidence from a natural experiment. *The Review of Financial Studies*, 24(4), 1337–1377.
- Jagannathan, R. (1988). Banking panics, information, and rational expectations equilibrium. *The Journal of Finance*, 43(3), 749–761.
- Jiménez, G., Ongena, S., Peydró, J., & Saurina, J. (2007). *Hazardous times for monetary policy: What do twenty-three million bank loans say about the effects of monetary policy on credit risk?* CEPR Discussion Papers 6514.
- de Jonghe, O. (2010). Back to the basics in banking? a micro-analysis of banking system stability. *Journal of Financial Intermediation*, 19(3), 387–417.
- Kahn, C., McAndrews, J., & Roberds, W. (2003). Settlement risk under gross and net settlement. *Journal of Money, Credit and Banking*, 35(4), 591–608.
- Kaminsky, G. L., & Reinhart, C. M. (1996). *The twin crises: The causes of banking and balance of payments problems*. Federal Reserve Board Discussion Paper No. 544.
- Kaminsky, G., Lizondo, S., & Reinhart, C. (1998). Leading indicators of currency crises. *IMF Staff Papers*, 45(1), 1–48.
- Kida, M. (2008). *A macro stress testing model with feedback effects*. Reserve Bank of New Zealand Discussion Paper Series DP2008/08, Reserve Bank of New Zealand.
- Kindleberger, C. (1978). *Manias, panics, and crashes: A history of financial crises*. New York: Wiley.
- Kiyotaki, N., & Moore, J. (2002). Balance sheet contagion. *American Economic Review*, 92(2), 46–50.
- Knedlik, T., & von Schweinitz, G. (2012). Macroeconomic imbalances as indicators for debt crises in Europe. *Journal of Common Market Studies*, 50(5), 726–745.
- Leinonen, H. (Ed.). (2005). *Liquidity, risks and speed in payment and settlement systems—a simulation approach*. Bank of Finland Studies, E:31.
- Leinonen, H. (Ed.). (2007). *Simulation studies of liquidity needs, risks and efficiency in payment networks*. Bank of Finland Studies, E:39.
- Leinonen, H. (Ed.). (2009). *Simulation analyses and stress testing of payment networks*. Bank of Finland Studies, E:42.
- Leitner, Y. (2005). Financial networks: contagion, commitment, and private sector bailouts. *The Journal of Finance*, 60(6), 2925–2953.
- van Lelyveld, I., & Liedorp, F. (2006). Interbank contagion in the dutch banking sector: a sensitivity analysis. *International Journal of Central Banking*, 2(2), 99–133.
- Lo Duca, M., & Peltonen, T. (2013). Assessing systemic risks and predicting systemic events. *Journal of Banking and Finance*, 37(7), 2183–2195.
- Malkiel, B. (2003). Efficient market hypothesis and its critics. *The Journal of Economic Perspectives*, 17, 59–82.
- Männasoo, K., & Mayes, D. (2009). Explaining bank distress in eastern european transition economies. *Journal of Banking and Finance*, 33(2), 244–253.
- Marghescu, D., Sarlin, P., & Liu, S. (2010). Early-warning analysis for currency crises in emerging markets: a revisit with fuzzy clustering. *Intelligent Systems in Accounting, Finance and Management*, 17(3–4), 143–165.
- McNeil, A., Frey, R., & Embrechts, P. (2005). *Quantitative risk management: Concepts techniques and tools*. Princeton: Princeton University Press.
- Mian, A., & Sufi, A. (2009). The consequences of mortgage credit expansion: evidence from the u.s. mortgage default crisis. *The Quarterly Journal of Economics*, 124(4), 1449–1496.
- Minsky, H. (1977). A theory of systemic fragility. In E. Altman & A. Sametz (Eds.), *Financial crises*. Wiley-Blackwell: New York.

- Minsky, H. (1982). *Can "it" happen again?: Essays on instability and finance*. Armonk: M.E Sharpe.
- Mishkin, F. (1991). Asymmetric information and financial crises: A historical perspective. In G. Hubbard (Ed.), *Financial markets and financial crises*. Chicago: University of Chicago Press.
- Mistrulli, P. (2011). Assessing financial contagion in the interbank market: maximum entropy versus observed interbank lending patterns. *Journal of Banking and Finance*, 35(5), 1114–1127.
- Nag, A., & Mitra, A. (1999). Neural networks and early warning indicators of currency crisis. *Reserve Bank of India Occasional Papers*, 20(2), 183–222.
- Oosterloo, S., & de Haan, J. (2004). Central banks and financial stability: a survey. *Journal of Financial Stability*, 1(2), 257–273.
- Padoa-Schioppa, T. (2003). Central banks and financial stability: Exploring a land in between. In V. Gaspar, P. Hartmann, & O. Sleijpen (Eds.), *The Transformation of the European financial system* (pp. 1–43). Frankfurt: European Central Bank.
- Peavy, J., & Hempel, G. (1988). The penn square bank failure: effect on commercial bank security returns—a note. *Journal of Banking and Finance*, 12(1), 141–150.
- Peltonen, T. (2006). *Are emerging market currency crises predictable? A test*. ECB Working Paper No. 571.
- Quagliariello, M. (Ed.). (2009). *Stress-testing the Banking System: Methodologies and applications*. Cambridge: Cambridge University Press.
- Repullo, R., Saurina, J., & Trucharte, C. (2010). Mitigating the pro-cyclicality of basel ii. *Economic Policy*, 25(64), 659–702.
- Resta, M. (2009). Early warning systems: An approach via self organizing maps with applications to emergent markets. In B. Apolloni, S. Bassis, & M. Marinaro (Eds.), *Proceedings of the 18th Italian workshop on neural networks* (pp. 176–184). Amsterdam: IOS Press.
- Van Rijckeghem, C., & Weder, B. (2003). Spillovers through banking centers: a panel data analysis of bank flows. *Journal of International Money and Finance*, 22(4), 483–509.
- Rochet, J., & Tirole, J. (1996). Interbank lending and systemic risk. *Journal of Money, Credit and Banking*, 28(4), 733–762.
- Sarlin, P. (2014a) Macroprudential oversight, risk communication and visualization. [arXiv:1404.4550](https://arxiv.org/abs/1404.4550).
- Sachs, J., Tornell, A., & Velasco, A. (1996). *Financial crises in emerging markets: The lessons from 1995*. Brookings Papers on Economic Activity 1, (pp. 147–218).
- Saunders, A., & Wilson, B. (1996). Contagious bank runs: evidence from the 1929–1933 period. *Journal of Financial Intermediation*, 5(4), 409–423.
- Scharfstein, D., & Stein, J. (1990). Herd behavior and investment. *American Economic Review*, 80(3), 465–479.
- Schinasi, G. (2004). *Defining financial stability*. IMF Working Paper No. 04/187.
- Schmidt, R. (1984). Early warning of debt rescheduling. *Journal of Banking and Finance*, 8(2), 357–370.
- Shiller, R. (1989). *Market volatility*. Cambridge: MIT Press.
- Slovin, M., Sushka, M., & Polonchek, J. (1993). The value of bank durability: borrowers as bank stakeholders. *The Journal of Finance*, 48(1), 247–266.
- Smirlock, M., & Kaufold, H. (1987). Bank foreign lending, mandatory disclosure rules and the reaction of bank stock prices to the mexican debt crisis. *The Journal of Business*, 60(3), 347–364.
- Soramäki, K., & Cook, S. 2012. *Algorithm for identifying systemically important banks in payment systems*. Economics Discussion Papers, No. 43/2012, Kiel Institute for the World Economy.
- Soramäki, K., Bech, M., Arnold, J., Glass, R., & Beyeler, W. (2007). The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications*, 379(1), 317–333.
- Sorge, M. (2004). *Stress-testing financial systems: An overview of current methodologies*. BIS Working Paper No 165.
- Stiglitz, J. (1993). The role of the state in financial markets. *Proceedings of the World Bank Annual Conference on Development Economics*. Washington: World Bank.

- Straetmans, S., Verschoor, W., & Wolff, C. (2008). Extreme us stock market fluctuations in the wake of 9/11. *Journal of Applied Econometrics*, 23(1), 17–42.
- Trichet, J. (2009). *Systemic risk: Clair distinguished lecture in economics and public policy*. Cambridge: Cambridge University. (December, 2009).
- Upper, C. (2007). *Using counterfactual simulations to assess the danger of contagion*. BIS Working Paper Series No 234.
- Upper, C., & Worms, A. (2004). Estimating bilateral exposures in the german interbank market: is there a danger of contagion? *European Economic Review*, 48(4), 827–849.
- Wall, L., & Peterson, D. (1990). The effect of continental illinois' failure on the financial performance of other banks. *Journal of Monetary Economics*, 26(1), 77–99.
- Wetherilt, A., Zimmerman, P., & Soramäki, K. (2010). *The sterling unsecured loan market during 2006–08: Insights from network theory*. Bank of England Working Papers No. 398.

Chapter 3

Macroprudential Data

I didn't have time to write a short letter, so here's a long one

– Mark Twain

An understanding of all elements in the macroprudential oversight process is obviously crucial for safeguarding financial stability. While providing a basis, such a framework is still highly dependent on the underlying data. Access to complete, accurate, and timely data is central not only for policymakers to make good economic policy, but also for businesses and investors alike to make good financial decisions. However, data for macroprudential purposes are, not surprisingly, as complex as the system they describe. Alas, complexity oftentimes implies challenges. Gathering, synthesizing, understanding and analyzing these data is hence not an entirely unproblematic task. With the aim of having a holistic view of the financial system to ensure system-wide stability, rather than only being concerned about individual financial institutions, a macroprudential approach to oversight has a wide range of data demands and needs. As early-warning models were at the core of the previous chapter's ending note, the key focus herein is also on input data for early-warning exercises. Yet, as macro stress-testing and contagion models will throughout this book be touched upon, this chapter will still provide a brief discussion on data needs for risk assessment tools as well.

The focus of this chapter is on attempting to clarify what macroprudential data consist of, from where they are derived, how complex they are and how their properties may or may not hinder analysis. Hence, after defining the concept of data considered herein, this chapter provides a brief overview of data for macroprudential oversight and untangles the data into a four-dimensional cube representation. Finally, this section discusses stylized challenges related to macroprudential data and summarizes the key implications for this book.

3.1 Data: What are They?

The tools and models for risk identification and assessment described in Sect. 2.3 (and summarized in Fig. 2.4) made use of four broad sources of information:

- (i) open financial and macroeconomic data and statistics,
- (ii) supervisory data and statistics,
- (iii) domain intelligence, and
- (iv) experience and judgment.

To enable a discussion of data in macroprudential oversight, we need to start by defining the notion of data. In a broad sense, everything that can be encoded may be seen as data. However, the early definitions differ significantly from the today's notion of data. Fry and Sibley (1976) follow its Latin origin by defining data as a collection of facts. As a fact by definition cannot be false, the imprecision and inaccuracies often found in data today suggest this definition to be unsound. Hence, later attempts have softened the definition of the notion, such as “*data are facts or are believed to be facts which result from the observation of physical phenomena*” (Yovits 1981). Lately, the definition has significantly broadened. Along the lines of one definition in a recent Delphi study of more than 50 leading scholars (Zins 2007), data can be defined *as everything that can be encoded and stored in a computer*.

The implication of this definition is that it includes two of the above presented sources, *publicly available* and *supervisory statistics*, while leaving out the “softer” notions of implicit domain intelligence and judgment. The former may, for instance, consist of numerical and textual data and oftentimes function as an input to tools and models for risk identification and assessment (see Sect. 2.3), or as stand-alone measures for monitoring various risks, vulnerabilities and imbalances. The latter sources are, on the other hand, used for interpreting the results of tools and models in particular, and during the overall process of macroprudential oversight in general. Domain intelligence comprises various dimensions, such as market, policy and institutional intelligence (e.g., understanding the role and risks of financial innovations), whereas experience and judgment may, likewise, relate to a wide variety of topics, ranging from the functioning of the financial system to statistical methods. These types of qualitative information are an important complement to quantitative data in assessing the soundness of financial systems. Schou-Zibell et al. (2010) divide the qualitative dimensions to the following elements:

- (i) institutional processes;
- (ii) legal infrastructures and regulatory frameworks governing financial operations;
- (iii) practices and standards with respect to disclosure and accounting;
- (iv) surveillance and supervision of banks and other financial institutions;
- (v) incentive structures; and
- (vi) safety nets to cover overexposure to international financial markets.

The authors exemplify the above six qualitative tasks with, for instance, compliance with the core principles of the Bank for International Settlements (BIS),

International Organization of Securities Commissions, and the International Association of Insurance Supervisors. Such qualitative information may aid in understanding the inner reasons for the behavior of banks and markets. It is not, however, straightforward to combine these types of qualitative information with current theory and historical experiences of financial crises, not to mention the multitude of quantitative models. Schou-Zibell et al. (2010) also stress the importance of structural information in assessing how a financial system works. While structures in banks may often be covered by data, this relates more to market intelligence and country surveillance than to automated approaches for identifying vulnerabilities and risks. Structural assessment may, for instance, be a combined analysis of the structure of banks and their relative size, business strategy, ownership, concentration, and competitive situation. However important qualitative information is, the key focus herein is on numerical data.

3.2 Data for Macroprudential Oversight

Data needs and demands for macroprudential oversight are set by a broad range of issues. First, the availability of data obviously restricts the types of inputs to tools and models used by policymakers. Second, the understanding of the financial system, its fragilities and instabilities and the general oversight process defines what a policymaker understands to demand. Third, the design of the tools and models used for the task at hand set their final nuance to the data needs.

Early-warning exercises commonly make use of a wide range of indicators, measuring various dimensions of risks, vulnerabilities and imbalances. In this book, macroprudential data are related to three different categories:

- (i) macroeconomic data,
- (ii) banking system data, and
- (iii) market-based data.

Generally, the key three sources of macroprudential data measure the behavior of three low-level entities: households, firms and assets. By grouping data for the entities, we may produce data on various levels of aggregation. While firm-level data may also be of interest in the case of systemically important financial institutions (SIFIs), the data for macroprudential analysis most commonly refer to high-levels or aggregations of three kinds [see, e.g., Woolford (2001)]: macroeconomic, banking system, and financial market behavior. Hence, for macroprudential purposes, low-level entities may be aggregated as follows: from data on individual households' actions to the macroeconomic, from data on banks to the banking system, and from data on individual assets to the financial market. For instance, an entity could be a country, which would be described by country-level aggregates of macroeconomic, banking system, and financial market behavior. It is still worth to note that a system-wide approach does not always necessitate aggregation, as an entire system may, for

instance, be viewed from the perspective of a network of entities. Further, the category aggregating banks to the banking system may likewise be defined in broader terms (e.g. financial intermediaries in general) or some other type of financial intermediaries (e.g., insurers).

Yet, these three categories do not perfectly cover all types of data relevant for macroprudential oversight, especially not novel unexplored sources. These data relate, for instance, to texts and discussions (e.g., news articles, blogs or discussion forums) and tracking human behavior (e.g., search-terms used in Google and buying behavior). Whereas text has, for instance, been utilized for mapping bank interrelations [see, e.g., Rönqvist and Sarlin (2013)], trends in Google searches have been used for nowcasting macroeconomic data with long publication lags [see, e.g., Carrière-Swallow and Labbé (2013)]. The focus herein is, however, on the above mentioned three categories of numerical data, and on an overview of their use as indicators. The below discussion is supported by a long, yet incomplete, list of indicators along all three categories in Table 3.1.

3.2.1 Macroeconomic Data

Macroeconomic data can be transformed to measure risks and vulnerabilities of economic activity on a country level, and may hence function as leading indicators. Rather than being narrow in definitions, many macroeconomic measures provide a broad picture of overall economic and financial activity, as well as general circumstances, in the entire economy or a particular area of it, such as economic and production growth, current account balance and inflation. Trends, and deviations from them, indicate not only broad economic development in general, but also whether quantities and prices are consistent with prospects, such as in credit markets. For instance, vulnerabilities and risks to financial stability may be represented through above-normal and sustained rates of growth or valuation of credit and investment.

The production of macroeconomic data involves a laborious and costly aggregation process to derive figures that represent all households in an economy. The data are obviously not only of interest for domestic analysis, but also for various cross-country comparisons. This has stimulated a wide range of attempts to harmonize macroeconomic measures. Explicitly aiming at standardizing macroeconomic data across countries, the United Nations have issued their System of National Accounts (SNAs) in 1953 and its revised versions in 1968, 1993 and 2008 (see United Nations (2008) for the latest version). Likewise, the International Monetary Fund (IMF) has issued a Balance of Payments Manual to provide an accounting standard for reporting of balance of payments statistics (see IMF (2008) for the latest version). The title of the version in 2008, in contrast to the versions in 1948, 1950, 1961, 1977 and 1993, has been amended to Balance of Payments and International Investment Position Manual to reflect that it now covers both transactions and stocks of the related financial assets and liabilities alike.

Table 3.1 Examples of macroprudential indicators

Macroeconomic data	Banking system data	Market-based data
<i>Macroeconomic indicators</i>	<i>Banking system indicators</i>	<i>Market-based indicators</i>
Internal indicators	Capital adequacy	Asset valuation
GDP growth	Equity to assets	Equity prices
Unemployment	Tier 1 and 2 ratio	Bond spreads
Inflation	Asset quality	Derivative valuation
Debt imbalances	Impaired assets	CDS prices
Credit imbalances	Non-performing loans	Option-adjusted spread
House prices	Loan loss provisions	Credit ratings
External indicators	Debt to equity	Sovereign ratings
Current account balance	Return on assets	Firm ratings
External investment position	Management	Credit spreads
Unit labor costs	Cost to income	Sovereign yield spread
Real exchange rate	Earnings	Default probabilities
Export market share	Return on equity	Distance-to-default
	Net interest margin	Bond default probabilities
	Liquidity	
	Liquid assets to liquid liabilities	
	Interest expenses to liabilities	
	Deposits to funding	
	Loans to deposits	
	Sensitivity to market risk	
	Share of trading income	
	Loans to assets	
	Net open position in foreign exchange to capital	
	Net open position in equities to capital	
<i>Macroeconomic linkages</i>	<i>Banking sector linkages</i>	<i>Market-based co-movements</i>
Equity and debt exposures	Equity and debt exposures	Asset and derivative interdependence

Notes The table draws upon compilations in Betz et al. (2014), Cihák (2006), IMF (2006), Woolford (2001). The table presents three types of indicators: macroeconomic, banking system and market-based. Macroeconomic indicators may be defined to describe different sectors, such as private and government sector. Banking system indicators are defined on the country level, but may also be measured per firm if needed, as oftentimes is for SIFIs. Likewise, market-based indicators may be used on an entity or aggregated market level, as needed. Following IMF (2006), credit ratings are classified as market-based indicators as they are produced mainly for use by market participants. The table does not discuss how the data may be transformed. Hence, each mentioned indicator may address different imbalances depending upon its transformation

Lately, multiple initiatives mainly run by the IMF have attempted and also prompted progress in data provision. In 1996, the IMF established the Special Data Dissemination Standard (SDDS) (see (IMF 2007c) for the latest version) to

guide member countries in providing national economic and financial statistics to the public. The SDDS is the first of a two-tier data standards initiative with the general aim of improving access to comprehensive, timely and accurate data to facilitate macroeconomic policies and the functioning of financial markets. To function as a development tool to prepare for SDDS subscription, the IMF established the second tier in 1997, called the General Data Dissemination System (GDSS) (see IMF (2007b) for the latest version). Likewise, the Data Quality Reference Site (DQRS) was established by the IMF in 2000 to foster a common understanding and importance of data quality.

The national accounts may further be complemented with balance-sheet exposures between aggregated entities, such as economies. These types of cross-border exposures represent crucial links in the global economy. Since 2001, the IMF has published data on bilateral portfolio investment positions among economies on an annual basis. The data have been collected through the annual Coordinated Portfolio Investment Survey. Likewise, the Coordinated Direct Investment Survey collects bilateral position data on direct investments among economies.

3.2.2 Banking System Data

Banking system data utilize, usually in the form of ratios, aggregated country-level information collected from balance sheets and income statements of individual financial institutions. The need for macroprudential assessment of financial conditions on the level of banking systems, rather than only a microprudential, or institution-level, approach, has been accentuated not only by the ongoing financial crisis, but also by the Asian financial crisis in the late 1990s. San Jose and Georgiou (2008) describe that vulnerabilities in Asia were related to international capital flow reversals, also involving shocks to the corporate and household sectors, whereas the recent wave of distress stemming from the sub-prime mortgage markets highlights the importance of balance-sheet exposures of financial institutions and vulnerabilities to credit and liquidity squeezes. Likewise, from a European viewpoint, the increasing integration of national financial systems has stimulated efforts to develop a common framework for financial stability analysis (Agresti et al. 2008).

The need for data to assess strengths and weaknesses in financial systems led to attempts to derive a commonly accepted list of financial stability indicators, not the least the financial soundness indicators (FSIs) developed at the IMF. Sundararajan et al. (2002) were the first to propose sets of so-called “core” and “encouraged” FSIs. The FSIs are measures of the current aggregated financial health and soundness of the financial institutions in an economy. A final list, with more precise definitions of the FSIs, was laid down in a set of indicators compiled by the IMF (2006) in the *Compilation Guide on Financial Soundness Indicators* (henceforth the *Guide*). IMF (2006) puts forward a handbook on concepts and definitions, as well as sources and techniques, for compiling and disseminating FSIs. For macroprudential surveillance, the key indicators are based upon aggregated information contained in

the balance sheets and income statements of individual financial institutions. The literature on individual bank failures draws heavily on the Uniform Financial Rating System, informally known as the CAMEL ratings system, introduced by U.S. regulators in 1979, where the letters refer to Capital adequacy (e.g., risk-based capital ratio), Asset quality (e.g., nonperforming loans to capital), Management quality (e.g., cost to income), Earnings (e.g., return on equity) and Liquidity (e.g., deposits to funding). Since 1996 the rating system also includes Sensitivity to Market Risk (e.g., net open position in equities to capital, which derives CAMELS). To implement the FSIs in the Guide, the IMF invited its members to participate in a Coordinated Compilation Exercise (CCE), which eventually led to 62 participating countries and regions (IMF 2007a).

In the European context, the European Central Bank (ECB), jointly with the Banking Supervision Committee (BSC) of the European System of Central Banks, have put efforts into developing their own financial stability indicators, called macroprudential indicators (MPIs) (see Mörttinen et al. (2005) for an overview of the methodology). The aim of the MPIs is defined to be to gauge conditions in the financial system and its resilience to stress situations. While differing in terms of the aim, the scope of FSIs and MPIs is analogous. The MPIs were reported and analyzed in the European Union (EU) Banking Sector Stability report prepared by the BSC until 2010, whereafter the data have only been reported in the Consolidated Banking Data, a dataset published in the ECB Statistical Data Warehouse.

Cross-border linkages among banking sectors is obviously a potential contagion channel (as also noted in Sect. 2.2), when assessing interdependence of the global economy. The BIS has been collecting international banking statistics with bilateral partner-country information on both a locational basis and on a consolidated group basis. Likewise, to assess the system-wide risk within countries, balance-sheet exposures between individual banks are of central interest.

3.2.3 Market-Based Data

Market-based data exploits aggregated information dispersed among financial market participants. The rationale for using market data is that prices of financial instruments, such as equities, bonds and options, capture forward-looking perceptions of financial market participants, not least related to vulnerabilities and risks in the financial system. Rather than being a substitute for the previous sources of information, market-based data complements analysis by conveying the view of financial market participants. Lately, joint efforts by the IMF, ECB and BIS have been put forward to assist the reporting and production of coherent, relevant and comparable securities statistics for use in financial stability analysis and monetary policy formulation. In a three-part series, the Handbook of Securities Statistics was published in 2009, 2010 and 2012 (BIS-ECB-IMF 2009, 2010, 2012).

Market-based data capture the perceptions of markets about vulnerabilities and risks in the financial system. The degree of system-wide risk may be measured by,

for instance, yields and spreads of financial instruments, asset prices, externally measured creditworthiness and sovereign ratings, interest rates, exchange rates and stock market volatility. Depending then on how these data are transformed, they function as forward-looking measures of the health of the financial system. They may be, for instance, changes in government or corporate bond spreads, relative stock-market prices, and indicators of volatility in share prices [e.g., Cihák (2006)].

Moreover, market-based data are oftentimes transformed into some more advanced stand-alone measures of default probability. One indicator that has gained large attention is Merton's (1974) distance-to-default, which uses a structural valuation model to compute the ratio of a firm's assets to debt. To be forward-looking, asset value and volatility is, however, estimated from equity data. Since supervisors commonly intervene before capital is depleted, Chan-Lau and Sy (2006), Nationalbank (2004) present two alternative, but similar, measures: distance-to-capital and distance-to-insolvency. Likewise, bond prices may be turned into a default probability by Fons' (1987) function of the additional required rate of return over default-free bonds. A more direct measure of default probability may be obtained from credit default swaps (CDSs). CDSs provide an insurance against default, where the seller guarantees protection by compensating the buyer in the event of a default of the reference obligor during the life of the contract and the buyer pays a quarterly fee (i.e., the CDS spread). The default probability is then calculated from the CDS spread, interest rate of default-free bonds and recovery rate (i.e., the amount recovered in event of a default). While being defined on the firm level, these measures can obviously be aggregated through simple or weighted averages or measures for entire portfolios. However, due to the existence of large co-movements in market-based data, aggregating these indicators from the entity level to the systemic level poses a number of challenges that still remain to be solved. One suggestion is the indicator by Cihák (2007) that attempts to account for correlation of defaults across institutions in an aggregate measure of financial stability.

These data may also be used to compute interdependence among economies. For this task, one can compute co-movements in country-specific market data, such as stock market indices, CDS spreads and bond spreads. Yet, the most common approach is to make use of firm-level data, in order to assess co-movements in their asset prices [see, e.g., Hautsch et al. (2011)].

3.3 A Four-Dimensional Data Cube

The previous section related macroprudential data to three key sources: macroeconomic data, banking system data and market-based data. Yet, the discussion provided little structure on the form and complexity of the data. Based upon the above discussion, macroprudential oversight can be said to utilize data that come from a so-called macroprudential data cube (henceforth data cube). The characteristics and challenges associated with macroprudential data can subsequently be paired

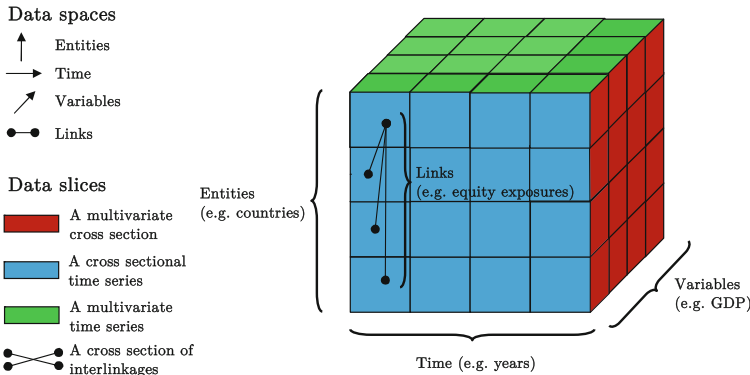


Fig. 3.1 A macroprudential data cube. *Notes* The figure represents the macroprudential data cube. It represents four spaces: entities (e.g., country), time (e.g., year), variables (e.g., GDP), and links (e.g., debt and equity exposures). Likewise, it illustrates four data slices: a multivariate cross section (*red side*), a cross section of time series (*blue side*), a multivariate time series (*green side*), and a cross section of interlinkage matrices (**black edges**)

with this data cube representation. Rather than three, the data cube in Fig. 3.1 is described by four dimensions:

- (i) entities (e.g., countries);
- (ii) time (e.g., years);
- (iii) variables [e.g., gross domestic product (GDP)];
- (iv) links (e.g., debt and equity exposures).

Each cell is hence defined by a specific cross-sectional entity, a specific time unit, a specific variable (or in computer science so-called input or feature vector), and a specific network of interlinkages. The value for each cell is the value for that particular variable and the related vector of links.

Following the four dimensions, the data cube can be described according to four types of slices. First, a multivariate cross section (red side) provides a view of multiple entities described by multiple variables at one point in time. Second, a cross section of time series (blue side) is a univariate view of multiple entities over time. Third, a multivariate time series (green side) provides a view of multiple variables over time for one entity. Finally, the fourth view is a cross section of interlinkage matrices (**black edges**) that represent links between multivariate entities at one point in time. As noted in the previous section, while links may be estimated from interdependence in the variable dimension (e.g., equity prices), a more common and less noisy measure is direct linkages and exposures between entities. To exemplify a macroprudential dataset in the data cube representation, the four dimensions could be defined as follows: countries as entities, quarterly frequency as time, indicators of various sources of risk and vulnerability as variables, and equity and debt exposures between economies as links.

3.4 Stylized Challenges in Macprudential Data

The more commonly used term ‘stylized fact’ refers to a broad generalization of a complex occurrence—which may be imprecise in the detail, but essentially true. This section presents stylized challenges related to the use of macroprudential data by not delving into atomic detail, but rather focusing on more general concerns and relative prominence of different data sources. The data discussed in Sect. 3.2 are problematic due to numerous reasons. Not only is it difficult to identify relevant data from the vast amounts available, but there are also challenges of their own in compiling the components needed for the macroprudential approach both within individual economies and across economies, not to mention challenges related to time. Schou-Zibell et al. (2010) pinpoint that the key concerns related to macroprudential data are excessive and frequently reoccurring delays, inaccuracies, inadequacies and incompleteness of data. The authors relate it to five major reasons:

- (i) spread of data in various databases and institutions;
- (ii) non-availability or non-applicability of some indicators;
- (iii) incomparability of indicators over time owing to the absence of or changes in accounting and prudential standards;
- (iv) lack of transparency and problems in the disclosure of data; and
- (v) late, incomplete, and inaccurate replies from participating institutions and agencies.

These five reasons can, however, be complemented. In that vein, we untangle the challenges according to two dimensions of the data cube: (i) temporal and (ii) cross sectional. The former relates to major challenges related to temporality and nonstationarity, whereas the latter relates to heterogeneity of countries in the cross section.

With regards to *temporality*, the indicators on the watchlist are prone to change over time, not least in the wake of the ongoing global financial crisis. A large number of papers and projects, often led by supervisory authorities, have viewed possible data gaps [see, e.g., Burgi-Schmelz (2009)]. Gaps seem to exist on individual, sectoral and market levels, where most frequently mentioned gaps are related to the real estate, corporate, and household sectors, as well as to nonbank financial institutions. This not only accentuates the importance of access to the data of latest relevance, but also imposes challenges in the application of analytical tools as the time dimension is commonly short for new types of data. In addition, with the aim to use historical data to infer about the future, nonstationarity may also easily become a problem. Bisias et al. (2012) note that the field of econometrics has provided a large number of techniques to address specific types of nonstationarities, such as deterministic and stochastic trends and cointegration relationships, whereas the type of nonstationarity that complicates risk assessment and identification, such as political institutional and cultural changes, is less easily dealt with through transformations or parametrizations. For instance, complex financial instruments, such as CDSS and collateralized debt obligations, as well as high-frequency trading in general, were not part of the risk assessment and identification agenda in the beginning of the 1990s, whereas they are in the core of today’s analysis.

Yet, an even more important and recurring challenge appears to be the deliberate migration of activities to areas which are not on the current watchlist. These included special purpose entities (e.g., in the context of loan securitization) and off-balance-sheet operations (e.g., in the context of hiding risky assets) during this crisis, but are likely to evolve into another form in the future. Another problem is the pace with which activities have moved to nonbank financial intermediaries not under the watchlist of regulatory and supervisory bodies. For instance, Feldman and Lueck (2007) show that the market share of “other financial intermediaries” has increased from less than 10% in the 1980s to about 45% in 2005, which does not yet include activities in hedge funds. The implications of the changing nature of risks and vulnerabilities relates obviously not only to data provision, but also in broad terms to macroprudential oversight and supervision in general. Likewise, due to advances in telecommunications, computer technology and financial innovation, Bisias et al. (2012) note that the intensity of activity in the financial sector has experienced a tremendous growth. More precisely, they pinpoint the challenges to the leisurely pace of quarterly financial reporting and annual examinations, as well as the failure of accounting standards in conveying all risk exposures due to light reporting requirements in unregulated markets.

The latter of the two challenges relates to heterogeneity in the *cross section*. The common problem of comparability is often cited as a cause of the lack of regular and uniform reporting of indicators for various types of financial institutions, such as nonbank financial institutions. For instance, Burgi-Schmelz (2009) provides a recent review of what has been achieved in the international collection, distribution and availability of statistical data, and highlights that a large number of gaps should be filled to further improve the coverage of statistical information, not the least dimensions accentuated by the ongoing global financial crisis. Another challenge is the identification of relevant indicators to signal risks and vulnerabilities in a particular financial system. Due to differences in financial and economic development, indicators useful in one country may not necessarily be useful for another. For instance, Schou-Zibell et al. (2010) relate cross-country differences in development to disparities in institutional and legal frameworks, the size and liquidity of financial markets and the versatility of financial instruments.

These remarks highlight challenges in overall use of these types of data for macroprudential oversight. The sequel of this section focuses on a comparative discussion of the identified three types of macroprudential data. The problematic nature of macroprudential data relates to a wide range of issues that hinder identification and assessment of risks and vulnerabilities. Hence, we look into challenges particular for each category of data.

3.4.1 Macroeconomic Data

Macroeconomic data have been harmonized through the SNA and are hence estimates of aggregates. Yet, statistical offices have had multiple reasons not to either participate

at all or to only fulfill partial requirements. Examples of reasons are, for instance, not being able to devote enough resources to implement the suggested harmonizations or preferring to stick to their own national income accounting rules [see, e.g., the case of the United States (US) in Mead et al. (2004)]. Hence, national accounting practices still have substantial differences. Further, Hartwig (2006) shows differences in the used deflators of Switzerland and the SNA. In a larger context, Hartwig (2007) illustrates that partial differences in economic growth in the SNA and Europe may be explained by different deflators, in particular a deflation method introduced in the SNA in 1997. Whilst National Income and Product Accounts, and the GDP, tend to be computed from the demand side in the US, the SNA uses the supply side, i.e., differences between gross output and intermediate inputs, to compute GDP. Most literature on accounting differences describe problems when comparing firm-level data in different countries using various reporting standards. The above mentioned issues illustrate, however, the existence of similar problems with the country-level aggregates.

Moreover, uncertainty in data are often caused through survey-based collection. For instance, an economy's value added and employment may be collected through household surveys. These, however, comprise only a small percentage of the population, may not always have reliable answers and are difficult to extrapolate to the macro level. Bruyère and Chagny (2002) exemplify the uncertainty in surveys by showing that in most of 8 Organisation for Economic Co-operation and Development (OECD) countries labor input growth according to household surveys exceed counterparts retrieved through establishment surveys.

An issue of crucial importance in early-warning exercises is to take into account publication lags for data. For instance, GDP, money and credit related indicators have an approximate lag from 1 to 2 quarters depending on the country. One seldom discussed, yet important, issue is to account for revisions of macroeconomic data. When evaluating early-warning models, the data that would have been available in real time (i.e., preliminary first-releases) should be used. Likewise, while relating to all categories of data, it is also important to only use the available information set when performing transformations of variables that are dependent on the historical data distribution (e.g., detrending or percentiles). Missing values, as well as outliers, are obviously an issue on their own.

3.4.2 Banking System Data

The aggregation procedure of banking system data does not involve equally comprehensive procedures as macro data. Still, the constraints faced by empirical analyses have illustrated challenges in availability and quality of banking system data. Cihák and Schaeck (2010) collected banking system indicators with the aim of testing the early-warning capabilities of the FSIs listed in the Guide by the IMF (2006). However, to collect a sufficiently large dataset, they had to narrow down the focus to three FSIs from the core set (i.e., regulatory capital, asset quality and profitability of

deposit taking institutions) and two proxies for FSIs from the encouraged set (i.e., profitability and leverage). The low number of indicators was driven by the limited availability of data. Likewise, availability leads to time series of annual frequency that only date back to 1994. Still, many economies deviate from definitions in the Guide (IMF 2006), which is likely to increase cross-country differences. For instance, minor errors in the reporting of one element of an indicator, such as non-performing loans, is likely to impact a large number of other indicators.

Between 2004 and 2007, the IMF conducted the Coordinated Compilation Exercise (CCE) for FSIs, after which they report that a total of 57 out of the 62 participating countries submitted their data and metadata (IMF 2007a). Yet, the final report of the CCE revealed in FSIs cross-country diversity on account of four issues (IMF 2007a): (i) accounting and supervisory practices; (ii) data availability; (iii) additional data collection costs for fully implementing the FSIs; and (iv) views on how to compile the FSIs. This indicates that cross-country standardization is still a goal to be achieved. However, the metadata compiled by the CCE facilitate more informed cross-country comparisons and unifications. While the 2008 update of the Guide improved the FSIs, Agresti et al. (2008) point out that the MPIs better follow international accounting and supervisory standards and thus requires only few adjustments to original national banking sector data.

Schou-Zibell et al. (2010) also note that the commonly used weighted averages of indicators may lower their accuracy. Likewise, relating indicators to asset size, or other data measuring size of banks, implies an implicit assumption that small banks are not contributing to systemic risk. Even though Schou-Zibell et al. (2010) suggest the use of qualitative information to complement the quantitative assessments, there are obvious measurement problems. Qualitative information, such as poor banking supervision, is challenging to quantify, whereas a qualitative assessment across countries may vary significantly. Other factors that are important in predicting a crisis, but are difficult to measure, include the quality of corporate governance, independence of the national central bank, reliability of the legal system, political stability, and other institutional qualities.

The aggregation procedure to derive data on the banking system is rather easy. However, while balance sheets and income statements of individual financial intermediaries are simple to add up, comparisons of them have limitations not only due to differences in individual firms' business models, but also due to cross-country variance in business models and accounting standards [see, e.g., Nobes (2006)]. They and their analysis can, however, be treated in close to similar manners as financial ratios for individual firms. Likewise, they oftentimes also exhibit outliers and skewed distributions [see, e.g., Deakin (1976)]. An issue of even more crucial importance than with macroeconomic data is to lag variables such that they take into account publication delays. Although some economies report quarterly banking system data, others report only on an annual basis, which means that data for a reference period are in most cases available only in the second quarter of the following year. That is, at the time of writing, in February 2013, the available data would refer to 2011. Pointing at the fact that in today's quick paced financial world, the frequency of financial reporting falls short in granularity.

3.4.3 Market-Based Data

Relying on prices of assets and other financial instruments, one can create a battery of financial stability indicators. These have been shown to have merits for some tasks, whereas they still exhibit a range of weaknesses [see, e.g., IMF (2007a), Cihák (2006)]. Advantageous features of market-based indicators is availability at high frequency and short publication lags, as well as the rarity of missing values and the lack of differences in accounting standards. They function as a measure of market participants' forward-looking assessment of risks and vulnerabilities, in contrast to some more backward-looking accounting measures (e.g., nonperforming loans and loan loss reserves). It is also worth noting that these data are publicly available and widely accessible, unlike supervisory data. Yet, market-based data also have their limitations. Availability is not only restricted to publicly traded institutions, but also to those with non-limited trading, where examples of unsuitable entities are government- or family-owned companies. Moreover, the quality of market-based data is directly linked to how efficient the financial markets are. If markets are not liquid, robust and transparent, price changes may reflect other factors than the health of the issuer. Likewise, if public information related to an institution is limited (e.g., loan classification data in some economies), prudential information collected by supervisors through other sources may be of higher value. Moreover, the forward-looking assessment of financial market participants only accounts for potential losses to their holdings (e.g., equities and bonds), rather than losses to depositors or systemic effects in general. It is also worth noting that some more advanced market-based indicators are based on distributional assumptions. For instance, measures based upon the distance-to-default methodology assume that asset values are drawn from a lognormal process, which implies the absence of extreme tail events relevant for systemic risk assessments.

With regards to market-based data, the aggregation procedure is simple and easily automated. Yet, it is obvious that market-based data exhibit cross-country differences. The financial market architecture and infrastructure, as well as many trading activities, differ depending on the state of financial development in the country, such as differences between advanced and emerging economies. When transforming market-based indicators, one should consider that they are oftentimes contemporaneous measures of financial stress, and hence lagged transformations (e.g., moving averages) or deviations from trends (e.g., Hodrick-Prescott filtering) may improve the early-warning capacity of the indicators. Due to the effortless aggregation and data collection procedure, problems related to missingness and non-complete data ought to be relatively rare.

3.5 Concluding Discussion

Today's world has already for some time experienced access to ever-increasing amounts of data. Yet, this chapter has highlighted that big data does not necessarily imply good data. Crises have more often than not exposed weaknesses in

data. The crises of the 1990s (e.g., Mexican and Asian crises) prompted progress in data provision along multiple fronts. Examples of standards and efforts to improve data provision are the establishment of SDDS, GDDS and DQRS, as well as updates of previous establishments. Likewise, this crisis has revealed weaknesses in data provision. Burgi-Schmelz (2009) pinpoints issues highlighted by the current crisis to be lack of data on who holds what, the balance-sheets on nonbanks and contingent risks and derivative positions, in addition to the longstanding need for more accurate, complete, frequent and timely data. The IMF is working on these issues in two projects that hold promise for improving macroprudential data. The Data Link Project is an internal project and aims at developing a set of timely and higher-frequency indicators, initially for a number of systemically important economies. Externally, the IMF is chairing an interagency group on national statistics with the aim of a global website of economic and financial indicators. In the US, the Dodd-Frank Wall Street Reform and Consumer Protection Act created, among many other things, the Office of Financial Research to support the Financial Stability Oversight Council and its member agencies by providing financial research and data. The improvements with respect to data concern collecting and providing financial data of higher quality and with better accessibility and transparency. The creation of the European Systemic Risk Board points to similar efforts in Europe. Hence, turning these current large-volume data also into high-quality data remains to be a critical objective for the available tools for safeguarding financial stability to bear fruit.

Another key challenge for creating tools for monitoring threats to financial stability has been the limited access to data. While mostly being available, some data are restricted to only specific supervisory authorities. This has not only an effect on monitoring, but also on research. One example is that the limited access to bilateral interbank exposures has stimulated research on methods for circumventing the use of such data. For instance, so-called maximum entropy (Mistrulli 2011) and stochastic block modeling (Halaj and Kok 2013) have been used to estimate interbank exposures from larger aggregates. Whereas research along these lines obviously improves the current state of monitoring, it still wastes resources that could be spent on advancing the state of the art, rather than on circumventing data accessibility.

In addition to a number of challenges that remain to be solved, this chapter has illustrated multiple characteristics of data that need to be acknowledged. The chapter described that the complexity and dimensionality exhibited by macroprudential data is large, and new plans on improving data provision have been established. Yet, the key question remains: *What should we do with these data?* Obviously, the data not only enable, but also motivate designing tools for risk identification and assessment with the ultimate aim of risk communication. Aggregating multidimensional information into crisis probabilities, systemic risk indicators and other quantitative measures capturing the functioning and interconnectedness of financial systems provide means for supporting decisionmaking in general and policymaking in particular. Yet, visualizing these complex data in easily understandable formats not only provides means for binary decisions, but also enables *disciplined and structured judgmental*

analysis based upon policymakers' experience, as noted by Mr. Trichet in the quote prior to Chap. 2. As this is also the key focus of this book, the next chapter digs deeper into possible approaches for such exploratory means to analysis.

References

- Agresti, A., Baudino, P., & Poloni, P. (2008). *The ECB and IMF indicators for the macro-prudential analysis of the banking sector: A comparison of the two approaches*, ECB, Occasional Paper No. 99/08.
- Betz, F., Oprica, S., Peltonen, T., & Sarlin, P. (2014). Predicting distress in European banks. *Journal of Banking & Finance* (forthcoming).
- BIS-ECB-IMF. (2009). *Handbook on securities statistics*, Part 1: Debt securities issues.
- BIS-ECB-IMF. (2010). *Handbook on securities statistics*, Part 2: Debt securities holdings.
- BIS-ECB-IMF. (2012). *Handbook on securities statistics*, Part 3: Equity securities.
- Bisias, D., Flood, M., Lo, A., & Valavanis, S. (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics*, 4, 255–96.
- Bruyère, M., & Chagny, O. (2002). *The fragility of international comparisons of employment and hours worked. An attempt to reduce data heterogeneity*, Observatoire français des conjonctures économiques, Working Paper No. 05/2002.
- Burgi-Schmelz, A. (2009). Data to the rescue. *Finance & Development*, 46(1), 31–3.
- Carrière-Swallow, Y., & Labbé, F. (2013). Nowcasting with google trends in an emerging market. *Journal of Forecasting*, 32(4), 289–298.
- Chan-Lau, J., & Sy, A. (2006). *Distance-to-default in banking: A bridge too far?* IMF Working Paper No. 06/215.
- Cihák, M. (2006). *How do central banks write on financial stability?* IMF Working Paper No. 06/163.
- Cihák, M. (2007). Systemic loss: A measure of financial stability. *Czech Journal of Economics and Finance*, 57(1–2), 5–26.
- Cihák, M., & Schaeck, K. (2010). How well do aggregate prudential ratios identify banking system problems? *Journal of Financial Stability*, 6(3), 130–44.
- Danmarks Nationalbank. (2004). Market-based risk measures for banks. Financial Stability Report, pp. 85–95.
- Deakin, E. (1976). Distributions of financial accounting ratios: Some empirical evidence. *The Accounting Review*, 51, 90–96.
- Feldman, R., & Lueck, M. (2007). Are banks really dying this time? an update of Boyd and Gertler. *The Region* (Sep), 6–51.
- Fons, J. (1987). The default premium and corporate bond experience. *The Journal of Finance*, 42(1), 81–97.
- Fry, J., & Sibley, E. (1976). Evolution of database management systems. *ACM Computing Surveys*, 8(1), 7–42.
- Halaj, G., & Kok, C. (2013). *Assessing interbank contagion using simulated networks*, ECB Working Papers, No. 1506.
- Hartwig, J. (2006). Messprobleme bei der ermittlung des wachstums der arbeitsproduktivität - dargestellt anhand eines vergleichs der schweiz mit den USA. *Jahrbücher für Nationalökonomie und Statistik*, 226(4), 418–435.
- Hartwig, J. (2007). Is the transatlantic gap in economic growth really widening? In J. McCombie & C. Rodriguez (Eds.), *The European union: Current problems and prospects* (pp. 68–83). Basingstoke: Palgrave-Macmillan.
- Hautsch, N., Schaumburg, J., & Schienle, M. (2011). *Financial network systemic risk contributions*, SFB 649 Discussion Paper 2011–072, Humboldt University zu Berlin.

- IMF. (2006). *Financial soundness indicators: Compilation guide*. Washington: International Monetary Fund.
- IMF. (2007a). *Financial soundness indicators—experience with the coordinated compilation exercise and next steps*. Washington: International Monetary Fund.
- IMF. (2007b). *The general data dissemination system: Guide for participants and users*. Washington: International Monetary Fund.
- IMF. (2007c). *The special data dissemination standard: Guide for subscribers and users*. Washington: International Monetary Fund.
- IMF. (2008). *Balance of payments and international investment position manual*. Washington: International Monetary Fund.
- Mead, C. I., Moses, K. E., & Moulton, B. R. (2004). The NIPAs and the system of national accounts. *Survey of Current Business*, 17–32.
- Merton, R. (1974). On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29, 449–470.
- Mistrulli, P. (2011). Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking & Finance*, 35(5), 1114–1127.
- Mörttinen, L., Poloni, P., Sandars, P., & Vesala, J. (2005). *Analysing banking sector conditions—how to use macro-prudential indicators*, ECB Occasional Papers No. 26/2005.
- Nobes, C. (2006). The survival of international differences under IFRS: Towards a research agenda. *Accounting and Business Research*, 36(3), 233–245.
- Rönnqvist, S., & Sarlin, P. (2013). *From text to bank interrelation maps*, Åbo Akademi University, Mimeo.
- San Jose, A., & Georgiou, A. (2008). Financial soundness indicators (FSIs): framework and implementation, *Proceedings of the IFC Conference on “Measuring financial innovation and its impact”*, Bank for International Settlements, August 26–27 2008, Basel, pp. 277–282.
- Schou-Zibell, L., Albert, J. & Song, L. (2010). *A macroprudential framework for monitoring and examining financial soundness*, ADB Working Paper Series on Regional Economic Intergration, Asian Development Bank, March 2010.
- Sundararajan, V., Enoch, C., San José, A., Hilbers, P., Krueger, R., & Moretti, M. G. S. (2002). *Financial soundness indicators: Analytical aspects and country practices*, IMF Occasional Paper No. 212.
- United Nations. (2008). System of national accounts 2008, commission of the European communities/eurostat, international monetary fund, organisation for economic co-operation and development, United Nations, World Bank, Brussels.
- Woolford, I. (2001). Macro-financial stability and macroprudential analysis. *Reserve Bank of New Zealand Bulletin*, 64(3), 29–43.
- Yovits, M. (1981). Information and data. In A. Ralston (Ed.), *Encyclopedia of Computer Science and Engineering* (pp. 714–717). New York: Van Nostrand Reinhold.
- Zins, C. (2007). Conceptual approaches for defining data, information, and knowledge. *Journal of the American Society for Information Science and Technology*, 58(4), 479–493.

Chapter 4

Data and Dimension Reduction

The eye, which is called the window of the soul, is the principal means by which the central sense can most completely and abundantly appreciate the infinite works of nature

– Leonardo da Vinci

Data and dimension reduction techniques hold promise for representing data in easily understandable formats, as has been shown by their wide scope of applications. Data reductions provide summarizations of data by compressing information into fewer partitions, whereas dimension reductions provide low-dimensional overviews of similarity relations in data. Thus, these techniques provide means for exploratory data analysis (EDA). From a broader perspective, EDA is only one approach out of many in data mining, and knowledge discovery includes data mining as only one of its steps. To provide a holistic view in a top-down manner, we start by the broader concepts, and end with discussions of data and dimension reductions and their combination. As the aim of Chap. 5 is to provide a comparison of early dimension reduction methods, the focus of this chapter is also on more detailed presentations of so-called first-generation methods, including Multidimensional Scaling (MDS), Sammon's mapping and the Self-Organizing Map (SOM).

Along these lines, this chapter first presents an overview of EDA, knowledge discovery in databases (KDD), information visualization and visual analytics, and then focuses on reviewing methods for both data and dimension reduction. As the focus of this book lies on dimension reductions, and data reductions are mainly used for enhancing the interpretation of the dimension reductions, this chapter also has a greater focus on dimension reductions.

This chapter is partly based upon previous research. Please see the following work for further information: Sarlin (2014a)

4.1 Some Key Concepts and Definitions

To better position data and dimension reduction, this section takes a broad perspective on EDA. EDA, as vaguely defined by Tukey (1977), is numerical, counting and graphical detective work. However, none of Tukey's works seems to provide a precise and concise definition of EDA. In a later work with two collaborators, he provided yet another broad definition, but still one with somewhat more precision (Hoaglin et al. 1983): "*Exploratory data analysis isolates patterns and features of the data and reveals these forcefully to the analyst*". The focus of the field may thus be related to representing data in easily understandable formats, which might involve summarizing characteristics of interest with descriptive statistics or visual examinations. Rather than being an approach to test hypotheses, EDA concerns tasks supporting the formulation of hypotheses, which may be tested with other methods, and the assessment of assumptions in data, on which statistical inference may rely. Hence, Tukey (1977) defines EDA to be exploratory or descriptive in nature, whereas it is not concerned with confirmatory or inferential tasks, in which the focus is on using data to confirm a number of assumptions or the validity of a hypothesis or model.

In the following, this section proceeds by first viewing EDA from above, i.e., knowledge discovery and data mining, and then zooming in on the most central parts of its core, i.e., information visualization and visual analytics. That is, EDA can be seen as a part of the broader concept of data mining. From a more narrow perspective, the increased importance of visual examinations strengthens the link between EDA and information visualization, which has lately burgeoned into a broad field on its own. Finally, the third topic of discussion is visual analytics, the combination of information visualization and data mining.

4.1.1 Knowledge Discovery and Data Mining

Data mining is an interdisciplinary subfield of computer science, of which EDA is one out of many approaches. While being a topic that has lately attracted broad attention in academia, industry and media alike, the original definition of data mining is only a step of the broader concepts of knowledge discovery (KD) and knowledge discovery in databases (KDD). The definitions of these three terms, and their variations, are not seldom confused. This motivates a further look into them.

KD, and later called KDD, concerns the broad knowledge discovery process applied to large databases. More precisely, KD was first defined as "*the nontrivial extraction of implicit, previously unknown, and potentially useful information from data*" (Frawley et al. 1992). Later, Fayyad et al. (1996a) revised the definition of KD to the following definition of KDD: "*the nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data*". Hence, KDD concerns the entire knowledge extraction process, including how to store and access data, how to develop efficient and scalable algorithms for analyzing large datasets, how to visualize and interpret results, and how to model and support human-machine

interaction (Fayyad et al. 1996b). Data mining, on the other hand, is most often only one of the steps in the KDD process. When untangling the concepts of KDD and data mining, Fayyad et al. (1996a) define the latter as follows: “*applying data analysis and discovery algorithms that, under acceptable computational efficiency limitations, produce a particular enumeration of patterns over the data*”. They further note that as the enumeration involves a search in the space of patterns, which oftentimes is infinite, computational constraints may restrict the subspace that is feasible to be explored through data mining. Although some define data mining as the process of automatically finding interesting facts in data [see, e.g., Fekete et al. (2008)], the notion of data mining is not in this work restricted to automated pattern recognition, but includes also various types of interactive exploratory approaches as a support to the KDD process in general and data mining in particular.

Data Mining in a KDD Process

So far, we have discussed and untangled the concepts of data mining and KDD, while the general process of KDD and the precise role of data mining is yet to be discussed. Alas, there is no one process of KDD. Table 4.1 provides four versions of the KDD process, as well as one generic example. Following a blend of all KDD processes in the table, especially the especially the Cross Industry Standard Process for Data Mining (CRISP-DM) process (Shearer 2000) and the generic example (Kurgan and Musilek 2006), we will focus on a simplified KDD process that corresponds to that applied in this book. The upper part of Fig. 4.1 summarizes the herein used KDD process into the following six steps:

- (i) Domain understanding
- (ii) Data understanding
- (iii) Data preparation
- (iv) Data mining
- (v) Performance evaluation
- (vi) Knowledge consolidation and deployment.

The key objective of *Step 1* is to have the necessary knowledge about the application domain for proceeding with further analysis. Only after that, it also involves formulating the key objectives of the project from the perspective of all stakeholders, translating these general objectives into a KDD problem, and then drafting a broad and preliminary plan for achieving the objectives. As in the CRISP-DM process (Shearer 2000), this type of a plan is the basis of the entire KDD process. Data exploration in *Step 2* involves achieving an understanding of the underlying data, whereas data preparation in *Step 3* addresses the observed deficiencies and properties. The understanding may involve exploring the existence of outliers and missing values, as well as distributions of data. To address these concerns, data preparation may comprise transforming, cleaning, imputing and preprocessing data. An initial task is to select and collect data as per relevance, availability, quality, and other

Table 4.1 Examples of KDD processes

Steps	Fayyad et al. (1996c)	Cabena et al. (1998)	Anand and Buchner (1998)	CRISP-DM (Shearer 2000)	Generic (Kurgan and Musilek 2006)
1.	1. Developing an understanding of the application domain	1. Business objectives determination	1. Human resource identification	1. Business understanding	1. Application domain understanding
2.	2. Creating a target dataset	2. Data preparation	2. Problem specification 3. Data prospecting 4. Domain knowledge elicitation	2. Data understanding	2. Data understanding
3.	3. Data cleaning and preprocessing		5. Methodology identification	3. Data preparation	3. Data preparation and choice of data mining method
4.	4. Data reduction and projection		6. Data preprocessing		
5.	5. Choosing the data mining task				
6.	6. Choosing the data mining algorithm				
7.	7. Data mining	3. Data mining			
8.	8. Interpreting mined patterns	4. Domain knowledge elicitation	7. Pattern discovery	4. Modeling	4. Data mining
9.	9. Consolidating discovered knowledge	5. Assimilation of knowledge	8. Knowledge post-processing	5. Evaluation	5. Evaluation
			6. Deployment	6. Knowledge consolidation and deployment	6. Knowledge consolidation and deployment

Notes The table relates five versions of a KDD process to each other. It is an adapted version of the compilation in Kurgan and Musilek (2006)

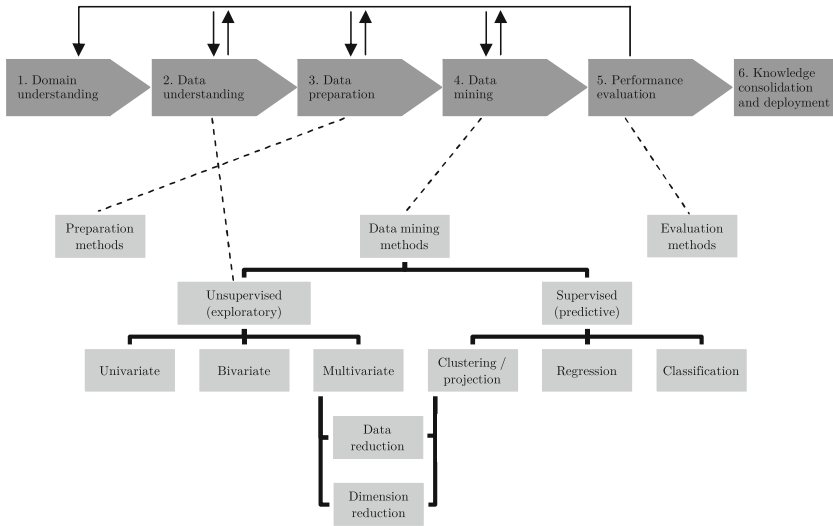


Fig. 4.1 KDD process and data mining techniques. *Notes* The figure represents the KDD process applied in this book, as well as a mapping of techniques for data mining, data preparation and performance evaluation to the process steps

domain-specific constraints and objectives. Another key task is to clean the selected data by identifying and correcting, replacing or removing data that are erroneous, irrelevant, incomplete or inaccurate. Further, it is also important to transform collected data and construct entirely new data, both as per the needs for the task at hand. Finally, one should not forget the oftentimes time consuming tasks of integrating and formatting data. There is some controversy concerning the order of *Steps 2 and 3*, as preparation is often needed before one can explore the data, whereas exploration affects preparation. Yet, it can mostly be assumed that the task of understanding has to logically precede that of preparation. Still, this is largely a detail of presentation as these two steps are most often iterated multiple times.

In *Step 4*, data mining techniques are applied to the selected data and parameterized for optimal performance. While data mining is only one step in the process, and the quality of work in *Steps 1–3* also significantly impact the results, data mining is the one step that has the largest influence on what the output or outcome is of the KDD process. Generally, multiple techniques are applied to the same problem for testing their suitability for the task at hand, including the stages of selecting a modeling technique, generating a test design, creating models and assessing the output of models [see, e.g., Shearer (2000)]. While also involving a ranking of models according to evaluation measures, as well as performance in fulfilling the objectives from the viewpoint of the domain in general and project in particular, the more thorough performance evaluation of the models is done in *Step 5*. This step involves a range of evaluation methods for determining the performance and robustness of the created models. The evaluation may not only relate to quantitative measures of performance,

such as accuracies (errors) and profits (costs) of classification tasks, but also to softer notions of achievement of the project objectives and domain requirements. Finally, *Step 6* involves a knowledge consolidation of the entire exercise and a deployment of the model for future use.

The lower part of Fig. 4.1 shows the methods that the KDD process makes use of and that support the overall process, such as in data preparation, understanding and mining, and performance evaluation. This is the following topic.

Methods in the KDD Process

The lower part of Fig. 4.1 points out the need for a variety of methods in the KDD process. The figure shows that data preparation in *Step 2* needs an own set of methods for preparing data into a format that lends to analysis. Likewise, performance evaluation in *Step 5* needs a wide range of methods for judging the quality of data mining models. In between these steps, the most important groups of methods are utilized for data understanding in *Step 2* and data mining in *Step 4*. These four separate groups of methods are discussed below.

First, in *Step 2*, methods for unsupervised exploration function as an aid in understanding data. Whereas methods in this category provide means for EDA in general, a focus at this stage is on methods that aid in exploring the existence of outliers and missing values, as well as distributions of data. Univariate and bivariate summary statistics (e.g., ranges, standard deviations and normality tests) and visual plots (e.g., box plots, histograms and time series) provide simple means for a range of descriptive assessments. Yet, having a multivariate viewpoint may be beneficial in the case of understanding structures in the high-dimensional space, such as exploring multivariate cluster structures and similarities in multivariate data. The former multivariate exploration may obviously be conducted with data reduction (or clustering) approaches and the latter with dimension reduction (or projection) approaches.

Second, data preparation involves the tasks of refining initial raw data such that they can be fed into the data mining methods. After selecting and collecting data needed, key tasks of preparation relate to cleaning, transforming and constructing data. A principal task is to clean the collected data by identifying, correcting, replacing or removing data that are erroneous, irrelevant, incomplete or inaccurate. For instance, missing values may be replaced with a wide range of imputation methods [e.g., the SOM (Cottrell and Letrémy 2005) and multiple imputation (Rubin 1987)] and identified outliers may be replaced or removed (e.g., modified boxplots, Winsorizing and other model-based methods like Chauvenet's criterion), both as per the needs for the task. Further, new data may be created by, for instance, translating symbolic fields to numerical data or deriving new variables based upon already collected data. Deriving new variables from already collected data may hence involve a wide variety of transformation methods, where some are simpler (e.g., levels, ratios, annual changes, logs and differencing) and others are more advanced (e.g., the X-12-ARIMA seasonal adjustment procedure and Hodrick-Prescott detrending). For some data mining methods, it is also of high importance to standardize or normalize

data with an appropriate pre-processing method (e.g., min-max, z-score, percentiles and sigmoids).

Third, the key group is the one that comprises data mining methods, where exploratory unsupervised and predictive supervised methods are split into various categories. While exploratory methods are divided as per the data (univariate, bivariate and multivariate), the predictive methods are divided as per the output and aim of methods (classification, regression and clustering/projection). It is hence obvious that there is an overlap between the two broad groups, in particular multivariate exploratory methods like data and dimension reduction and clustering and projection, which in essence can be seen as the same methods. The categories overlap in that not only can projection and clustering methods be used for prediction of unlabeled data, but also as they can be semi-supervised. Projection methods, such as MDS and its variants [see, e.g., Cox and Cox (2001)], map multidimensional data into a lower dimension. The MDS methods do not, however, reduce the amount of presented data. Clustering techniques attempt to find clusters in the data, and thus reduce the amount of data by enabling analysis of a smaller number of profiles or partitions.

Fourth, the evaluation of the modeling mainly relates to quantitative measures of performance, but also includes softer notions of achievement. The quantitative goodness-of-fit measures are most often chosen based upon the applied data mining methods. If the key aim is classification, then measures of classification performance are used, where one might want to account for imbalanced class size and misclassification costs, both on the level of classes and entities. That is, cost-sensitive evaluations focus on estimated profits or costs of a model. Likewise, if the aim is regression or time-series forecasting, one should choose the evaluation methods that best measure conducted errors (e.g., mean square and mean absolute error, mean absolute deviation or root mean squared error). Unsupervised multivariate methods, such as data and dimension reduction, focus on how well they can preserve the structures in original data. In data reduction, a common quality measure is the quantization error, whereas dimension reduction oftentimes uses pairwise distances between data as a measure of preservation of similarity relations. As Shearer (2000) points out, one should also qualitatively evaluate the adequacy of the overall process and whether there are any important factors or tasks missing. The quality assurance also includes controlling whether the model was correctly built and whether only attributes available for future use were utilized in the models. Likewise, a more general-level check covers the extent to which project objectives and domain requirements have been achieved.

4.1.2 Information Visualization

Information visualization has lately emerged as one of the key fields to support EDA. Mainly, it has its origin in the fields of human-computer interaction, computer science, graphics and visual design (Bederson and Shneiderman 2003). A more precise definition of information visualization is “*the use of computer-supported, interac-*

tive, visual representations of abstract data to amplify cognition” (Card et al. 1999), where the aim is to improve human understanding of the data with graphical presentations or graphics. Thereby, the aim resembles what we discussed as intelligence amplification (IA) in the introduction. Generally, tools for information visualization are mainly and best applied for EDA tasks, which oftentimes includes browsing a large space of information. The identification of situations where browsing is useful aids in determining when information visualization is of value. Lin (1997) lists browsing to be useful when:

- (i) there is a good underlying structure and when related items are located close to one another;
- (ii) users are unfamiliar with the contents of the collection;
- (iii) users have little understanding of the organization of a system and prefer to use a method of exploration with a low cognitive load;
- (iv) users have difficulty in articulating or verbalizing the specific information need; and
- (v) users search for information that is easier to recognize than describe.

Above, we see five situations when browsing information visualizations is useful, yet we still need to discuss the elements of information visualization in depth. The rest of this subsection focuses on three subtopics of information visualization: human perception and cognition, data graphics and visualization techniques.

Human Perception and Cognition

An essential part of visual communication relates to the attempt to match the design according to the capabilities and limits of the human information and visual system. The visual system comprises the human eye and brain and can be seen as an efficient parallel processor with advanced pattern recognition capabilities [see, e.g., Ware (2004)]. The focus of human perception is the understanding of sensory information, where the most important form is the visual perception. The final IA of information visualization can be viewed as a type of cognitive support. The mechanisms of cognitive support are, however, multiple. Hence, visualization tools should be targeted to exploit advantages of human perception.

Mostly, arguments about the properties and perception capabilities of the human visual system rely on two grounds: (i) information theory (Shannon and Weaver 1963), and (ii) psychological findings. *Information theory* states that the visual canal is best suited to carry information to the brain as it is the sense that has the largest bandwidth. Ware (2004) asserts that there are two main *psychological theories* for explaining how to use vision to perceive various features and shapes: preattentive processing theory (Triesman 1985) and gestalt theory (Koffa 1935). Prior to focused attention, preattentive processing theory relates to simple visual features that can be perceived rapidly and accurately and processed effectively at the low level of the visual system. Whereas more complex visual features require a much longer process of sequential scanning, preattentive processing is useful in information visualization

as it enables rapid dissemination of the most relevant visual queries through the use of suitable visual features, such as line orientation, line length or width, closure, curvature and color (Fekete et al. 2008). At a higher cognitive level, gestalt theory asserts that our brain and visual system follow a number of principles when attempting to interpret and comprehend visuals. Ware (2004) summarizes the principles as follows:

Proximity: Items close together are perceptually grouped together.

Similarity: Elements of similar form tend to be grouped together.

Continuity: Connected or continuous visual elements tend to be grouped.

Symmetry: Symmetrical elements are perceived as belonging together.

Closure: Closed contours tend to be seen as objects.

Relative size: Smaller components of a pattern tend to be perceived as objects.

The principles of gestalt theory can easily be related to some more practical concepts. For instance, most dimension reduction methods, when aiming at visualizing data, may be seen to relate to the proximity principle, as they locate data with high proximity close to each other, whereas others are pushed far away. Likewise, a time trajectory may be paired with continuity. More related to the cognition of visualizations, Fekete et al. (2008) explain that the core benefit of visuals is their functioning as a frame of reference or temporary storage for human cognitive processes. The authors assert that visuals are external cognition aids in that they augment human memory, and thus enable allocating a larger working set for thinking and analysis. In the above stated definition of information visualization by Card et al. (1999), visuals are presented as a means to “amplify cognition”. Following that definition, the authors also list a number of ways how well-perceived visuals could amplify cognition:

- (i) by increasing available memory and processing resources;
- (ii) by reducing the search for information;
- (iii) by enhancing the detection of patterns and enabling perceptual inference operations;
- (iv) by enabling and aiding the use of perceptual attention mechanisms for monitoring; and
- (v) by encoding the information in an interactive medium.

Examples of the *first* way to amplify cognition, the increase in available resources, are parallel perceptual or visual processing and offloading work from the cognitive system to the perceptual system (Larkin and Simon 1987). *Second*, visuals facilitate the search procedure by the provision of a large amount of data in a small space (i.e., high data density) (Tufte 1983) and by grouping information used together in general and information about one object in particular (Larkin and Simon 1987). *Third*, abstraction and aggregation aid in the detection of patterns and operations for perceptual inference (Card et al. 1991). *Fourth*, perceptual monitoring is enhanced, for instance, through the use of pop-out effects created by appearance or motion (Card et al. 1999). Likewise, Card et al. (1999) exemplify the *fifth* way to amplify

cognition, the use of a manipulable medium, by allowing the user to explore a wide range of parameter values to interactively explore properties of data.

Yet, matters concerning human perception and cognition also constitute a large set of issues that may hinder, disturb or generally negatively affect how visualizations are read. A key starting point is to take into account the deficiencies and limitations of human perception. Preattentive processing, for instance, becomes a deficiency if visuals are not designed properly. Patterns a user is supposed to identify quickly—or give visual but not conscious attention to—should hence be made distinct from the rest by using features that can be preattentively processed. Likewise, visual attention functions as a filter in that only one pattern is brought into working memory (Baddeley and Logie 1999). Hence, if provided with multiple patterns, we only see what we need or desire to see by tuning out other patterns. Ware (2005) also mentions the fact that humans process simple visual patterns serially at a rate of one every 40–50 ms, and a fixation lasts for about 100–300 ms, meaning that our visual system processes 2–6 objects within each fixation, before we move our eyes to visually attend to some other region. In addition, one important factor to account for is how perception of visuals is affected by properties of the human eye, such as acuities, contrast sensitivity, color vision, perception of shape or motion with colors, etc. Another aspect of crucial importance is obviously to pay regard to human perceptions of shapes in visuals, such as distances, sizes and forms. Cognitive deficiencies should also be accounted for when designing visuals, such as the limited amount of working memory. For instance, Haroz and Whitney (2012) show that the effectiveness of information visualizations is severely affected by the capacity limits of attention, not the least for detecting unexpected information. Hence, an understanding of the functioning of the human visual system aids in producing effective displays of information, where data are presented such that the patterns are likely to be correctly perceived.

Data Graphics

The literature on data graphics has its focus on the principles for visual representations of data. Herein, the focus is on the early, yet brilliant, work by Tufte (1983) and Bertin (1983). Their works, while being principles for graphics design, are to some extent also valid to overall computer-based visualizations. Tufte's set of principles are called a theory of data graphics, whereas Bertin's work is most often denoted a framework of the planar and retinal variables. However, rather than an exact theory, Tufte and Bertin provide a set of rules of thumb to follow.

The following overview is included to provide concrete guidelines, in addition to the above discussion of human perception and cognition. Herein, we will discuss only the key components of frameworks and theories by Bertin and Tufte. We start from Bertin's (1983) framework called the Properties of the Graphic System, which consists of two planar and six retinal variables. The two planar variables are the x and y dimensions of a visual, whereas the six retinal variables describe the following visual marks on the plane: size, value, texture, color, orientation and shape. The

eight variables can be categorized according to the following levels of organization, or so-called perceptual properties:

- (i) **Associative** (\equiv): If elements can be isolated as belonging to the same category, but still do not affect visibility of other variables and can be ignored with no effort.
- (ii) **Selective** (\neq): If elements can immediately and effortlessly be grouped into a category, and formed into families, differentiated by this variable, whereas the grouping cannot be ignored.
- (iii) **Ordered** (**O**): If elements can perceptually be ordinally ranked based upon one visually varying characteristic.
- (iv) **Quantitative** (**Q**): If the degree of variation between elements can perceptually be quantified based upon one visually varying characteristic.

When having an understanding of the four levels of organization, we can return to discussing Bertin's (1983) eight visual variables. Bertin describes the *plane*, and its two dimensions (x , y), as the richest variables. They fulfill the criteria for all levels of organization by being selective, associative, ordered and quantitative. The retinal variables, on the other hand, are always positioned on the plane, and can make use of three types of implantation: a point, line, or area. First, *size* is ordered, selective but not associative, and the only quantitative retinal variable. Second, *value* is the ratio of black to white on a surface, according to the perceived ratio of the observer, and is also sometimes called brightness. The usage of value in this case is close to the one in the HSV (hue, saturation and value), cylindrical-coordinate representation of points in an RGB (red, green and blue) color space. It is an ordered and selective retinal variable. Third, *texture* represents the scale of the constituent parts of a pattern, where variation in texture may occur through photographic reductions of a pattern of marks. That is, it may range from null texture with numerous but tiny elements that are not identifiable to large textures with only few marks. Texture as a retinal variable can be ordered and is both selective and associative. Fourth, variation may occur in *color*. The variation of two marks with the same value or brightness is thus more related to changes in hue of HSV. Color as a retinal variable is selective and associative, but not ordered. Fifth, the *orientation* variable enables variation in the angle between marks. In theory, this opens up an infinite set of alternatives of the available 360° , whereas Bertin suggests the use of four steps of orientation. The orientation variable is associative and selective only in the cases of points and lines, but has no direct interpretation of order. Finally, the sixth variable of *shape*, while being a retinal variable on its own, also partly incorporates aspects of size and orientation. It is associative, but neither selective nor ordered.

A complement to Bertin's framework is the Theory of Data Graphics by Tufte (1983), which consists of a large number of guidelines for designing data graphics. The two key, broad principles are *graphical excellence* and *graphical integrity*. In addition to these, Tufte provides two more focused principles: *data-ink maximization* and *data density maximization*.

Tufte (1983) defines *graphical excellence* as something that "gives to the viewer the greatest number of ideas in the shortest time with the least ink in the smallest

space”. The principle of graphical excellence summarizes a number of his guidelines that encourage graphical clarity, precision, and efficiency: (i) avoid distortions of what the data have to say; (ii) aid in thinking about the information rather than the design; (iii) encourage the eye to compare the data; (iv) make large data sets coherent; (v) present a large number of data in a small space; (vi) reveal data at multiple levels of detail ranging from a broad overview to fine detail; (vii) and closely integrate statistical and verbal descriptions of the data.

The second of Tufte’s (1983) principles, *graphical integrity*, relates to telling the truth about data. To follow this principle, Tufte provides six key guidelines: (i) visual representations of numbers should be directly proportional to the quantities which the visuals represent; (ii) clear and detailed labeling should be used to avoid ambiguity; (iii) show data variation, not design variation; (iv) deflate and standardize units when dealing with monetary values; (v) the number of dimensions depicted should not exceed the number of dimensions in data; and (vi) data should not be showed out of context. The aim of these principles is to avoid deception and misinterpretation.

Third, the principle of *data-ink maximization* proposes that data graphics should focus on the data, and nothing else. Hence, a good graphical representation focuses on data-ink maximization with minimum non-data-ink. The data-ink ratio is calculated by 1 minus the proportion of the graph that can be erased without loss of data information. Tufte (1983) puts forward the following five guidelines related to data ink: (i) above all else, show data; (ii) maximize the data-ink ratio; (iii) erase non-data-ink; (iv) erase redundant data-ink; and (v) revise and edit.

The fourth of Tufte’s (1983) principles, *data density maximization*, relates to the share of the area of the graphic dedicated to showing the data. For too low densities, Tufte suggests to either reduce the size of the graphic (shrink principle) or the use of a table. In particular, he claims that graphs can oftentimes be shrunk in size without losing legibility or information. In terms of concrete design, he proposes the small multiples, a design for showing varying data onto a series of the same small graph repeated in one visual.

Bertin’s and Tufte’s principles provide a guiding set of rules of thumb to follow when spanning the space of two-dimensional visualizations. Yet, visualizations, not the least interactive visualizations, go beyond a static two-dimensional space by including additional visual variables, such as depth and time. The next part discusses a range of visualization techniques and tools, where interaction becomes essential.

Visualization Techniques

The literature has provided a long list of techniques for creating visual representations. Herein, we will mainly focus on a rough overview, as well as a brief and simple taxonomy, of methods, rather than a detailed survey of methods. Obviously, a key issue of information visualization is what formats and features the methods will help to organize and visualize, as well as how that relates to the use of human visual capabilities. Following Zhang et al. (2012), the techniques supporting information visualization can be divided into two groups: graphical representations of data and

interaction techniques. The former group refers to the visual form in which the data or model is displayed, such as standard visualization techniques like bar and line charts. Yet, visualizations may often refer to manipulable graphical displays of data. The latter group of interaction techniques refers to how the user can interact with or manipulate the graphical displays, such as zooming or panning. These oftentimes have their basis in one or more graphical displays such that they enable more freedom and flexibility to explore the data.

From the viewpoint of the underlying data, rather than the formats of visual displays, Zhang et al. (2012) categorize visualization techniques into four groups: numerical data, textual data, geo-related data and network data. First, *numerical data* can be visualized by a vast number of approaches, such as standard visualization techniques like bar and pie charts and scatter plots. These focus most often on the visualization of low-dimensional numerical data. On the other hand, visualization techniques like parallel coordinates, heatmaps and scatter plot matrices provide means to display data with higher dimensionality. Second, visualization of *textual data* is a new, growing field. Recent techniques include word cloud (Kaser and Lemire 2007) and theme river (Havre et al. 2000), for instance. Likewise, the availability of the third type of data, *geo-tagged data*, has caused a soar in the demand for geo-spatial visualizations. Geo-related univariate or multivariate information is oftentimes projected into conventional two-dimensional and three-dimensional spaces. Fourth, graph visualizations provide means for displaying patterns in *network data* with relationships (i.e., edges) between entities (i.e., nodes). They most often consist of a technique for positioning, such as force-based algorithms, as well as coloring or thickness of edges to display the size of a relationship. Graph or network visualizations have been increasingly applied in a wide range of emerging fields like social and biological network analysis, not to mention financial network analysis.

A categorization of visualization techniques as per the types of data does not, however, differentiate all possibilities of techniques. While being some years old, Keim and Kriegel (1996) groups visualization techniques into five categories: geometric, icon-based, pixel-oriented, hierarchical, and graph-based techniques. First, *geometric techniques* provide means for visualization of geometric transformations and projections of data. Examples of the methods are scatterplot-matrices, parallel-coordinate plots and projection methods. Second, *icon-based techniques*, as already the name states, visualize data as features of icons. The methods include, for instance, Chernoff-faces and stick figures, of which the former visualize multidimensional data using the properties of a face icon and the latter use stick figures. Third, *pixel-oriented techniques* map each attribute value to a colored pixel and present attribute values belonging to each attribute in separate subwindows. For instance, query-independent techniques arrange data from top-down in a column-by-column fashion or left to right in a line-by-line fashion, while query-dependent techniques visualize data in the context of a specific user query. Four, *hierarchical techniques* provide means to illustrate hierarchical structures in data. Most often, hierarchical methods focus on dividing an n -dimensional attribute space by “stacking” two-dimensional subspaces into each other. Finally, the fifth category, *graph-based techniques*, focus on the visualization of large graphs, or networks, to illustrate the properties of the network, as was above

discussed. In addition, Keim and Kriegel also illustrate the existence of a wide range of hybrids that make use of multiple categories.

The above discussion obviously illustrates the importance of choosing a suitable type of display format, given the data and the task at hand. Yet, it not only illustrates, but also guides in the choice. Following the above paragraphs, we can use the data and display categories. The first factor to define the nature of the chosen visualization technique is the properties of the data, such as the form of data, dimensionality of data, data structures and size of data. The second factor to determine is the expected output and purpose of use, where the variation of purposes is large, such as predictive versus exploratory, temporal versus cross-sectional, and univariate versus multivariate analysis and similarity versus dissimilarity matching, as well as other purposes related to a focus on geo-spatial visualizations and network relationships, for instance. While there is no one way to choose the correct technique, considering the two dimensions of data and display, as well as other restrictions, demands and needs for the task, provides an adequate basis.

Given a technique, a critical factor of information visualization is, however, the possibility to interact with the visuals. Like the KDD process in the entire knowledge extraction process, a common guideline for interactions with visualizations is the visual information seeking mantra (Shneiderman 1996): “*Overview first, zoom and filter, then details-on-demand*”. Whereas Shneiderman (1996) characterizes the mantra with seven abstract tasks, we focus only on the following four explicitly mentioned ones: First, a user should gain an *overview* of the entire collection through a high-level representation. Second, users should have the possibility to *zoom* in on a portion of items that are of particular interest. Third, there should exist the possibility to *filter* out or to eliminate uninteresting and unwanted items, such as allowing users to specify which items to display. Fourth, the user should have the option to select an item or group of items to get further *details-on-demand*, such as clicking a group or individual items.

This provides a starting point to data visualization and user interaction, but does still not address the role of analytical or data mining techniques in visualization. The next step is to combine graphical representations of data and interaction techniques with analytical methods.

4.1.3 Visual Analytics

By adding data mining to the ingredients of information visualization, we end up with the original definition of visual analytics (Thomas and Cook 2005): “*the science of analytical reasoning facilitated by interactive visual interfaces*”. Hence, the field of visual analytics has strong roots in information visualization. Likewise, visual analytics is also strongly related to the KDD process. The term visual data mining descends from the integration of the user in the KDD process through visualization techniques and interaction capabilities [see, e.g., Keim (2001)]. This has taken visual analytics to be applied in areas with challenging problems that were unsolvable using

standalone automatic or visual analysis [see, e.g., Keim et al. (2009)]. In particular, automated analysis enables scaling to larger and more challenging tasks, whereas visualizations may be used to effectively communicate the outcome to the user in particular or a broad audience in general.

Since we derive visual analytics from three above presented concepts—graphical representations of data, interaction techniques and data mining techniques—there is no need to repeat the discussion of each component. Yet, the above presented information seeking mantra only mentions visualizations in the KDD process, while not integrating the two concepts. Keim et al. (2006) propose combining the KDD process and information seeking mantra for a visual analytics mantra: “*Analyze first, show the important, zoom, filter and analyze further, details on demand*”. The authors exemplify the visual analytics mantra with analysis of large network security data. As graphical representations of raw data is infeasible and seldom reveals deep insights, the data need to first be analyzed, such as computing changes and intrusion detection analysis. Then, the outcome of the automated analysis is visualized. Out of the displayed results, the user filters out and zooms in to choose a suspicious subset of all recorded intrusion incidents for further, more careful analysis. Thus, the mantra involves automated analysis before and after the use of interactive visual representations. Following the mantra, an adapted version of the visual analytics process in Keim et al. (2010) is presented in Fig. 4.2. The key steps in the process are data preparation, visual and automatic analysis, and knowledge consolidation. Whereas the step of data preprocessing and transformations is similar to that discussed in the previous section, the user selects in the following step between visual or automatic analysis methods. The user might prefer to start from whichever of the two tasks, but it is likely that several iterations of data visualization and interaction, and automatic analysis is needed. Finally, after alternating between visual and automatic methods, the thus far gained knowledge is not only gathered, but also transferred through a feedback loop to support future analysis.

Visual analytics in general and the visual analytics mantra in particular link to the core of this book, the use of data and dimension reduction methods to support human cognition. While clustering and data reduction methods provide overviews or summarizations of data by compressing information into fewer profiles, projection and dimension reduction methods lead to the visualization of high-dimensional spaces in a low-dimensional mapping by preserving similarity structures. Following the above four categories of data, they lead to analysis of numerical, textual and network data, given adequate preprocessing. Likewise, they may also be paired with geo-spatial visualizations. The sequel of this chapter focuses on aims of and methods for data and dimension reduction, where the underlying data are assumed to be numerical.

4.2 Dimension Reduction

Before defining the concept of dimension reduction, it is worth to note that it goes by multiple other names. Some call it a projection, vector projection or projection pursuit, and others call it a mapping. Whereas the key aim of dimension reduction is to

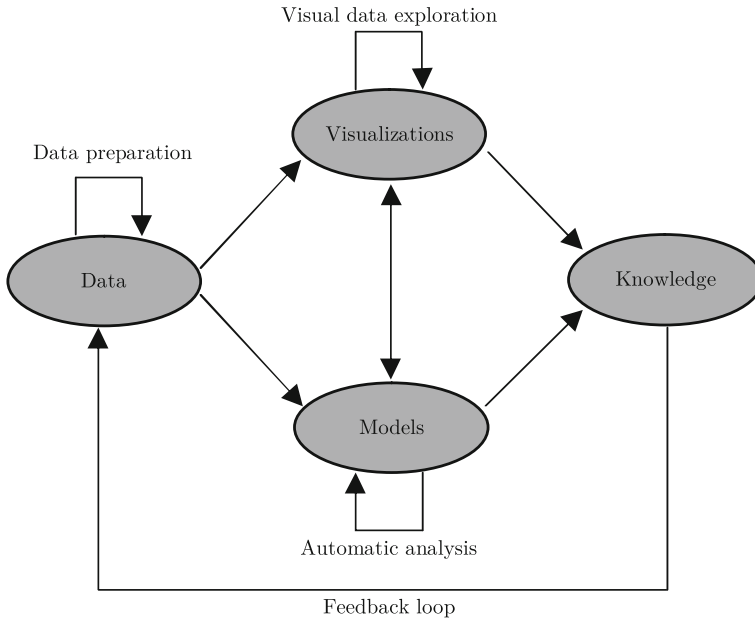


Fig. 4.2 Visual analytics process. *Notes* The figure represents the visual analytics process. The figure is adapted from Keim et al. (2010)

provide a low-dimensional overview of similarity relations in data, slight variation in purposes of use lead to different preferences of preserved structures. In this section, we will first broadly discuss variations in aims and purposes of use with respect to dimension reduction and then review first-generation and second-generation methods, as well as position them in a taxonomy.

4.2.1 Aims of Dimension Reduction

Reducing dimensionality may be motivated by a large number of reasons. For instance, Zhang and Liu (2005) relate the desire of dimension reduction to enhancing the understanding of data, reducing the complexity of the system, and avoiding the curse of dimensionality. Yet, different purposes of use and applications have different aims and preferences of the preserved properties. Without going into the details of how to preserve structures, and what are the similarity relations of crucial interest in various tasks, we can categorize dimension reductions by relating them to three broad aims [see, e.g., Lee and Verleysen (2007), Zhang and Liu (2005)]: (i) visualization and exploration, (ii) regression, and (iii) classification. The first aim of *visualization* relates to embedding high-dimensional data into a low-dimensional space by preserving their intrinsic dimensions. The second aim is *regression*, in which the focus is on reducing the dimensionality of the predictor vector (i.e., explanatory variables),

but at the same time minimizing the loss in inferences about the predicted variable (i.e., explained variable). Third, the aim may also relate to *classification*, in which case the goal is to find a low-dimensional space with the minimum classification error.

Another approach to categorization is to view the problem of dimension reduction as supervised and unsupervised [see, e.g., Zhang and Liu (2005), Gisbrecht et al. (2012)]. Whereas the standard approach to dimension reduction is an unsupervised search for a low-dimensional representation of similarity relations in high-dimensional data, one may also opt to supervise the mapping by integrating class information. The supervised version, also called a discriminative dimension reduction, may be thought of as having two parts in its cost function, where one consists of the preserved structures in the observed variables and the other of the preserved structures (e.g., distances) of the labels. The use of labels to steer the dimension reduction can oftentimes be a useful addition when interpreting the low-dimensional output, as users oftentimes have a direct understanding of the classes. In addition to amplifying the understanding of the underlying class structure, a central task of supervised dimension reductions is also to aid in visually classifying data and communicating the results of a classification.

This boils into the following question: *What are the tasks that the methods perform and the functionalities that are needed?* Lee and Verleysen (2007) describe that the key functionalities of dimension reductions are to be able to: (i) estimate the number of latent variables, (ii) reduce dimensionality by embedding data, and/or (iii) recover latent variables by embedding data. First, to judge the number of *latent variables*, one needs to perform an estimation of the intrinsic dimensionality. Yet, only few methods provide means for such an estimation. For instance, with the two latter of the above aims (i.e., regression and classification), one might be interested in reducing dimensionality only up to a point that captures variations of the latent variables, whereas the ones capturing noise and other imperfections are disregarded. Second, a natural next step is to *re-embed the high-dimensional data* into a better filled lower dimension. The aims may be to achieve a compact representation and/or to facilitate subsequent processing of data, where the former aids in visualizing the data and the latter supports a further data compression (or data reduction). Third, the task of *latent variable separation* also involves means for recovering the variables, in order to fulfill an aim beyond only a reduction of the dimensionality. One intuitive approach to recovering latent variables is to model the observed variables as linear combinations of the latent ones. However, it is worth noting that the same method seldom performs the second and third tasks of reducing dimensionality and separating latent variables. In the following, while presenting a broad palette of methods, the focus herein is on embedding data into a lower dimension to support data compression and visualization.

4.2.2 An Overview of Methods

The first dimension reduction methods date back to the early 20th century. However, only since the 1990s has there been a significant soar in the number of developed methods. This can be used as a cutting point for dividing the methods into *first-generation* and *second-generation methods*, are then described in a tree-structured *taxonomy*.

First-Generation Methods

The first generation consists of the well-known classical methods that are still broadly used and accepted in a wide range of domains. Drawing upon the first introduced, but still commonly used, variance-preserving Principal Component Analysis (PCA) (Pearson 1901), an entire family of distance-preserving MDS-based methods have been developed. The MDS counterpart to PCA, classical metric MDS, which attempts to preserve pairwise distances, was proposed by Young and Householder (1938), Torgerson (1952). Non-linear versions are the first introduced non-metric MDS by Shepard (1962), Kruskal (1964) and the later developed Sammon's (1969) mapping. As said, the key aim of these methods is to project high-dimensional data x_j (i.e., the input space) to a two-dimensional data vector y_j (i.e., the output space) by preserving distances. To start with, we look at the functioning of the distance-preserving counterpart of PCA, classical metric MDS. Let the distance in the input space between x_j and x_h be denoted $d_x(j, h)$ and the distance in the output space between y_j and y_h be denoted $d_y(j, h)$. This gives us the objective function of metric MDS:

$$E_{MDS} = \sum_{j \neq h} (d_x(j, h) - d_y(j, h))^2. \quad (4.1)$$

Due to the simple linear form of metric MDS, we also explore the functioning of a non-linear MDS-based method, Sammon's (1969) mapping. It is an MDS method in that it also attempts to preserve pairwise distances between data but differs by focusing on local distances relative to larger ones. The square-error objective function for Sammon's mapping is

$$E_{SAM} = \frac{1}{\sum_{j \neq h} d_x(j, h)} \sum_{j \neq h} \frac{(d_x(j, h) - d_y(j, h))^2}{d_x(j, h)}, \quad (4.2)$$

and shows that it considers all pairs (j, h) normalized by the input space distance $d_x(j, h)$ and weighted with $1/d_x(j, h)$. The objective functions of MDS-based methods are most often optimized with an iterative steepest-descent process.

The topology-preserving family of methods was launched through the introduction of the SOM (Kohonen 1982). The SOM differs by reducing both dimensions and

data through a neighborhood-preserving vector quantification. A further discussion is given separately in Sect. 4.4, as it falls into both data and dimension reduction methods.

Second-Generation Methods

The second generation is a less homogeneous group of methods ranging from so-called spectral techniques to graph embedding. Due to the large number of methods, we will discuss only broadly some of the most recent methods, in order to later present their location in a taxonomy. A soar in developed methods at the turn of the century lead to several innovative approaches, such as Curvilinear Component Analysis (CCA) and Curvilinear Distance Analysis (CDA) (Demartines and H erault 1997), Local MDS (LMDS) (Venna and Kaski 2006; Chen and Buja 2009), Generative Topographic Mapping (GTM) (Bishop et al. 1998), Locally Linear Embedding (LLE) (Roweis and Saul 2000), Isomap (Tenenbaum et al. 2000), Laplacian Eigenmaps (LE) (Belkin and Niyogi 2001) and Maximum Variance Unfolding (MVU) (Weinberger and Saul 2005). Some more recent methods are, for instance, t-distributed Stochastic Neighbor Embedding (t-SNE) (van der Maaten and Hinton 2008) and Exploration Observation Machine (XOM) (Wism uller 2009). These methods, while differing along multiple properties, generally aim at the above three tasks of non-linear dimension reduction. A detailed mathematical treatment of them is, however, outside the scope of this book.

A Taxonomy

In addition to two generations, dimension reduction methods can also be illustrated in a tree-structured taxonomy. The tree-structure in Fig. 4.3 is a non-exhaustive taxonomy of dimension reduction methods based upon that in Lee and Verleysen (2007, p. 234). While the focus herein is on methods based upon geometrical concepts, there exists also other methods, such as the Auto Associative Neural Networks (AANNs). The tree structure ends with some exemplifying methods, where first-generation methods are differentiated from second-generation methods through a gray background. Methods can roughly be divided into those aiming at distance and topology preservation. The distance-preserving methods can still be divided into different distances, such as spatial (e.g., PCA, MDS, Sammon's mapping and CCA), graph (e.g., Isomap and CDA) and other (e.g., MVU). Topology-preserving methods can be divided into those with a predefined grid shape (e.g., SOM, GTM and XOM) and those without (e.g., LLE, LE and t-SNE). It is worth considering that the example methods in the taxonomy are only a subset consisting of the most commonly used ones.

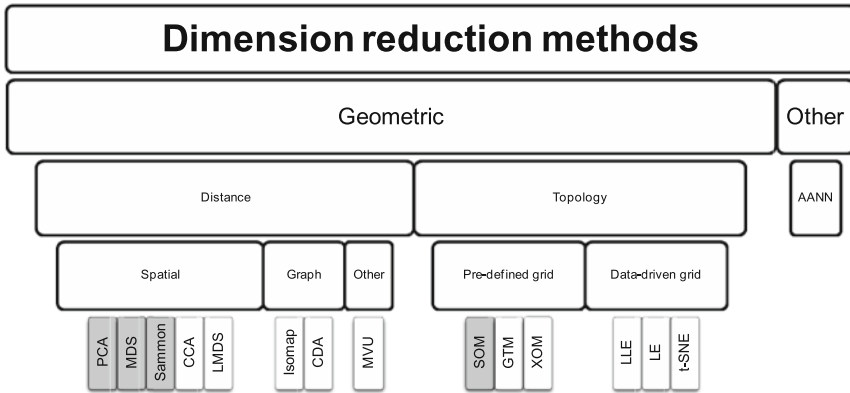


Fig. 4.3 A taxonomy of dimension reduction methods. *Notes* The figure represents a non-exhaustive taxonomy of dimension reduction methods adapted from Lee and Verleysen (2007, p. 234). The lowest level associates methods to their families, where a *gray background* indicates first-generation methods and *white* second generation. *Acronyms* Auto Associative Neural Network (AANN), Principal Component Analysis (PCA), Local MDS (LMDS), Multidimensional Scaling (MDS), Curvilinear Component Analysis (CCA), Curvilinear Distance Analysis (CDA), Maximum Variance Unfolding (MVU), Self-Organizing Map (SOM), Generative Topographic Mapping (GTM), Exploration Observation Machine (XOM), Locally Linear Embedding (LLE), Laplacian Eigenmaps (LE) and t-distributed Stochastic Neighbor Embedding (t-SNE)

4.3 Data Reduction

Data reduction, as also dimension reduction, goes by multiple names, such as data compression, data clustering and cluster analysis. In today’s information rich world with vast amounts of available unlabeled data it is not enough to decrease dimensionality, oftentimes one also needs to focus on reducing the number of data. Moreover, unsupervised approaches are frequently the only feasible approach to form an understanding of the data, not the least when they are unlabeled. Clustering methods provide means for exploring tendencies and structures in data by reducing data to fewer partitions, mostly with the aim of having small intra-cluster distances and/or large inter-cluster distances. In this section, we first discuss how overall aims of data reduction may differ and then provide a brief classification of methods and introduction to the ones essential for this book.

4.3.1 Aims of Data Reduction

A key aim of data reduction is the exploratory task of organizing data into sensible groupings to find structure in them. Whereas data reduction methods can be used for dividing data into homogeneous groups, however those are defined, aims and

objectives of these methods may be inherently different. Jain (2010) presents three overall aims of data reduction: (i) to explore the underlying structure, (ii) to find a natural classification, and (iii) to perform a compression of data. The first aim of *exploration* concerns gaining insight into data, generating hypotheses, detecting anomalies, and identifying salient features in data. The second type of aim involves the attempt to derive a natural *classification* that identifies the degree of similarity among objects. Finally, the third aim is to *compress* data such that they can easily be organized and summarized, as well as utilized as an input for additional analysis, through the use of a smaller set of representative cluster prototypes.

It is also worth noting that data reduction, while most often being unsupervised in that it does not use class information, has lately also been used in a semi-supervised manner [see, e.g., Chapelle et al. (2006)]. There are three key arguments for having a semi-supervised approach to data reduction (Chapelle et al. 2006; Jain 2010): (i) sometimes only a small portion of class information is available to the user, in which case one can let it partly guide the process through limited supervision; (ii) the available class information may also be too far from an ideal target variable for pure supervised learning to be feasible; and (iii) the user might possess pair-wise must-link and cannot-link constraints that can be used to guide two objects to be or not to be assigned into the same cluster.

While having three overall aims, Jain (2010) notes that data reduction is prevalent in any type of discipline involving the analysis of high-dimensional data. As there is no exhaustive list of all scientific fields and application areas utilizing some form of data reduction, Jain et al. (1999), Jain (2010) provide a number of examples: (i) image segmentation to facilitate computer vision, (ii) clustering views of two-dimensional and three-dimensional objects to aid in object and character recognition, (iii) clustering of text documents to automatically provide segments and hierarchies and improve efficiency, (iv) customer segmentation to aid in marketing campaigns, (v) to group genome-wide expression data to arrange genes according to similarity of gene expression patterns in biology and (vi) overall data mining to facilitate predictive modeling, exploratory segmentation, and visualization of large databases. To support the tasks in this book, the focus is on compressing data to support their organization, summarization and visualization, and the final category of applications.

4.3.2 An Overview of Methods

Although there is no common taxonomy of data reduction methods, several properties can be used for differentiating between methods: soft versus hard clustering, hierarchical versus non-hierarchical methods and monothetic versus polythetic goals, for instance. While soft clustering reduces data by assigning them to each cluster to a certain degree, hard clustering either assigns data to a cluster or not. Fuzzy c-means (FCM) clustering exemplifies the difference by being a soft counterpart of the classical k -means (or c -means to be consistent) clustering algorithm. Hierarchical methods [e.g., Ward's (1963) method] produce a taxonomy of cluster struc-

tures, in which small child clusters are also nested within larger parent clusters, and may be divided into agglomerative (bottom-up) and divisive (top-down) approaches. Non-hierarchical methods approach data reduction from numerous different viewpoints, and may roughly be divided into centroid-based [e.g., k -means clustering (MacQueen 1967) and Vector Quantization (VQ) (Linde et al. 1980)], distribution-based [e.g., Expectation-maximization algorithm (Dempster et al. 1977)] and density-based clustering [e.g., DBSCAN (Ester et al. 1996)]. The SOM may also be seen as a spatially constrained form of centroid-based clustering. The differences between monothetic versus polythetic methods relate mainly to hierarchical clustering, where the former uses the inputs one by one and the latter all the inputs at once. To illustrate the above described differences, the below discussion will focus on the functioning of a number of classical methods in more detail.

Centroid-Based Clustering

This part introduces two centroid-based clustering methods: VQ (Linde et al. 1980) and k -means clustering (MacQueen 1967). They can be seen as counterparts of the SOM, where the former relates to the sequential and the latter to the batch SOMs. VQ attempts to model the probability density functions in data x_j by reference vectors m_i (where $i = 1, 2, \dots, M$). It uses $\min(\|x_j - m_b\|)$ for finding the best-matching unit (BMU) m_b for x_j , and then updates sequentially only the BMU towards the data vector. Hence, it attempts to minimize the standard squared error function, or quantization error:

$$J_{VQ} = \sum_{j=1}^N d_x(j, b(j))^2, \quad (4.3)$$

where $d_x(j, b)$ is the input space distance between the data x_j and reference vector $m_{b(j)}$ and $b(j)$ denotes that b is the BMU of data j .

K -means is a similar least-square partitioning algorithm that pairs each data x_j to a cluster k (where $k = 1, 2, \dots, C$) and then updates the centroids c_k to averages of all attracted data. Thus, the aim is again to minimize the squared error function:

$$J_{km} = \sum_{j=1}^N \sum_{k=1}^C d_x(j, k)^2, \quad (4.4)$$

where $d_x(j, k)$ is the input space distance between the data x_j and cluster centroid c_k .

Fuzzy Clustering

We illustrate the functioning of the FCM algorithm, developed by Dunn (1973) and improved by Bezdek (1981), that assigns a degree of membership of each data in each of the clusters. The FCM algorithm implements an objective function-based fuzzy clustering method. The objective function J_θ is defined as the weighted sum of the Euclidean distances between each data and each cluster center, where the weights are the degree of memberships of each data in each cluster, and constrained by the probabilistic requirement that the sum of memberships of each point equals 1:

$$J_\theta = \sum_{j=1}^N \sum_{k=1}^C u_{jk}^\theta \|x_j - c_k\|^2, \quad \sum_{k=1}^C u_{jk} = 1, \quad (4.5)$$

where $\theta \in (1, \infty)$ is the fuzzy exponent, u_{jk} is the degree of membership of data x_j (where $j = 1, 2, \dots, N$) in the cluster center c_k (where $k = 1, 2, \dots, C$, and $1 < C < N$), and $\|x_j - c_k\|^2$ is the squared Euclidean distance between x_j and c_k . It operates through an iterative optimization of J_θ by updating the membership degree u_{jk} :

$$u_{jk} = 1 / \left(\sum_{s=1}^C \left[\frac{\|x_j - c_k\|}{\|x_j - c_s\|} \right]^{\frac{2}{\theta-1}} \right), \quad (4.6)$$

where s are the iteration steps, and by updating the cluster centers c_k :

$$c_k = \left[\sum_{j=1}^N u_{jk}^\theta x_j \right] / \left[\sum_{j=1}^N u_{jk}^\theta \right], \quad (4.7)$$

The algorithm proceeds as follows. First, the cluster centers are initialized randomly. Thereafter, each data x_j is assigned a membership grade u_{jk} in each cluster k . Then the so-called Picard iteration through Eqs. (4.6) and (4.7) is run to adjust the cluster centers c_k and the membership values u_{jk} . The algorithm stops when the minimum amount of improvement between two consecutive iterations is less than a small positive number ε or after a specified number of iterations.

Hierarchical Clustering

The third type of data reduction is hierarchical clustering. The following Ward's (1963) criterion is used as a basis for agglomerating clusters with the shortest distance:

$$d_{kl} = \frac{n_k n_l}{n_k + n_l} d_c(k, l)^2, \quad (4.8)$$

where k and l represent clusters, n_k and n_l the cardinality of clusters k and l , and $d_c(k, l)^2$ the squared Euclidean distance between the cluster centers of clusters k and l . When clusters k and l are merged to cluster h , the cardinality n_h is the sum of n_k and n_l and the centroid c_h the mean of c_k and c_l weighted by n_k and n_l . Hence, this specification accounts for cluster size. A particularly advantageous feature of hierarchical methods is that agglomeration can be restricted to some specific property of the underlying relations between clusters. For instance, the distance between non-adjacent clusters can be set to infinite, where adjacency needs to be defined (e.g., neighborhoods of data or clusters). Again, clusters can be said to agglomerate as to minimize the Euclidean distance to the centroids, or the squared error function. The algorithm starts with each data as its own cluster and merges units for all possible numbers of clusters using the minimum Ward distance (1, 2, ..., N).

4.4 The Self-Organizing Map (SOM)

The SOM (Kohonen 1982, 2001) is an inherently different, yet not unique, method in that it performs a simultaneous data and dimension reduction. It differs from non-linear projection techniques like multidimensional scaling by attempting to preserve the neighborhood relations in a data space Ω on a k -dimensional array of units (represented by reference vectors m_i) instead of attempting to preserve absolute distances in a continuous space. On the other hand, it differs from standard VQ by also attempting neighborhood preservation of the m_i . The VQ capability of the SOM performs this data reduction into mean profiles (i.e., units m_i). It models from the continuous space Ω , with a probability density function $p(x)$, to the grid of units, whose location depend on the neighborhood structure of the data Ω .

There exists two commonly used versions of the basic SOM algorithm: the sequential and the batch SOM. In this book, the batch training algorithm is employed, and thus data are processed simultaneously instead of in sequences. Important advantages of the batch algorithm are the reduction of computational cost and reproducible results. Reproducibility is, given the same initialization, independent of the order of data. Before training, initial values are assigned to the reference vectors (e.g., random, sample or linear initializations). Following the initialization, the batch training algorithm operates a specified number of iterations t (where $t = 1, 2, \dots, T$) in two steps. In the first step, each input data vector x_j is assigned to the BMUs m_b :

$$d_x(j, b) = \min_i d_x(j, i), \quad (4.9)$$

where $d_x(j, b)$ is the input space distance between data x_j and reference vector m_b (i.e., BMU) and $d_x(j, i)$ is the input space distance between data x_j and each reference vector m_i . Hence, data are projected to an equidimensional reference vector m_b , not

a two-dimensional vector as in MDS. In the second step, each reference vector m_i (where $i = 1, 2, \dots, M$) is adjusted using the batch update formula:

$$m_i(t+1) = \frac{\sum_{j=1}^N h_{ib(j)}(t)x_j}{\sum_{j=1}^N h_{ib(j)}(t)} \quad (4.10)$$

where index j indicates the input data vectors that belong to unit b , N is the number of the data vectors, and $h_{ib(j)}$ is some specified neighborhood function. In comparison to the update formula of the k -means algorithm in Eq. (4.4), the batch update of the SOM can be seen as a spatially ($h_{ib(j)}$) constrained version.

Mathematical treatment of the SOM has, however, shown to be difficult. Despite an extensive discussion of the form and existence of an objective function, the literature has still not provided one for the general case [see, e.g., Yin (2008)]. It has, however, been noted that a decomposed distortion measure illustrates the learning of the SOM [a discrete form with a fixed neighborhood of that suggested in Lampinen and Oja (1992)]:

$$E_{SOM} = \sum_{j=1}^N \sum_{i=1}^M d_x(j, i)^2 h_{ib} d_x(i, b), \quad (4.11)$$

where $d_x(j, i)$ is the input space distance between data x_j and reference vector m_i and $d_x(i, b)$ is the input space distance between reference vectors m_i and m_b .

4.4.1 Parametrizing the SOM

Setting the parameters, or parametrizing, a SOM involves a number of choices by the user.¹ While Kohonen (2001) has noted that the selection of all parameters is not

¹ There are several software implementations of the SOM. The seminal packages—SOM_PAK, SOM Toolbox for Matlab, Nenet, etc—are not regularly updated or adapted to their environment. Out of the newer implementations, Viscovery SOMine provides the needed means for interactive exploratory analysis. The most recent addition to the list of implementations is the interactive, web-based implementation provided by infolytika (<http://risklab.fi/demo/macropu/>). For a description, see Sarlin (2014a). For a practical discussion of SOM software and an early version of the implementation in Viscovery SOMine, see Deboeck (1998a, b). See also Moehrmann et al. (2011), for a comparison of SOM implementations. The first analyses of this book were performed in the Viscovery SOMine 5.1 package due to its easily interpretable visual representation and interaction features, not the least when introducing it to practitioners in general and policymakers in particular. Recently, the packages available in the statistical computing environment R have significantly improved, in particular regarding the visualization of SOM outputs. Thus, the final parts of the research in this book, including the figures, have been produced in R. Moreover, the above mentioned interface by infolytika provides an interactive implementation of the R-based models.

crucial if map size is small, it is needless to say that Kohonen does not generally overlook the importance of finding an adequate specification, as it indeed impacts the final result. The choices are defined to be on three different levels: (i) architecture, and (ii) internal and (iii) external specifications. The framework put forward herein draws upon the discussion in Kohonen (2001), Vesanto et al. (2000). First, one needs to make decisions related to the *architecture* of the SOM, that is, the form of the array of units created before training. In practice, the array may have one to three dimensions, of which the most common choice is the two-dimensional array. Moreover, the array is associated with a lattice, where hexagonal and rectangular forms are the most often used. In a rectangular lattice, a unit has four neighbors and in the hexagonal six. While most often being two-dimensional, the lattice shape may also vary from the standard sheet to toroids and cylinders, for instance.

Second, one has to decide upon *internal* specifications of the SOM algorithm. Over the years, a large number of variations to the specifications of the standard SOM have been provided. To start with, one needs to choose the initial values for the reference vectors using, for instance, random, sample or linear initialization. One commonly modified parameter is the neighborhood function $h_{ib(j)}$ that could take the form of a bubble, Gaussian, cut Gaussian and Epanechicov, for instance. Whichever function is chosen to be used, it is also common to implement the neighborhood function to be decreasing over training iterations as per a specified scheme.

Third, the *external* parameters, which most often are specified outside the SOM machinery, are defined as the true free parameters. The neighborhood function, whichever form it takes, commonly has a radius of the neighborhood parameter to be specified. The user also has to decide the number of training iterations. Finally, one also has to decide upon the number of units and the map shape (ratio of X and Y dimensions).

At this point of the book, we only discuss the choices regarding architecture and internal specifications, whereas external parameters outside the SOM machinery have to be specified during each training phase. Kohonen (2001), Vesanto et al. (2000) provide a range of solutions, hints and tips, not to say rules of thumbs, related to the specifications when designing and training a SOM. First, specifications relating to the *architecture* are set as follows. The use of a hexagonal lattice is not only common practice for its visual appeal, but also advisable for six neighbors of a unit being at the same distance, rather than only four in the case of a rectangular lattice. For the purpose of this book, the output of the SOM is chosen to be a two-dimensional sheet. The rationale for not using a one-dimensional array is to better represent general detail, particularly differences within clusters, whereas a three-dimensional map, while adding a further dimension, impairs the interpretability of data visualizations, not the least visualizations displayed on static paper. Kohonen's (2001) general suggestion is to set the shape of the lattice to correspond to the shape of the data manifold. Given a two-dimensional SOM, a common recommendation is thus to set the side length along each dimension to equal the PCA eigenvalues of the training data.

Second, the *internal* specifications are as follows. The training process starts with a linear initialization of the reference vectors set to the direction of the two

principal components of the input data. The principal component initialization not only further reduces computational cost and enables reproducible results, but has also been shown to be important for convergence when using the batch SOM (Forte et al. 2002). Following Kohonen (2001), this is done in three steps:

- (i) Determine two eigenvectors, v_1 and v_2 , with the largest eigenvalues from the covariance matrix of all data Ω .
- (ii) Let v_1 and v_2 span a two-dimensional linear subspace and fit a rectangular array along it, where the two dimensions are the eigenvectors and the center coincides with the mean of Ω . Hence, the direction of the long side is parallel to the longest eigenvector v_1 with a length of 80% of the length of v_1 . The short side is parallel to v_2 with a length of 80% of the length of v_2 .
- (iii) Identify the initial value of the reference vectors $m_i(0)$ with the array points, where the corners of the rectangle are $\pm 0.4v_1 \pm 0.4v_2$.

The implementations in this book make use of the commonly utilized Gaussian neighborhood function. Its properties are desired as it gives a non-linearly increasing weight to data the closer they are to the updated unit, highlighting the importance of close-by neighbors. The neighborhood function $h_{ib(j)} \in (0, 1]$ is defined as the following Gaussian function:

$$h_{ib(j)} = \exp\left(-\frac{d_r(b, i)^2}{2\sigma^2(t)}\right) \quad (4.12)$$

where $d_r(b, i)$ is the distance between the coordinates r_b and r_i of the reference vectors m_b and m_i on the two-dimensional grid. Moreover, the radius of the neighborhood $\sigma(t)$ is a monotonically decreasing function of time t . Here, Kohonen stresses that special caution is required in the choice of the starting radius to achieve global ordering, as otherwise one risks ending up with mosaic-like patterns. The radius of the neighborhood begins as half the diagonal of the grid size $((X^2 + Y^2)/2)$, and decreases towards a user-specified radius σ . As above mentioned, external parameters outside the SOM machinery need to be specified and discussed during each training phase.

4.4.2 Supervision of the SOM

The standard SOM may be used in a semi-supervised manner [see, e.g., Kohonen's (1991) Hypermap]. It is most common to use the SOM for unsupervised learning, where input data are used for learning previously unknown patterns in data. However, if one possesses class information (e.g., labels), they can be used to supervise learning for a classification task. The main rationale for using the SOM over more traditional methods for classification is its inherent local modeling property and topology preservation of units that enhances the understanding of the problem, as well as the availability of, for instance, growing SOMs that facilitate the choice of

parsimony [for a thorough review see Barreto (2007)]. While unsupervised versions use only the explanatory variables in matching (Eq. 4.9), the supervision of the semi-supervised versions is introduced by the use of both the explanatory and the class variables in matching. Both unsupervised and supervised versions may or may not include the classes in the batch update (Eq. 4.10) without affecting the general learning procedure.² An additional possibility with the semi-supervised SOM is the use of multiple classes, rather than only binary. Multi-class supervision may be thought of as a way to separate all, say four, classes in data, which involves an even better understanding of the dimension reduction as the classes might be associated with separately interpretable properties.

4.4.3 Qualities of the SOM

When using the SOM for a classification task where class labels are known, a direct and obvious measure of quality is the classification performance (see Chap. 7 for a further discussion of measuring classification performance). However, rather than the quality of a classification, the SOM is most often measured in terms of the quality of the unsupervised data and dimension reduction. The literature has provided a large number of metrics for measuring different qualities of the SOM [see Pözlbauer (2004) for a review]: quantization error, topographic product, topographic error, trustworthiness, neighborhood preservation and the distortion measure. Herein, the three most common goodness measures are illustrated: quantization error, distortion measure and topographic error.

The fit of the SOM to the data distribution can be measured with the standard quantization error and distortion measure. The quantization error ε_{qe} computes the average distance between x_j and m_b :

$$\varepsilon_{qe} = \frac{1}{N} \sum_{j=1}^N \|x_j - m_{b(j)}\|, \quad (4.13)$$

The distortion measure ε_{dm} indicates, similarly, the fit of the map to the shape of the data distribution, but also accounts for the radius of the neighborhood:

$$\varepsilon_{dm} = \frac{1}{N} \frac{1}{M} \sum_{j=1}^N \sum_{i=1}^M h_{ib(j)} \|x_j - m_{b(j)}\|, \quad (4.14)$$

² In the literature, learning of the SOM has been defined through the entire spectrum of supervision. For instance, van Heerden and Engelbrecht (2008) define semi-supervised SOMs as similar to the supervised ones, except for them not being included in the matching phase (Eq. 4.9), whereas the semi-supervised version herein is their supervised SOM. However, as the SOM is never fully supervised, we stick to the definition of an unsupervised and a semi-supervised version.

The topology preservation of the SOM can be measured using the standard topographic error ε_{te} :

$$\varepsilon_{te} = \frac{1}{N} \sum_{j=1}^N u(x_j), \quad (4.15)$$

where $u(x_j)$ measures the average proportion of $x_j \in \Omega$ for which first and second BMUs are non-adjacent units.

4.5 Concluding Summary

This chapter has provided a necessary overview of not only data and dimension reduction methods, but also their relation to the KDD process, information visualization and visual analytics. Thereby, the emphasis is clearly on dimension reduction methods and their relation to the above mentioned topics. Likewise, a greater focus has been on first-generation dimension reduction methods to support a subsequent comparison of methods. To this end, this chapter has provided a basis for a more thorough comparison of data and dimension reduction methods, as well as their combination for data-dimension reduction, for financial performance analysis and macroprudential oversight. This is a crucial task as the choice of method is not always a straightforward, quantitative decision to make.

References

- Anand, S., & Buchner, A. (1998). *Decision support using data mining*. London: Financial Time Management.
- Baddeley, A., & Logie, R. (1999). Working memory: The multiple-component model. In A. Miyake & P. Shah (Eds.), *Models of working memory* (pp. 28–61). New York: Cambridge University Press.
- Barreto, G. (2007). Time series prediction with the self-organizing map: A review. In P. Hitzler & B. Hammer (Eds.), *Perspectives on neural-symbolic integration*. Heidelberg: Springer-Verlag.
- Bederson, B., & Shneiderman, B. (2003). *The craft of information visualization: Readings and reflections*. San Francisco, CA: Morgan Kaufman.
- Belkin, M., & Niyogi, P. (2001). Laplacian eigenmaps and spectral techniques for embedding and clustering. In T. Dietterich, S. Becker & Z. Ghahramani (Eds.), *Advances in neural information processing systems* (Vol. 14, pp. 586–691). Cambridge, MA: MIT Press.
- Bertin, J. (1983). *Semiology of graphics*. Madison, WI: The University of Wisconsin Press.
- Bezdek, J. (1981). *Pattern recognition with fuzzy objective function algorithms*. New York: Plenum Press.
- Bishop, C., Svensson, M., & Williams, C. (1998). GTM: The generative topographic mapping. *Neural Computation*, 10(1), 215–234.
- Cabena, P., Hadjinian, P., Stadler, R., Verhees, J., & Zanasi, A. (1998). *Discovering data mining: From concepts to implementation*. New Jersey: Prentice Hall.

- Card, S., Mackinlay, J., & Schneidermann, B. (1999). *Readings in information visualization, using vision to think*. San Diego, CA: Academic Press.
- Card, S., Robertson, G., & Mackinlay, J. (1991). The information visualizer, an information workspace. In *Proceedings of CHI '91, ACM Conference on Human Factors in Computing Systems, New Orleans* (pp. 181–188).
- Chapelle, O., Schölkopf, B., & Zien, A. (Eds.). (2006). *Semisupervised learning*. Cambridge, MA: MIT Press.
- Chen, L., & Buja, A. (2009). Local multidimensional scaling for nonlinear dimension reduction, graph drawing and proximity analysis. *Journal of the American Statistical Association*, 104, 209–219.
- Cottrell, M., & Letrémy, P. (2005). Missing values: Processing with the Kohonen algorithm. In *Proceedings of Applied Stochastic Models and Data Analysis (ASMDA 05), Brest, France* (pp. 489–496).
- Cox, T., & Cox, M. (2001). *Multidimensional scaling*. Boca Raton, Florida: Chapman & Hall/CRC.
- Deboeck, G. (1998a). Best practices in data mining using self-organizing maps. In G. Deboeck & T. Kohonen (Eds.), *Visual explorations in finance with self-organizing maps* (pp. 201–229). Berlin: Springer-Verlag.
- Deboeck, G. (1998b). Software tools for self-organizing map. In G. Deboeck & T. Kohonen (Eds.), *Visual explorations in finance with self-organizing maps* (pp. 179–194). Berlin: Springer-Verlag.
- Demartines, P., & Héroult, J. (1997). Curvilinear component analysis: A self-organizing neural network for nonlinear mapping of data sets. *IEEE Transactions on Neural Networks*, 8, 148–154.
- Dempster, A., Laird, N., & Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society (Series B)*, 39(1), 1–38.
- Dunn, J. (1973). A fuzzy relative of the isodata process and its use in detecting compact, well-separated clusters. *Cybernetics and Systems*, 3, 32–57.
- Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In E. Simoudis, J. Han & U. Fayyad (Eds.), *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD 96)* (pp. 226–231). AAAI Press.
- Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996a). From data mining to knowledge discovery: An overview. In U. Fayyad, G. Piatetsky-Shapiro, P. Smyth & R. Uthurusamy (Eds.), *Advances in knowledge discovery and data mining* (pp. 1–34). Menlo Park, CA: AAAI Press / The MIT Press.
- Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996b). The KDD process for extracting useful knowledge from volumes of data. *Communications of the ACM*, 39(11), 27–34.
- Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996c). Knowledge discovery and data mining: Towards a unifying framework. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining, Portland, OR* (pp. 82–88).
- Fekete, J.-D., van Wijk, J., Stasko, J., & North, C. (2008). The value of information visualization. In *Information visualization: Human-centered issues and perspectives* (pp. 1–18). Springer.
- Forte, J., Letrémy, P., & Cottrell, M. (2002). Advantages and drawbacks of the batch Kohonen algorithm. In *Proceedings of the European Symposium on Artificial Neural Networks (ESANN 02), Bruges, Belgium* (pp. 223–230).
- Frawley, W., Piatetsky-Shapiro, G., & Matheus, C. (1992). Knowledge discovery in databases: An overview. *AI Magazine*, 13(3), 57–70.
- Gisbrecht, A., Hofmann, D., & Hammer, B. (2012). Discriminative dimensionality reduction mappings. In *Proceedings of the International Symposium on Intelligent Data Analysis* (pp. 126–138). Helsinki, Finland: Springer-Verlag.
- Haroz, S., & Whitney, D. (2012). How capacity limits of attention influence information visualization effectiveness. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), 2402–2410.
- Havre, S., Hetzler, B., & Nowell, L. (2000). Themeriver: Visualizing theme changes over time. In *Proceedings of the IEEE Symposium on Information Visualization* (pp. 115–123).

- Hoaglin, D., Mosteller, F., & Tukey, J. (1983). *Understanding robust and exploratory data analysis*. New York: Wiley.
- Jain, A. (2010). Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8), 651–666.
- Jain, A., Murty, M., & Flynn, P. (1999). Data clustering: A review. *ACM Computing Surveys*, 31(3), 264–323.
- Kaser, O., & Lemire, D. (2007). Tag-cloud drawing: Algorithms for cloud visualization. In *Proceedings of the Tagging and Metadata for Social Information Organization Workshop, Banff, Alberta, Canada*.
- Keim, D. (2001). Visual exploration of large data sets. *Communications of the ACM*, 44(8), 38–44.
- Keim, D., Kohlhammer, J., Ellis, G., & Mannsmann, F. (2010). *Mastering the information age. Solving problems with visual analytics*. Goslar: Eurographics Association.
- Keim, D., & Kriegel, H.-P. (1996). Visualization techniques for mining large databases: A comparison. *IEEE Transactions on Knowledge and Data Engineering*, 8(6), 923–938.
- Keim, D., Mansmann, F., Schneidewind, J., & Ziegler, H. (2006). Challenges in visual data analysis. In *Proceedings of the IEEE International Conference on Information Visualization (iV 13)* (pp. 9–16). London, UK: IEEE Computer Society.
- Keim, D., Mansmann, F., & Thomas, J. (2009). Visual analytics: How much visualization and how much analytics? *SIGKDD Explorations*, 11(2), 5–8.
- Koffa, K. (1935). *Principles of gestalt psychology*. London: Routledge & Kegan Paul.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59–69.
- Kohonen, T. (1991). The hypermap architecture. In T. Kohonen, K. Mäkisara, O. Simula & J. Kangas (Eds.), *Artificial neural networks* (Vol. II, pp. 1357–1360). Amsterdam, Netherlands: Elsevier.
- Kohonen, T. (2001). *Self-organizing maps* (3rd ed.). Berlin: Springer-Verlag.
- Kruskal, J. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29, 1–27.
- Kurgan, L., & Musilek, P. (2006). A survey of knowledge discovery and data mining process models. *The Knowledge Engineering Review*, 21(1), 1–24.
- Lampinen, J., & Oja, E. (1992). Clustering properties of hierarchical self-organizing maps. *Journal of Mathematical Imaging and Vision*, 2(2–3), 261–272.
- Larkin, J., & Simon, H. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65–99.
- Lee, J., & Verleysen, M. (2007). *Nonlinear dimensionality reduction*. Information science and statistics series. Heidelberg, Germany: Springer-Verlag.
- Lin, X. (1997). Map displays for information retrieval. *Journal of the American Society for Information Science*, 48(1), 40–54.
- Linde, Y., Buzo, A., & Gray, R. (1980). An algorithm for vector quantizer design. *IEEE Transactions on Communications*, 28(1), 702–710.
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* (pp. 281–297). Berkeley, CA: University of California Press.
- Moehrmann, J., Burkovski, A., Baranovskiy, E., Heinze, G., Rapoport, A., & Heideman, G. (2011). A discussion on visual interactive data exploration using self-organizing maps. In J. Laaksonen & T. Honkela (Eds.), *Proceedings of the 8th International Workshop on Self-Organizing Maps* (pp. 178–187). Helsinki, Finland: Springer-Verlag.
- Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(6), 559–572.
- Pözlbauer, G. (2004). Survey and comparison of quality measures for self-organizing maps. In *Proceedings of the 5th Workshop on Data Analysis (WDA 2004), Sliezsky dom, Vysoké Tatry, Slovakia* (pp. 67–82).
- Roweis, S., & Saul, L. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290, 2323–2326.

- Rubin, D. (1987). *Multiple imputation for nonresponse in surveys*. New York: Wiley & Sons.
- Sammon, J. (1969). A non-linear mapping for data structure analysis. *IEEE Transactions on Computers*, 18(5), 401–409.
- Sarlin, P. (2014a). Macroprudential oversight, risk communication and visualization. [arXiv:1404.4550](https://arxiv.org/abs/1404.4550).
- Shannon, C., & Weaver, W. (1963). *A mathematical theory of communication*. Champaign: University of Illinois Press.
- Shearer, C. (2000). The CRISP-DM model: The new blueprint for data mining. *Journal of Data Warehousing*, 15(4), 13–19.
- Shepard, R. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. *Psychometrika*, 27(125–140), 219–246.
- Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages, Boulder, CO* (pp. 336–343).
- Tenenbaum, J., de Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290, 2319–2323.
- Thomas, J., & Cook, K. (2005). *Illuminating the path: Research and development agenda for visual analytics*. Los Alamitos: IEEE Press.
- Torgerson, W. S. (1952). Multidimensional scaling: I. theory and method. *Psychometrika*, 17, 401–419.
- Triesman, A. (1985). Preattentive processing in vision. *Computer Vision, Graphics and Image Processing*, 31(2), 156–177.
- Tufte, E. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.
- Tukey, J. (1977). *Exploratory data analysis*. Reading, PA: Addison-Wesley.
- van der Maaten, L., & Hinton, G. (2008). Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research*, 9, 2579–2605.
- van Heerden, W., & Engelbrecht, A. (2008). A comparison of map neuron labeling approaches for unsupervised self-organizing feature maps. In *Proceedings of the IEEE International Joint Conference on Neural Networks* (pp. 2139–2146). Hong Kong: IEEE Computer Society.
- Venna, J., & Kaski, S. (2006). Local multidimensional scaling. *Neural Networks*, 19, 889–899.
- Vesanto, J., Himberg, J., Alhoniemi, E., & Parhankangas, J. (2000). *SOM toolbox for Matlab 5*. Technical Report: Helsinki University of Technology. A57.
- Ward, J. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58, 236–244.
- Ware, C. (2004). *Information visualization: Perception for design*. San Francisco, CA: Morgan Kaufman.
- Ware, C. (2005). Visual queries: The foundation of visual thinking. In S. Tergan & T. Keller (Eds.), *Knowledge and information visualization* (pp. 27–35). Berlin, Germany: Springer.
- Weinberger, K., & Saul, L. (2005). Unsupervised learning of image manifolds by semidefinite programming. *International Journal of Computer Vision*, 70(1), 77–90.
- Wismüller, A. (2009). A computational framework for non-linear dimensionality reduction and clustering. In J. Principe & R. Miikkulainen (Eds.), *Proceedings of the Workshop on Self-Organizing Maps (WSOM 09)* (pp. 334–343). St. Augustine, Florida, USA: Springer.
- Yin, H. (2008). The self-organizing maps: Background, theories, extensions and applications. In J. Fulcher & L. Jain (Eds.), *Computational intelligence: A compendium* (pp. 715–762). Heidelberg, Germany: Springer-Verlag.
- Young, G., & Householder, A. S. (1938). Discussion of a set of points in terms of their mutual distances. *Psychometrika*, 3, 19–22.
- Zhang, J., & Liu, Y. (2005). SVM decision boundary based discriminative subspace induction. *Pattern Recognition*, 38(10), 1746–1758.
- Zhang, L., Stoffel, A., Behrisch, M., Mittelstädt, S., Schreck, T., Pompl, R., et al. (2012). Visual analytics for the big data era—a comparative review of state-of-the-art commercial systems. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST), Seattle, WA* (pp. 173–182).

Chapter 5

Data-Dimension Reductions: A Comparison

Data and dimension reduction techniques, and particularly their combination for DDR, have in many fields and tasks held promise for representing data in an easily understandable format. However, comparing methods and finding the most suitable one is a challenging task. Above, we discussed the aim of dimension reduction in terms of three tasks. For the third task of visualization, the most popular method has been the SOM, which is oftentimes asserted as an artifact of its simplicity and intuitive formulation [e.g., Lee and Verleysen (2007), Trosset (2008)]. Yet, being well-known or simple, while being an asset, is not a proper validation of relative goodness. The focus of this chapter is to challenge the superiority of the SOM by comparing it to alternative methods.

To capture the most suitable methods for visual financial performance analysis according to the needs for the task, this chapter assesses the suitability of three classical, or so-called first-generation, dimension reduction methods: metric MDS (Torgerson 1952), Sammon's mapping (Sammon 1969) and the SOM (Kohonen 1982). Rather than being the most recent methods, the rationale for comparing these is to capture the suitability of well-known dimension reduction methods with inherently different aims: global and local distance preservation and topology preservation, respectively. For DDR, and due to access to overabundant amounts of data, we look into test serial and parallel combinations of the projections with three data reduction or compression methods: VQ (Linde et al. 1980), *k*-means clustering (MacQueen 1967), Ward's (1963) hierarchical clustering. While conceptually being similar, the functioning of the SOM differs from the other DDR combinations as the two tasks of data and dimension reduction are treated as concurrent subtasks. In serial combinations, the dimension reduction is always subordinate to the data reduction, whereas parallel combinations deal separately with the initial dataset.

This chapter compares DDR combinations to financial performance analysis as follows. After a general review of the literature on comparisons of data and dimension

This chapter is partly based upon previous research. Please see the following work for further information: Sarlin (2014)

reduction methods, we discuss the aims and needs of DDR combinations in general and for the task at hand in particular. That is, building low-dimensional mappings from high-volume and high-dimensional data that function as displays for additional information, be it individual data (e.g., time series of entities) or general structural properties of data (e.g., qualities, distance structures and densities). The relative goodness of methods for financial performance analysis will then be discussed from a qualitative perspective. Further, experiments on a dataset of annual financial ratios for European banks is used to illustrate the general applicability of the DDR combinations for the task. After illustrating some approaches to link information to the visualization displays, results of these comparisons are then projected to the second generation of dimension reduction methods for a final discussion on the superiority of methods for overall visual financial performance analysis, including tasks for macroprudential oversight, as well as the general applicability of this comparison. These discussions also include an information visualization perspective to dimension reductions.

5.1 The Optimal Method: A Literature Review

When reviewing the literature on method comparisons, we first focus on dimension reduction methods and then on data reduction methods. The focus is on neutral evaluations of methods rather than evaluations in papers presenting novel methods. While papers presenting new methods generally include an evaluation and conclude at least partial superiority of it, such as some of those found in Sect. 4.3, they may be biased to a lesser or greater extent towards data and evaluation measures suitable for that particular approach.

5.1.1 *A Comparison of Dimension Reductions*

The large number of methods has obviously also stimulated a large number of performance comparisons between them. The comparisons mainly vary in terms of used data and evaluation measures, whereas there may still be some variation in the precise utilization of methods. For instance, Flexer (1997, 2001) used Pearson correlation, Duch and Naud (1996) hypercubes in 3–5 dimensions and Bezdek and Pal (1995) the metric topology preserving index to show that MDS outperforms the SOM. Trosset (2008) argues that a serial combination of clustering and MDS is superior to the SOM. Venna and Kaski (2001) and Nikkilä et al. (2002) show superiority of the SOM and GTM in terms of trustworthiness of neighborhood relationships, while later Himberg (2004) and Venna and Kaski (2007) show superiority of CCA in terms of the same measure. Not surprisingly, de Vel et al. (1996) show, using Procrustes analysis and Spearman rank correlation coefficients on various datasets, that the superiority of

a method depends on the used evaluation measures and data. Hence, despite many attempts, inconsistent comparisons do not indicate the superiority of one method.

Lately, Lee and Verleysen (2009) proposed a unified measure based upon a co-ranking matrix for evaluating dimension reductions, an adequate ground for generic evaluations. Lueks et al. (2011) further developed the measure by letting the user specify the properties that are more important to be preserved. While being useful aids in comparing methods, they neither show nor propose existence of one superior method for every type of data and preferences of similarity preservation.

5.1.2 A Comparison of Data Reductions

When reviewing the literature on methods for data reduction, one can easily observe that neither is there a unanimity on the best available method. Herein, the focus is on comparisons between the SOM and stand-alone data reduction methods. Bação et al. (2005) show that the SOM outperforms k -means clustering with 3 evaluation measures and 4 datasets. Flexer (1997, 2001) show that k -means clustering outperforms the SOM using a Rand index and 36 datasets. Waller et al. (1998) show on 2,580 datasets that the SOM performs equally well as k -means clustering and better than other methods. Balakrishnan et al. (1994) show that k -means outperforms the SOM on 108 datasets, but do not decrease the SOM neighborhood to zero at the end of learning [as, e.g., Kohonen (2001) proposes]. Vesanto and Alhoniemi (2000) showed on 3 datasets that two-level clustering of the SOM is equally accurate as agglomerative and partitive methods, while being computationally cheaper and having merits in visualizing relations in data. Ultsch and Vetter (1994) compare the SOM with hierarchical and k -means clustering and conclude that the SOM not only provides an equally accurate result, but also an easily interpretable output. Despite no unanimity on superiority, the literature still indicates that the SOM, and its adaptations, are equally considerable alternatives for data reduction as other methods, such as centroid-based and hierarchical clustering.

5.1.3 Why is the Literature so Divided?

While the quality of data reductions can be quantified by common evaluation measures like quantization error, assessing the superiority of one dimension reduction method over others with a quantitative measure is more difficult. And there is still no unanimity on the superiority of one data reduction method over others. What varies in the above discussed studies is mainly the underlying data, which indicates that methods show different performance on different types of data. One reason might be that clusters in the SOM topology learn from and are provided guidance by neighboring data as well, which aids the analysis of noisy data, whereas accuracy suffers on well-behaving toy data. This is supported by the findings of de Bodt et al. (1999)

and Bação et al. (2005), where they propose that the SOM better spans the search space as neighborhood relations force units to follow each other. This is, however, only speculative reasoning about the above lack of unanimity.

Since the mid-20th century, the overload of available data has stimulated a soar in the development of dimension reduction methods with inherent differences (as reviewed in Chap. 4). However, most differences in the quality of dimension reductions, as all structural information can impossibly be preserved in a lower dimension, derive from variations in preserved similarity relations, such as pairwise distances or topological relationships. The performance, and choice of model specification, of one method can generally be motivated by its own quantitative quality measure. However, the relative goodness of different methods depend strongly on the correspondence between the particular quality measure and the objective function.

Despite the fact that the large number of dimension reduction methods has stimulated quality comparisons along different measures, inconsistency of the comparisons has lead to no unanimity on the superiority of one method [see, e.g., Flexer (1997, 2001) and Venna and Kaski (2001)]. This also indicates that the goodness of methods depends to a large extent on the correspondence between the measure and the objective function, and confirms that the quality measure is a user-specified parameter depending on the task at hand. While recent advances in unified measures for evaluating dimension reductions have included a parameter for the user to specify properties that are more important to be preserved (Lee and Verleysen 2009; Lueks et al. 2011), quantitative measures still have difficulties in including qualitative differences in properties of methods, such as differences in flexibility for difficult data and the shape of the low-dimensional output. This motivates assessing the suitability of data and dimension reduction methods for a specific task from a qualitative perspective.

5.2 DDR Combinations for the Task at Hand

This section discusses specific aims, needs and restrictions of DDR combinations for visual financial performance analysis. Based upon this discussion, we look into dimensions of DDR combinations relevant for measuring the suitability of methods for the task herein.

5.2.1 *Aims and Needs for the Task*

So, what is the so-called task at hand? The aim of models for visual financial performance analysis, including tasks for macroprudential oversight, is to represent high-volume and high-dimensional data of financial entities on low-dimensional displays. The data for such a task are derived from a data cube, as the one represented in Fig. 3.1 (see Sect. 3.3). Data and dimension reductions hold promise for the task, but the form of the models still set some specific needs and restrictions. While recent advances

in information technology have enabled access to databases with nearly endless amounts of macroeconomic and financial information (e.g., Bankscope, Bloomberg, Standard & Poor's and Capital IQ), as well as provision and integration of multiple sources (e.g., Haver Analytics), data are oftentimes problematic in being incomplete and non-normal (e.g., Deakin 1976). For instance, in the case of representing a financial entity with its balance-sheet information, it is more common than not that some items of the balance sheet are missing. Due to changes in reporting rules and financial innovation, data might be missing or start in the latter part of a time series. An example of skewed distributions is the commonly appearing power-law distribution and Benford's law, as well as the particularly fat tails of market-based data. While there exist a multitude of preprocessing methods for transforming, normalizing and trimming data, the tails of financial ratio distributions are oftentimes of high interest. This derives two necessities: the computational cost of the method needs to be considerably low and scalable and the method needs to be flexible for problematic data.

The main aim of the low-dimensional mappings is to use them as displays for additional information, in particular for: (i) individual data, (ii) structural properties of data, and (iii) qualities of the models. This is due to three respective reasons:

- (i) the two-dimensional plane should function as a basis or display for visual performance comparisons of financial entities (i.e., observation-level data) and their time series;
- (ii) for the human visual system to recognize patterns in data, we need to provide guidance for interpreting general data structures, and oftentimes also possess this types of linkable information; and
- (iii) qualities of a dimension reduction may vary across mappings and locations in mappings as all information cannot be correctly preserved in a lower dimension.

The main aim of these mappings is hence not to be an ending point, but rather to function as a basis for a wide range of additional visualizations.

5.2.2 Aims and Needs of DDR Combinations

When evaluating or comparing performance of data and dimension reduction methods, particularly DDR combinations, quantitative measures have difficulties in accounting for qualitative differences in properties of methods. Hence, as the performed comparison is qualitative, the needs for visual financial performance analysis are suppressed into four qualitative criteria for evaluating DDR combinations: form of structure preservation, computational cost, flexibility for problematic data and shape of the output. Next, we discuss these criteria in more detail.

Form of Structure Preservation

As all relations in a high dimensional space can obviously not be preserved in a lower dimension, there are differences in what locations are stressed when preserving the structure. Given these differences, the main characteristics of structure preservation should obviously match important desires of the particular task at hand. The key question is thus: *Which relations are of central importance for visual financial performance analysis?* With a main focus on visualizing individual financial entities on a low-dimensional display, correctly locating neighboring data becomes essential. This leads to trustworthiness of neighborhood relationships being more important than precision on the exact distance to those far away. Noise and erroneous data as well as comparability issues related to reporting differences, for instance, also motivate attempting this type of a local order-preserving mapping rather than focusing on global detail.

Computational Cost

We oftentimes have access to vast amounts of macro-financial data in today's databases, including high-dimensional data for a large number of entities with a high frequency over long periods (i.e., a large data cube along all three dimensions), not the least if the used data are based upon market sources. This obviously sets some restrictions on computational cost and scalability of methods. While computation time is not entirely a qualitative property, it has still not been incorporated in quantified evaluation measures. As also noted by van der Maaten and Hinton (2008), the practical applicability of a dimension reduction method relies upon its computational complexity, as application becomes infeasible if the computational resources needed are too large. In addition to the properties of data, computational cost of a method is set by the dimensionality of the output, the definition of a neighborhood in the case of neighborhood preservation and for iterative techniques the number of iterations, not to mention the form of input data (e.g., pairwise distance matrices or high-dimensional data points). It is also worth to consider that computational expense is not only a one-off cost when creating a dimension reduction, but also when updating it. Combinations with data reduction methods may also affect the computational cost of a dimension reduction. Still, it is important to acknowledge that a cut-off between computationally costly and non-costly methods is difficult. Yet, the differences between methods oftentimes tend to be significant.

Flexibility for Problematic Data

Methods differ in flexibility for non-normal and incomplete data, something more common than not in real-world macro-financial settings. Hence, desired properties of dimension reduction methods are flexibility for incomplete and non-normal data. While the former can be defined in terms of treatment of missing values, the latter

depends largely on the task at hand. Most often data are preprocessed for ideal results, including treatment of skewed distributions. Yet, preprocessing seldom does, and is most often not desired to, compress the data into uniform density. Oftentimes, the most extreme values of data are among the most interesting states of financial performance. Hence, one type of tolerance towards outliers can be derived from the output of methods. A method is judged to be tolerant towards outliers and skewed distributions if problematic data do not significantly impair the intelligibility of an output or display (e.g., stretch towards outliers).

Shape of the Output

One of the main aims is to use a dimension reduction as a display to which additional information is linked. In particular, the low-dimensional mappings are used as displays for individual data, structural properties and qualities. This turns the focus to the shape of the outputs of dimension reduction mappings. They can take a wide range of forms. The interrelated properties of the shape can be considered to be the following: continuous versus discrete mappings, optional versus mandatory data reductions and predefined versus data-driven grid shapes. While a mandatory data reduction is generally not desirable, it is not considered a significant disadvantage. Rather the opposite, due to the large amounts of available data. This leads also to restricting mappings to discrete rather than continuous, whereas continuous mappings would obviously be desirable from the perspective of detail and accuracy. The largest difference for interpretation, especially in terms of linking visualizations, is between predefined and data-driven grid shapes. While methods with data-driven grid shapes may better adapt to data, the methods with predefined regular shapes are superior in functioning as a regularly formed display for additional information. This is a key property as the mappings are starting points rather than ending points of the analysis, where additional information may be individual data, structural properties of data and qualities of the models.

5.3 A Qualitative Comparison

This section presents a qualitative discussion of DDR combinations for visual performance analysis and relates it to the four identified criteria: form of structure preservation, computational cost, flexibility for problematic data and shape of the output. Below, we discuss MDS, Sammon's mapping and the SOM from the viewpoint of the task at hand and the four criteria.

Form of Structure Preservation

The main difference between DDR combinations is how the dimension reduction methods differ in the properties of data they attempt to preserve. For the task of

visual financial performance analysis, the focus is on one question: *Which methods better assure trustworthy neighbors?* MDS-based methods with objective functions attempting distance preservation, while potentially being better at approximating distance structures, may end up with skewed errors across the projection. To this end, Venna and Kaski (2001) and Nikkilä et al. (2002) have shown that the SOM, which stresses neighborhood relations, better assures trustworthy neighbors. That is, data found close-by each other on a SOM display are more likely to be similar in terms of the original data space as well. The conceptual difference in structure preservation between distance- and topology-preserving methods is illustratively described by Kaski (1997) with an experiment on a curved two-dimensional surface in a three-dimensional space: the former methods may follow the surface in data with two dimensions, whereas the latter require three dimensions to describe the structure.

Computational Cost

Expensive computations is obviously an issue when dealing with large-volume financial data. Generally, computing pairwise distances between data is costly with an order of magnitude of N^2 . The topology preservation of the SOM relates instead to the grid size M with an order of magnitude of M^2 (Kaski 1997). This implies that the complexity of the methods are similar if the grid size M equals the number of data N , but more importantly that the SOM allows for adjusting M for cheaper complexity. Further, parallel DDR combinations suffer from an additional computational cost as the clustering is performed on the initial dataset rather than on a reduced number of units. The computational cost of MDS-based methods motivates serial DDR combinations. Another issue related to computational cost is the lack of an explicit mapping function for the MDS-based methods. Hence, when including new samples, the projection needs to be recomputed. While new samples can be visualized via projection to their best-matching data, each update requires recomputing the projection.¹ In contrast, the SOM can cheaply be updated with individual data using the sequential algorithm (i.e., an online version of the batch SOM).

Flexibility for Problematic Data

The methods significantly differ in flexibility for problematic data. Methods dealing with distance preservation have obvious difficulties with incomplete data. However, the SOM, and its self-organization, can be seen as tolerant to missing values by only considering the available ones in matching (Samad and Harp 1992). In practice, the SOM has been shown to be robust when up to approximately 1/3 of the variables in a row (i.e., data vector x_j) are missing (Kaski and Kohonen 1996; Kohonen 2001; Denny Squire 2005; Sarlin 2012b). Indeed, the SOM has even been shown to be

¹ While Relative MDS (Naud and Duch 2000) allows to add new data to the basis of an old MDS, it does still not update all distances within the mapping.

effective for imputing missing values (e.g., Cottrell and Letrémy 2005). Tolerance towards outliers is measured in terms of representation of skewed distributions. An MDS-based mapping becomes difficult to interpret if it is stretched towards directions of outliers and extreme tails. While the processing of the SOM does not per se treat outliers, its regularly shaped grid of units facilitates visualizing data with non-uniform density functions. This provides a hint of the final criterion.

Shape of the Output

A key to using a dimension reduction as a display, and linking information to it, is the shape of its output. Whereas the SOM has a discrete mapping, mandatory data reduction and predefined grid shape, MDS-based methods are its contrasts by having continuous mappings, optional data reduction and data-driven lattice (if combined with data reduction). The predefined SOM grid, while also having drawbacks for representing structural properties of data, facilitates the interpretation of linked information. Today, it is standard that the SOM comes with a wide set of linked extensions for visual analytics, such as the so-called feature planes, U-matrix and frequency plots (Vesanto 1999). Even though visual aids for showing distance structure and density compensate for constraints set by the grid shape, there is a large group of other aids that enhance the representation of available information in data. The visual aids, while not always being even applicable, have generally not been explored in the context of MDS-based projections. Feature planes (see Sect. 5.4), for instance, are difficult to visualize due to the lack of a reduced number of units. Even DDR combinations with serial VQ, i.e., processing similar to that of the SOM, would still lack the concept of neighborhood relations of a regularly shaped grid.

5.4 Illustrative Experiments

The qualitative discussion of properties of DDR combinations for financial performance analysis still lacks illustrations of the above discussed properties of methods. This section presents experiments with these methods. Dimension reduction is performed with the SOM, metric MDS and Sammon's mapping and data reduction with Ward's hierarchical clustering, *k*-means clustering and VQ. We explore various combinations for DDR with the aim of achieving easily interpretable models for visual financial performance analysis. The methods are chosen and combined as to their suitability for data reduction of dimension reductions, and *vice versa*.

Data

The dataset used in these examples consists of annual financial ratios for banks from the EU, including all provided financial ratios in the Bankscope database from

Bureau van Dijk. Initially, the dataset consisted of 38 annual financial ratios for 1,236 banks spanning from 1992:12–2008:12. A large concern in the dataset is the share of missing values, due to which 24 ratios were chosen by dropping those with more than 25 % missing data. Observations with missing values for more than 1/3 of the ratios were removed. Finally, we are left with a resulting 9,655 rows of data, and a total of 855 banks. Yet, the dataset still includes missing values. Although the SOM is tolerant to missing data, we need to impute them in this work as distance-preserving methods require complete data. For simplicity, the SOM is used for imputing missing values. A SOM allows mapping incomplete data to their best-matching units (BMUs) by only considering the available variables. Hence, complete data were used for training a SOM, incomplete data were mapped to their BMUs and the missing values were imputed from their BMUs. Moreover, although outliers are not a problem *per se*, they may still affect the interpretability of the models, in particular MDS-based models. Not to lose significant amounts of data, modified boxplots are used for trimming with replacement. The modified boxplot is preferred over Winsorizing, for instance, as it accounts for variable-specific distributions, resulting in replacement of a total of 7.39 % of the data, distributed as needed per variable and tail. In the following experiments, we use the entire dataset, in particular when creating displays with data and dimension reduction methods. Further, a sample of trajectories are used to illustrate the visualization of individual data on the created displays. The trajectories consist of all input variables spanning from 2002 to 2008 for Deutsche Bank, ABN Amro and Société Général.

Parallel DDR

Figure 5.1 shows parallel DDR combinations on the entire dataset. Sammon's mapping is combined with k -means clustering, and MDS and the SOM are combined with Ward's clustering.² Ward's clustering of the SOM is, however, performed on its units rather than on the dataset and restricted to agglomerate only adjacent clusters in the SOM topology. This option is not, however, considered for MDS-based projections as there is no natural definition of adjacency. On top of all three mappings, we can observe a superimposed cluster color coding and performance comparison of trajectories from 2002–2008 for three large European banks. Cluster memberships are visualized through a qualitative color scheme from ColorBrewer (Harrower and Brewer 2003), where groups are differentiated in hue contrast with nearly constant saturation and lightness. The projections of MDS and Sammon's mapping on this large dataset are very similar, whereas k -means clustering has less overlapping cluster memberships in the mapping than Ward's clustering. The trajectories as well as the underlying variables confirm that, while the orientations of the two MDS-based

² When training SOMs, one has to set a number of free parameters. A set of quality measures is used to track the topographic and quantization accuracy as well as clustering of the map. Given the purpose herein, details about the parametrization of the models in the experiments are not presented in depth.




Fig. 5.1 Parallel DDR combinations. *Notes* The figures show parallel DDR combinations on the entire financial dataset; Sammon’s mapping is combined with *k*-means clustering, and MDS and the SOM are combined with Ward’s clustering. Color codes on each mapping correspond to clusters and the superimposed trajectories to a performance comparison of three large European banks from 2002–2008

projections are somewhat different from those of the SOM model, their structure is still inherently similar. Yet, the computational cost differs significantly. While it takes on an ordinary personal computer only a few seconds to train SOM-based models on these data, the MDS-based projections require several hours on a dedicated server.

Serial DDR

For cheaper complexity, we further explore possibilities of MDS by testing serial combinations. Figure 5.2 shows a Sammon’s mapping of the *k*-means cluster centroids as well of the second-level centroids of the SOM, where size represents the number of data in each cluster. This type of usage of MDS-based methods was already proposed by Sammon (1969) due to their high computational cost, and later applied by Flexer (2001), for instance. It is, indeed, a cheap way to illustrate relations between the cluster centroids, but lacks detail for structural as well as individual analysis.

k-means clustering and Sammon's mapping



SOM, Ward's clustering and Sammon's mapping




Fig. 5.2 Serial DDR combinations. *Notes* The figure shows serial DDR combinations on the entire financial dataset; Sammon's mapping is combined with k -means clustering, and the SOM with second-level Ward's clustering. Color codes on each mapping correspond to clusters. Not to clutter the display, trajectories are not displayed in this figure

Serial and Parallel DDR

Costly, yet detailed, MDS-based projections in Fig. 5.1 and cheap, yet crude, projections in Fig. 5.2 motivate finding a compromise solution. For reducing computational expense, it is still necessary to rely on a serial DDR combination. For more detail, however, the initial dataset is reduced to a smaller but representative dataset. This type of data compression can, for instance, be achieved with standard VQ that approximates probability density functions of data. The compressed reference vectors can then be used as an input for a parallel DDR. Conceptually, while still lacking the interaction between the tasks as well as the regular grid shape, we come close to what is achieved using a SOM in Fig. 5.1 by relying on both serial and parallel DDR combinations. The left plot in Fig. 5.3 shows a VQ of the initial dataset and then a subsequent Sammon's mapping and k -means clustering on the VQ reference vectors. The right plot in Fig. 5.3 shows a corresponding Sammon's mapping of SOM units with a superimposed cluster color coding. However, the figure illustrates two issues: the ordered SOM units have less overlap of cluster memberships and the importance of naturally defined topological relations. The former issue is partly a result of interaction between the tasks of data and dimension reduction and partly of the inclusion of neighborhood relations when agglomerating clusters. The latter issue of a regularly shaped grid is particularly useful when attempting to visualize as much of the available information as possible through linked visualizations.

5.5 The SOM and Its Visualization Aids

This section first briefly reviews visualization aids for the SOM and then illustrates the use of the regularly shaped SOM grid, and its visualization aids. Figure 5.1 showed the two-dimensional SOM grid, and trajectories for three large European banks from




Fig. 5.3 Serial and parallel DDR combinations. *Notes* The figures show serial and parallel combinations on the entire financial dataset; Sammon's mapping is combined with VQ and k -means clustering, and the SOM with Ward's clustering and Sammon's mapping. Color codes on each mapping correspond to clusters and the net-like representation illustrates neighborhood relations

2002–08, but a central question remains: *How should we interpret the map?* The possibility of linking additional information to the SOM grid has stimulated the development of a wide scope of visualization aids [see Vesanto (1999) for an early overview]. These can be classified into three groups:

- (i) those compensating structural properties inherent in data that the regular grid shape eliminates;
- (ii) those extending the visualization of properties inherent in data but not normally accessible in dimension reductions; and
- (iii) those linking the SOM grid with other methods or data to further enhance the understanding of the task.

The first group includes means to represent the distance structure and density on a SOM, something missing due to the VQ and grid shape. Densities on the SOM are generally assessed with frequency plots and the Pareto density estimation matrix (P-matrix) (Ultsch 2003a). Examples of aids for assessing distance structures are Sammon's mapping, the Unified distance matrix (U-matrix) (Ultsch and Siemon 1990) and cluster connections (Merkl and Rauber 1997). Moreover, some methods attempt to account for both structures and densities, such as the U^* -matrix (Ultsch 2003b), the sky metaphor visualization (Latif and Mayer 2007), the neighborhood graph (Pözlbauer et al. 2005), smoothed data histograms (Pampalk et al. 2002), and cluster coloring (Kaski et al. 2001; Sarlin and Rönqvist 2013).

The second group consists of visualizations that enhance the representation of the high-dimensional information. Feature planes are a standard method for visualizing the spread of values of individual dimensions on the SOM, but they have been further enhanced in several aspects. For instance, Vesanto and Ahola (1999) use a SOM for reorganizing the feature planes according to correlations and Neumayer et al. (2007) introduced the metro map discretization to summarize all feature planes onto one plane. Kaski et al. (2001) have developed a visualization of the contribution of each variable to distances between units, that is, the cluster structure. Another extension, while partly also belonging to the other groups, is visualization of vector fields

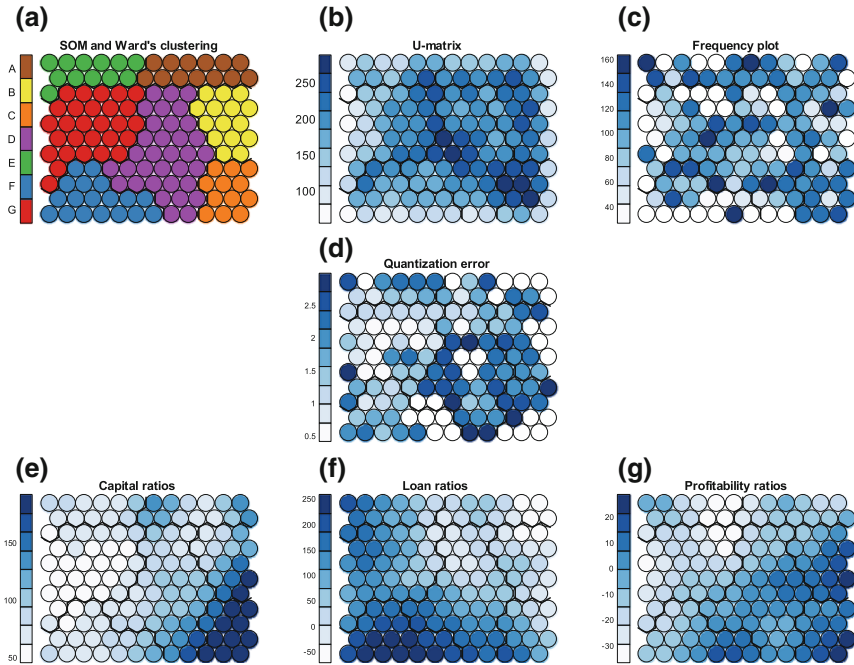


Fig. 5.4 An exemplification of information linked to a SOM. **a** SOM and Ward's clustering, **b** U-matrix, **c** Frequency plot, **d** Quantization error, **e** Capital ratios, **f** Loan ratios, **g** Profitability ratios. *Notes* The figures link additional information to the regularly shaped SOM grid. Charts **(a–c)** illustrate structural properties of the model: **(a)** shows cluster memberships of the second-level clustering, **(b)** shows average distances between units, or the so-called U-matrix, and **(c)** shows the frequency distribution on the SOM grid. Chart **(d)** shows qualities of the model, whereas charts **(e–g)** show the spread of three subdimensions of financial performances on the SOM grid: capital, loan and profitability ratios

(Pözlbauer et al. 2006) for assessing contributions to the cluster structure and for finding correlations and dependencies in the underlying data.

The third group uses other methods or data for further enhancing the understanding of the task. One common way to represent cluster structures in a SOM is applying a second-level clustering on the units, and visualizing it through color coding (Vesanto and Alhoniemi 2000). The reference vectors have been used as an input for other predictive methods, such as a neural network in Serrano-Cinca (1996), whereafter the prediction may be visualized on the SOM grid.

Next, we look at some examples of how visualizations from the above three groups can be linked to the SOM. The previously presented SOM in Fig. 5.1 already showed a financial performance comparison over time of three large European banks using labels and trajectories. Figure 5.4 uses the regular shape of the SOM grid as a basis for seven different representations of additional information. Whereas cluster memberships are visualized through a qualitative color scheme, the rest of

the visualizations are shown through variation in luminance (light to dark to represent low to high values) in a blue hue. It is worth noting that a complicating factor in using luminance is that perceived lightness is dependent on context (Purves et al. 2004), namely the lightness of surrounding colors. For this reason, color scales ought to be presented with a consistent reference color to be comparable in lightness. The units of the SOM are in this book represented with circles rather than hexagons to leave space for reference coloring.

First, Fig. 5.4a, b, c illustrate structural properties of the model: (a) shows crisp cluster memberships of the second-level clustering, (b) shows distance structures using a U-matrix visualization, and (c) shows the frequency distribution on the SOM grid. While Fig. 5.4a, b show similar characteristics of cluster structures, Fig. 5.4c shows no specific patterns in density, except for borders being comparatively less dense. Second, Fig. 5.4d shows qualities of the model, where larger quantization errors cluster around the lower right corner. Third, Fig. 5.4e, f, g enable assessing correlations and distributions by showing the spread of three financial performance measures on the SOM grid: capital, loan and profitability ratios. Here, one can observe that, generally, the right part represents well-performing and the left part poor banks, which gives a direct interpretation to the trajectories in Fig. 5.1.

So, how does the SOM relate to information visualization? Following the discussion about data graphics in Sect. 4.1, the SOM can be related to Bertin's (1983) framework. The plane, and its two dimensions (x , y), are described as the richest variables, which can be perceived at all levels of organization. On the SOM, they represent discrete neighborhood relations. This corresponds also to the key aim of the SOM, that is, to preserve neighborhood relations, whereas global distance structures are of secondary importance. The retinal variables, and their three types of implantation (point, line and area), are thus positioned on the grid. The six retinal variables may be used to represent properties of the SOM grid, particularly properties of the units. To refresh memory, they are as follows (where the parenthesis refers to Bertin's levels of organization): *size* (ordered, selective and quantitative), *value* (ordered and selective), *texture* (ordered, selective and associative), *color* (selective and associative), *orientation* (associative, and selective only in the cases of points and lines), and *shape* (associative). The choice of retinal variable should be based upon the purpose of the visualization and the type of data to be displayed. For instance, variation in size has been used to represent frequency of data in units [see, e.g., Resta (2009)]. Value, or brightness, has been used to visualize the spread of univariate variable values (i.e., feature planes) on the SOM (see, e.g., Fig. 5.4). Likewise, texture has been used for representing cluster memberships [see, e.g., Sarlin (2012a)]. Orientation is commonly applied to represent high-dimensional reference vectors by the means of arrows [see, e.g., Kohonen (2001, p. 117)]. Variation in color (or hue) has been used for illustrating crisp cluster memberships (see, e.g., Fig. 5.4) and for a coloring that reveals multivariate cluster structures [see, e.g., Kaski et al. (2001) and Sarlin and Rönqvist (2013)]. Variation in shape is commonly used on the SOM by the means of labels, such as phoneme strings and phonemic symbols (Kohonen 2001, pp. 208–210).

5.6 Discussion

This chapter has considered data and dimension reduction methods, as well as their combination, for visual financial performance analysis. The discussions and illustrations in this chapter, while being at times somewhat trivial, are motivated by inconsistency of argumentation for and application of various methods. The main conclusion of the comparison is that the SOM has several useful properties for financial performance analysis. In particular, this chapter has noted the following advantages of the SOM over alternative distance-preserving methods:

- (i) trustworthy neighbors,
- (ii) low computational cost,
- (iii) flexibility for problematic data, and
- (iv) a regularly shaped grid.

So, is the superiority of the SOM supported by information visualization theories? Indeed, the SOM representation can be related to Tufte's (1983) advise and principles on graphical clarity and precision. Due to a potential loss of information when projecting from a high-dimensional space to one of a lower dimension, trustworthy neighbors clearly relates to Tufte's advise on avoiding distortions of data (given some losses in detail). Furthermore, the regular, predefined grid shape of the SOM enables and facilitates many types of information linking to the same grid structure. This functions as an aid in thinking about the information rather than the design and encourages the eye to compare data. The SOM's property of approximating the probability density functions of data also facilitates presenting vast amounts of data in a small space, as units will be located in dense areas of the data space, which could also be thought of as an aid in making large data sets coherent. On the SOM, data may be revealed at multiple levels of detail ranging from overview of multivariate structures on the grid, to illustration of individual data on the grid (e.g., trajectories located in their BMUs), which also integrates statistical and verbal descriptions. Along these lines, Tufte's six guidelines on telling the truth about data are also supported. For instance, showing data variation, not design variation, and not showing data out of context relates to, and is supported by, the use of a regular grid shape. Likewise, an example of visuals being directly proportional to the quantities they represent is the adjustment of color scales used for the linked visualizations, such as normalizations of feature plane scales in order for all variables to be comparable (see, e.g., Sect. 6.2.2), and the use of perceptually uniform color scales, such as CIELab (1986).

It is, however, worth noting that the relative goodness of a method depends always on the task in question. That said, the SOM is obviously far from a *panacea* for all sorts of data and dimension reduction. When only attempting stand-alone tasks, it is indeed very likely that there exists better methods than the SOM. Similarly, when attempting DDR, the superiority of one method over others depends entirely on the aims of the task in question.

Even though the SOM has been assessed as advantageous for visual financial performance analysis, it is worth to carefully consider its limitations:

- (i) The SOM performs a crude mapping. Rather than data points, the SOM attempts to embed the reference vectors, a significant constraint if detail is of central importance and/or if only projecting a few data points.
- (ii) The regular grid shape sets some restrictions on the SOM. For instance, it may cause interpolating sparse locations with idle units, it may lead to an analyst overinterpreting the regular-like y and x axes, and leads to the need for additional visual aids to fully represent structures.
- (iii) Mathematical treatment of the SOM has shown to be problematic. The lack of an objective function, as well as a general training schedule for or proof of convergence, complicates parametrizing a SOM.

The comparison in this section has covered classical first-generation dimension reduction methods. This leads to one key question: *Can the results of this comparison be generalized to all available methods?* As reviewed in Sect. 5.1, CCA has been shown to outperform the SOM in terms of trustworthiness of neighborhood relations (Himberg 2004; Venna and Kaski 2007). Likewise, two more recent local versions of MDS, denoted LMDS, by Venna and Kaski (2006) and Chen and Buja (2009) adapt the functioning of standard MDS to preserve local relations. These methods, while holding promise for one criterion, fall short in other, not the least in the shape of the output. It is thus important to consider methods from the second generation with the key properties of the SOM. There are two conceptually similar topology-preserving methods that possess the capabilities of the SOM and a predefined grid shape: GTM and XOM. GTM mainly differs from the SOM by relying on well-founded statistical properties. It is based upon Bayesian learning with an objective function, namely the log-likelihood, which is optimized by the Expectation-maximization algorithm. The objective function directly facilitates assessing convergence of the GTM. Even though Bishop et al. (1998) originally stated that the GTM is computationally comparable to the SOM, it has later been shown that the SOM is cheaper (e.g., Rauber et al. 2000). This may result from the number of developed algorithmic shortcuts for computing SOMs, such as fast-winner search (Kaski 1999). Both methods are flexible for problematic data, i.e., outliers and missing values, through a similar predefined grid shape and an extension of the GTM for treating missing values (Carreira-Perpiñan 2000; Sun et al. 2001). However, while choosing parameters for the SOM may be a tedious task, given adequate initializations and parametrization, convergence has seldom appeared to be a problem in practice (see, e.g., Yin 2008). A decade after the introduction of the GTM, neither it nor its variants, such as the S-Map (Kiviluoto and Oja 1997), have displaced the standard SOM.

The XOM is a computational framework for data and dimension reduction. By inverting the functioning of the SOM, the XOM systematically exchanges functional and structural components of topology-preserving mappings by self-organized model adaptation to the input data. It has two main advantages compared to the SOM: (i) reduced computational cost, and (ii) applicability to non-metric data as there is no restriction on the distance measures. Even though the use of non-metric dissimilarity measures is of little use on the data in these particular examples, while still having potential for other pairwise financial data, the reduced computational cost is

particularly beneficial for large financial datasets in general. The XOM has, however, been recently introduced and is thus still lacking thorough tests in relation to other methods, such as comparisons to SOMs with algorithmic shortcuts. Yet, the XOM should be considered as a valid alternative to the SOM paradigm.

The key message is thus that all four criteria are fulfilled by three methods that perform a topology-preserving mapping to a regularly shaped grid: the SOM, GTM and XOM. It is worth noting, as widely suggested (e.g., Lee and Verleysen 2007; Trosset 2008), that one of the main reasons for the SOM being very popular for a broad range of tasks, such as classification, clustering, visualization, prediction, missing value imputation, etc, might be because it produces an intuitive output using a simple and easily understandable principle. This simplicity, while being beneficial for a method to be widely accepted, applied and understood, should still not be used for assessing relative goodness. One should, nevertheless, note that when introducing dimension reductions to the general public, such as policy- or decision-makers in general, simplicity is definitely an asset. To this end, the most suitable method for financial performance analysis is one from the family of methods that perform a topology-preserving mapping to a regularly shaped and predefined grid. In the work in this book, out of the above described family of methods, the choice of the SOM is motivated by the simplicity of and large number of extensions provided to the SOM.

5.7 Concluding Summary

The literature shows a lack of unanimity on the superiority of one dimension reduction method over others. Yet, every task has its own needs. Data and dimension reduction for financial performance analysis should thus be performed with methods that have the best overall suitability for the performed task, not the best processing capabilities for some other objective. To this end, this chapter has addressed the choice of method for visual financial performance analysis from a qualitative perspective. We have first discussed the properties of three inherently different classical first-generation dimension reduction methods, and their combination with data reduction, and illustrated their performance in a real-world financial application to benchmarking European banks. The conclusions drawn from the comparison of classical methods was then prolonged to second-generation methods. The qualitative discussion and experiments showed superiority of the SOM for financial performance analysis in terms of four criteria: form of structure preservation, computational cost, flexibility for problematic data and shape of the output. When considering second-generation methods, the recently introduced GTM and XOM have clear potential for similar tasks. GTM improves the SOM paradigm with its well-defined objective function, but is computationally more costly, whereas XOM is a recently introduced promising method, but lacks still thorough comparisons.

From the discussions in this chapter, an obvious conclusion is that the family of methods that perform a topology-preserving mapping to a regularly shaped and predefined grid provides means for visual financial performance analysis. The aims




Fig. 5.5 An exemplification of linking information to a Sammon’s mapping. *Notes* The figures link additional information to the coordinates of the Sammon’s mapping. All three plots show the spread of three individual variables measuring financial performance (i.e., feature planes): capital, loan and profitability ratios. They are comparable to the feature planes of the SOM grid shown in Fig. 5.4d–f. The reader is referred to these scales for an interpretation of the color scale

and needs for the task at hand, where the main focus lies on using the output as a display for additional information in general and individual data in particular, are neither rare objectives in other fields. While not being generalizable to their full extent, parts of the conclusions herein will also apply in other fields, domains and tasks. The methods advocated in this book do obviously not provide a *panacea* for visual financial performance analysis. They should be paired with other methods, not least visualizations of different kinds, that compensate for missing properties when having, for instance, a regularly shaped grid. To this end, the chapter also motivates exploring the information commonly linked to the SOM in not only the same family of methods with predefined grid shapes, but also other dimension reduction paradigms in general. Figure 5.5 exemplifies how “feature planes” for a Sammon’s mapping visualize the spread of individual variables for the Sammon’s mapping coordinates.

To sum up, the SOM was found to hold most promise for the task performed in this book, which also sets the direction in the sequel of this book. Yet, the standard SOM as such is not always enough for the task at hand. In the following chapter, we will discuss how the SOM can be extended to better meet the aims and needs for the tasks and data at hand.

References

- Bação, F., & Sousa Lobo, V. (2005). Self-organizing maps as substitutes for k-means clustering. *Proceedings of the International Conference on Computational Science (ICCS 02)* (pp. 476–483). Amsterdam: The Netherlands.
- Balakrishnan, P., Martha, C., Varghese, S., & Phillip, A. (1994). A study of the classification capabilities of neural networks using unsupervised learning: a comparison with k-means clustering. *Psychometrika*, 59, 509–525.
- Bertin, J. (1983). *Semiology of graphics*. WI: The University of Wisconsin Press.
- Bezdek, J. C., & Pal, N. R. (1995). An index of topological preservation for feature extraction. *Pattern Recognition*, 28(3), 381–391.
- Bishop, C., Svensson, M., & Williams, C. (1998). Developments of the generative topographic mapping. *Neurocomputing*, 21(1–3), 203–224.

- Carreira-Perpiñan, M. (2000). Reconstruction of sequential data with probabilistic models and continuity constraints. In S. Solla, T. Leen, & K. Müller (Eds.), *Advances in neural information processing systems* (Vol. 12, pp. 414–420). MIT Press MA: Cambridge.
- Chen, L., & Buja, A. (2009). Local multidimensional scaling for nonlinear dimension reduction, graph drawing, and proximity analysis. *Journal of the American Statistical Association*, 104, 209–219.
- CIELab. (1986). Colorimetry. *CIE Publication, No. , 15, 2*.
- Cottrell, M., & Letrémy, P. (2005). Missing values: processing with the kohonen algorithm. *Proceedings of Applied Stochastic Models and Data Analysis (ASMDA 05)* (pp. 489–496). France: Brest.
- de Bodt, E., Cottrell, M., & Verleysen, M., (1999). Using the Kohonen algorithm for quick initialization of simple competitive learning algorithms. In *Proceedings of the European Symposium on Artificial Neural Networks (ESANN 99)*. Bruges, Belgium.
- de Vel, O., Lee, S., & Coomans, D. (1996). Comparative performance analysis of non-linear dimensionality reduction methods. In D. Fischer & L. H-J. (Eds.), *Learning from data: Artificial intelligence and statistics* (pp. 320–345). Heidelberg, Germany: Springer.
- Deakin, E. (1976). Distributions of financial accounting ratios: some empirical evidence. *The Accounting Review*, 51, 90–96.
- Denny Squire, D., 2005. Visualization of cluster changes by comparing self-organizing maps. In: *Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD 05)*. Hanoi, Vietnam, pp. 410–419.
- Duch, W., & Naud, A. (1996). Multidimensional scaling and kohonen's self-organizing maps. *Proceedings of the Conference on Neural Networks and their Applications (CNNA 16)* (pp. 138–143). Poland: Szczyrk.
- Flexer, A. (1997). Limitations of self-organizing maps for vector quantization and multidimensional scaling. In M. Mozer (Ed.), *Advances in Neural Information Processing Systems* (Vol. 9, pp. 445–451). Cambridge, MA: MIT Press.
- Flexer, A. (2001). On the use of self-organizing maps for clustering and visualization. *Intelligent Data Analysis*, 5(5), 373–384.
- Harrower, M., & Brewer, C. (2003). Colorbrewer.org: an online tool for selecting color schemes for maps. *The Cartographic Journal*, 40(1), 27–37.
- Himberg, J. (2004). From insights to innovations: data mining, visualization, and user interfaces. Ph.D. thesis, Helsinki University of Technology, Espoo, Finland.
- Kaski, S. (1997). Data exploration using self-organizing maps. Ph.D. thesis, Helsinki University of Technology, Espoo, Finland.
- Kaski, S., & Kohonen, T. (1996). Exploratory data analysis by the self-organizing map: structures of welfare and poverty in the world. *Proceedings of the International Conference on Neural Networks in the Capital Markets* (pp. 498–507). London: World Scientific.
- Kaski, S. (1999). Fast winner search for som based monitoring and retrieval of high dimensional data. *Proceedings of the IEEE International Conference on Artificial Neural Networks (ICANN 99)* (pp. 940–945). London, UK: IEEE Press.
- Kaski, S., Venna, J., & Kohonen, T. (2001). Coloring that reveals cluster structures in multivariate data. *Australian Journal of Intelligent Information Processing Systems*, 60, 2–88.
- Kiviluoto, K., & Oja, E. (1997). S-map: a network with a simple self-organization algorithm for generative topographic mappings. In M. I. Jordan, M. J. Kearns, & S. A. Solla (Eds.), *Advances in Neural Information Processing Systems* (Vol. 10, pp. 549–555). MIT Press MA: Cambridge.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59–69.
- Kohonen, T. (2001). *Self-organizing maps* (3rd ed.). Berlin: Springer.
- Latif, K., & Mayer, R. (2007). Sky-metaphor visualisation for self-organising maps. In *Proceedings of the International Conference on Knowledge Management (I-KNOW 07)*. Graz, Austria.
- Lee, J., & Verleysen, M. (2007). *Nonlinear dimensionality reduction*. Heidelberg, Germany: Springer, Information Science and Statistics Series.

- Lee, J., & Verleysen, M. (2009). Quality assessment of dimensionality reduction: rank-based criteria. *Neurocomputing*, 72(7–9), 1431–1443.
- Linde, Y., Buzo, A., & Gray, R. (1980). An algorithm for vector quantizer design. *IEEE Transactions on Communications*, 28(1), 702–710.
- Lueks, W., Mokbel, B., Biehl, M., & Hammer, B. (2011). How to evaluate dimensionality reduction? In B. Hammer & T. Villmann (Eds.), *Proceedings of the Workshop on New Challenges in Neural Computation*. Machine Learning Reports: University of Bielefeld, Department of Technology, Frankfurt, Germany.
- van der Maaten, L., & Hinton, G. (2008). Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9, 2579–2605.
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* (pp. 281–297). Berkeley, CA: University of California Press.
- Merkel, D., & Rauber, A. (1997). Alternative ways for cluster visualization in self-organizing maps. In *Proceedings of the Workshop on Self-Organizing Maps (WSOM 97)*. Helsinki, Finland.
- Naud, A., & Duch, W. (2000). Interactive data exploration using MDS mapping. *Proceedings of the Conference on Neural Networks and Soft Computing* (pp. 255–260). Poland, Zakopane.
- Neumayer, N., Mayer, R., Poelzlbauer, G., & Rauber, A. (2007). The metro visualisation of component planes for self-organising maps. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN 07)*. Orlando, FL, USA: IEEE Computer Society.
- Nikkilä, J., Törönen, P., Kaski, S., Venna, J., Castrén, E., & Wong, G. (2002). Analysis and visualization of gene expression data using self-organizing maps. *Neural Networks*, 15(8–9), 953–966.
- Pampalk, E., Rauber, A., & Merkl, D. (2002). Using smoothed data histograms for cluster visualization in self-organizing maps. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN 02)* (pp. 871–876). Madrid, Spain.
- Pözlbauer, G., Rauber, A., & Dittenbach, M. (2005). Advanced visualization techniques for self-organizing maps with graph-based methods. *Proceedings of the International Symposium on Neural Networks (ISNN 05)* (pp. 75–80). Chongqing, China: Springer.
- Pözlbauer, G., Dittenbach, M., & Rauber, A. (2006). Advanced visualization of self-organizing maps with vector fields. *Neural Networks*, 19(6–7), 911–922.
- Purves, D., Augustine, G., Fitzpatrick, D., Hall, W., LaMantila, A., McNamara, J., et al. (Eds.). (2004). *Neuroscience*. Massachusetts: Sinauer Associates.
- Rauber, A., Paralic, J., & Pampalk, E. (2000). Empirical evaluation of clustering algorithms. *Journal of Information and Organizational Sciences*, 24(2), 195–209.
- Resta, M. (2009). Early warning systems: an approach via self organizing maps with applications to emergent markets. In B. Apolloni, S. Bassis, & M. Marinaro (Eds.), *Proceedings of the 18th Italian Workshop on Neural Networks* (pp. 176–184). Amsterdam: IOS Press.
- Samad, T., & Harp, S. (1992). Self-organization with partial data. *Network: Computation in Neural Systems*, 3, 205–212.
- Sammon, J. (1969). A non-linear mapping for data structure analysis. *IEEE Transactions on Computers*, 18(5), 401–409.
- Sarlin, P. (2012a). Chance discovery with self-organizing maps: discovering imbalances in financial networks. In Y. Ohsawa & A. Abe (Eds.), *Advances in Chance Discovery* (pp. 49–61). Heidelberg, Germany: Springer.
- Sarlin, P. (2012b). Visual tracking of the millennium development goals with a fuzzified self-organizing neural network. *International Journal of Machine Learning and Cybernetics*, 3, 233–245.
- Sarlin, P. (2014). Data and dimension reduction for visual financial performance analysis. *Information Visualization* (forthcoming). doi:[10.1177/1473871613504102](https://doi.org/10.1177/1473871613504102)
- Sarlin, P., & Rönnqvist, S. (2013). Cluster coloring of the self-organizing map: An information visualization perspective. In *Proceedings of the International Conference on Information Visualization (iV 13)*. London, UK: IEEE Press.

- Serrano-Cinca, C. (1996). Self organizing neural networks for financial diagnosis. *Decision Support Systems*, 17, 227–238.
- Sun, Y., Tino, P., & Nabney, I. (2001). *GTM-based data visualisation with incomplete data*. Technical Report. Birmingham, UK: Neural Computing Research Group.
- Torgerson, W. S. (1952). Multidimensional scaling: i. theory and method. *Psychometrika*, 17, 401–419.
- Trosset, M. (2008). *Representing clusters: K-means clustering, self-organizing maps, and multidimensional scaling*. Technical Report 08–03. Department of Statistics, Indiana University.
- Tufte, E. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.
- Ultsch, A. (2003b). *U*-matrix: A tool to visualize clusters in high dimensional data*. Technical Report No. 36. Germany: Department of Mathematics and Computer Science, University of Marburg.
- Ultsch, A., & Siemon, H. (1990). Kohonen's self organizing feature maps for exploratory data analysis. In *Proceedings of the International Conference on Neural Networks (ICNN 90)* (pp. 305–308). Dordrecht, the Netherlands.
- Ultsch, A., & Vetter, C. (1994). *Self-organizing feature maps versus statistical clustering methods: A benchmark*, University of Marburg. Research Report. FG Neuroinformatik & Kuenstliche Intelligenz. 0994.
- Ultsch, A. (2003a). Maps for the visualization of high-dimensional data spaces. *Proceedings of the Workshop on Self-Organizing Maps (WSOM 03)* (pp. 225–230). Kitakyushu, Japan: Hibikino.
- Venna, J., & Kaski, S. (2001). Neighborhood preservation in nonlinear projection methods. an experimental study. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN 01)* (pp. 485–491). Vienna, Austria: Springer.
- Venna, J., & Kaski, S. (2006). Local multidimensional scaling. *Neural Networks*, 19, 889–899.
- Venna, J., & Kaski, S. (2007). Comparison of visualization methods for an atlas of gene expression data sets. *Information Visualization*, 6(2), 139–154.
- Vesanto, J. (1999). Som-based data visualization methods. *Intelligent Data Analysis*, 3(2), 111–126.
- Vesanto, J., & Ahola, J. (1999). Hunting for correlations in data using the self-organizing map. *Proceeding of the International ICSC Congress on Computational Intelligence Methods and Applications (CIMA 99)* (pp. 279–285). Rochester, NY, USA: ICSC Academic Press.
- Vesanto, J., & Alhoniemi, E. (2000). Clustering of the self-organizing map. *IEEE Transactions on Neural Networks*, 11(3), 586–600.
- Waller, N., Kaiser, H., Illian, J., & Manry, M. (1998). A comparison of the classification capabilities of the 1-dimensional kohonen neural network with two partitioning and three hierarchical cluster analysis algorithms. *Psychometrika*, 63, 5–22.
- Ward, J. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58, 236–244.
- Yin, H. (2008). The self-organizing maps: background, theories, extensions and applications. In J. Fulcher & L. Jain (Eds.), *Computational intelligence: A compendium* (pp. 715–762). Heidelberg, Germany: Springer.

Chapter 6

Extending the SOM

*As the present now
Will later be past [...]
And the first one now
Will later be last
For the times they are a-changin'.*
–Bob Dylan

The standard Self-Organizing Map (SOM), while having merit for the task at hand, may be extended in multiple directions, not the least to better meet the demands set by macroprudential oversight and data. Chapters 2 and 3, spell out the needs and demands for the task at hand, to be used as a basis for the applications and extensions of methods. As discussed in Chaps. 4 and 5, the method of preference for the purposes in this book is the SOM. A particular focus of the extensions is related to two tasks that not only meet the demands of macroprudential oversight and data, but have also been stated to be in need of future research in the fields of information visualization and dimension reductions. First, Chen (2005) and Wong et al. (2012) highlight a paradigm shift from only visualizing structures to visualizing dynamics. An even further step is to assess dynamics of structures. Second, to be aware of the quality and distortions of dimension reductions, Wismüller et al. (2010) and Wong et al. (2012) stress that they are not an end, but provide only a means to display useful information on top of them, such as evidence, uncertainty and individual data.

Along these lines, with a key focus on temporality, this chapter first discusses the literature on time in SOMs. This is followed by extensions to the standard SOM paradigm. In general, the chapter presents extensions to the SOM paradigm for processing data from the cube representation, i.e., along multivariate, temporal and

This chapter is partly based upon previous research. Please see the following works for further information: Sarlin et al. (2012), Sarlin and Eklund (2011), Sarlin (2013b,d,e), Sarlin and Yao (2013)

cross-sectional dimensions, where a focus of emphasis is on a better processing and visualization of time. The motivation and functioning of the extensions is demonstrated with a number of illustrative examples.

6.1 Time in SOMs: A Brief Review

There is a wide range of literature adapting and extending the standard SOM for temporal processing. While the literature on time in SOMs has been thoroughly reviewed in Barreto (2007), Barreto et al. (2003), Barreto and Araújo (2001), Guimarães et al. (2003) and Hammer et al. (2005), a unanimous classification dividing it into distinct groups of studies is not clear-cut. Drawing upon the above reviews, the literature related to time in SOMs is reduced into four groups of works: (i) those with an implicit consideration of time, (ii) those adapting the learning or activation rule, (iii) those adapting the topology, and (iv) those combining SOMs with other visualization techniques.

The first group *implicitly considers time* by applying the standard SOM algorithm and illustrates the temporal dimension either as a pre- or post-processing step. The pre-processing concerns embedding a time series into one input vector, such as so-called tapped delay [e.g., Kangas (1990)]. A time-related visualization through post-processing is, however, more common. A connected time series of best-matching units (BMUs), i.e., a trajectory, has been used in the literature to illustrate temporal transitions [e.g., Kohonen (1988) and Martín-del Brío and Serrano-Cinca (1993)]. By exploiting the topological ordering of the SOM, visualization of the current and past states enables visual tracking of the dynamics in multivariate data (i.e., process dynamics). However, while temporal patterns require large datasets for generalization and significance, trajectories can only be visualized for a limited set of data. Thus, strengths and actual directions of the patterns can be obtained by probabilistic modeling of state transitions between SOM units [e.g., Sulkava and Hollmén (2003), Luysaert et al. (2004) and Fuertes et al. (2010)].

The second group of works *adapts the standard SOM activation or learning rule*. Those decomposing the learning rule of the standard SOM into two parts, past and future, for time-series prediction have their basis in the Hypermap (Kohonen 1991). The past part is used for finding BMUs, while the entire input vector is used within the updates of the reference vectors. For predicting out-of-sample data, the past part is again used for finding BMUs while the future part of that unit is the predicted value. This type of learning has been used for standard time-series prediction [e.g., Principe and Wang (1995) and Ultsch et al. (1996)] and predictions through non-linear regression [e.g., Sarlin and Marghescu (2011)]. As noted in Sect. 4.4, the latter type of decomposition can still be divided into semi-supervised and unsupervised SOMs, where the difference depends on whether or not the present part is used for matching in training. Instead of considering the context explicitly in SOM training, it can be treated as the neighborhood of the previous BMU. Kangas (1992), for

instance, constrains the choice of a BMU to the neighborhood of the previous BMU and thus has a behavior that resembles the functioning of SOMs with feedback in the next group.

The *third* group deals with *adaptations of the standard SOM network topology* through feedback connections and hierarchical layers. The feedback SOMs have their basis in the seminal Temporal SOM (TSOM) (Chappell and Taylor 1993), also called the Temporal Kohonen Map, that performs leaky integration to the outputs of the SOM. The Recurrent SOM (RSOM) (Varsta et al. 1997; Koskela et al. 1998) differs by moving the leaky integration from the output units to the input vectors. A recent recurrent model is the Merge SOM (MSOM) (Strickert and Hammer 2005) whose context combines the current pattern with the past by a merged form of the properties of the BMU. The Recursive SOM (RecSOM) (Voegtlin 2002) keeps information by considering the previous activation of the SOM as part of the input to the next time unit, while the Feedback SOM (FSOM) (Horio and Yamakawa 2001) differs by integrating an additional leaky loop onto itself. The SOM for structured data (SOMSD), on the other hand, labels directed acyclic graphs to regular (Hagenbuchner et al. 2003) and arbitrary (Strickert et al. 2005) grid structures. Finally, Hammer et al. (2004) define a general formal framework and show that a large number of SOMs with feedback can be recovered as special cases of the framework. The hierarchical network architectures, on the other hand, use at each layer one or more SOMs operating at different time scales. The next level in the hierarchy can either use the lower level SOM as input vectors without any processing, such as two-level clustering commonly does, or use transformed input vectors by computing distances between units or concatenating a time series to one input vector, for instance. Kangas (1990) introduced hierarchical network architectures to SOMs, and shows that a hierarchical SOM without any additional processing outperforms SOMs with backwards averaged and concatenated input vectors.

The *fourth* group of studies attempts to create *SOM-based visualization tools* for exploratory analysis of data by combining the SOM with other methods and interactive interfaces. The particular focus of these tools is to provide means for dealing with spatiotemporal data. Standard SOMs using both cross-sectional and temporal data have, in addition to trajectory and state-transition analysis, been paired with stand-alone visualization aids for a spatial mapping [e.g., Kaski et al. (2001)]. Guo et al. (2006) introduces an integrated approach of computational, visual and cartographic methods for visualizing multivariate spatiotemporal patterns, where parallel coordinate plots and reorderable matrices enhance the information products of the SOM. The visualization tool created by Andrienko et al. (2010) extends the one in Guo et al. (2006) by not only grouping spatial situations as per time units, but also spatial locations as per temporal variations. Further, a SOM-based visualization tool for temporal knowledge discovery is introduced in Guimarães (2000) and Guimarães and Ultsch (1999). The tool presents a hierarchical SOM to handle complexity, and includes a U-matrix visualization, trajectory analysis and a transformation of data into linguistic knowledge.

These four groups of extensions, while covering a wide range of temporal processing, leave room to provide better means for visualizing patterns in macroprudential data. This is the topic of the following sections.

6.2 Extensions for Exploiting the SOM

Directed by the task at hand, this section takes the standard SOM as a basis and then aims at extending it, as well as combining it with other methods, to better meet the needs and demands of macroprudential oversight and data. The section proposes three extensions: a fuzzification; transition probabilities; and shock propagation assessment. The functioning of all three extensions is demonstrated with simple examples on the bank SOM used in Chap. 5.

6.2.1 Fuzzification of the SOM

In the early days, information extraction on the SOM was mainly facilitated by visual analysis of some form of the U-matrix [e.g., Ultsch and Siemon (1990)], where a color code between all neighboring units indicates their average distance. The SOM units have also been used as input for a second stage (or hierarchy) of two-level clustering. However, one source of ambiguity with the SOM clustering is that the degree of membership in a particular cluster is not always easy to judge. Although location on the SOM represents closeness, the distance structure on the SOM is most often not uniform. In some cases, it might be beneficial to judge the degree to which a particular area of a cluster (i.e., one or more units) differs from the rest of the cluster, and what its closest match among the other clusters is. While the SOM is commonly partitioned using a crisp clustering technique [e.g., Vesanto and Alhoniemi (2000)], one solution to judging membership degrees is to fuzzify the SOM. The motivation for using a fuzzification is threefold:

- (i) For monitoring belongingness of individual data (e.g., over time),
- (ii) For assessing distance structures of the SOM units with respect to clusters; and
- (iii) For assessing topological ordering of the SOM.

The fuzzification can take various forms. Below, we look at three possible approaches: Fuzzy c-means (FCM) clustering, distance-based fuzzification and class and distance-based fuzzification.

Fuzzy c-Mean Clustering

The FCM algorithm, developed by Dunn (1973), Bezdek (1981), may be employed for assigning a degree of membership of each unit in each of the clusters, as suggested

in Sarlin and Eklund (2011). This provides a fuzzified representation of the SOM. Following the presentation of FCM clustering in Sect. 4.3, the objective function J_θ can also be applied to SOM units. The objective function J_θ is thus defined as the weighted sum of the Euclidean distances between each unit and each cluster center, where the weights are the degree of memberships of each unit in each cluster, and again constrained by the probabilistic requirement:

$$J_\theta = \sum_{i=1}^M \sum_{k=1}^C u_{ik}^\theta \|m_i - c_k\|^2, \quad \sum_{k=1}^C u_{ik} = 1, \quad (6.1)$$

where $\theta \in (1, \infty)$ is the fuzzy exponent, u_{ik} is the degree of membership of reference vector m_i (where $i = 1, 2, \dots, M$) in the cluster center c_k (where $k = 1, 2, \dots, C$, and $1 < C < M$), and $\|m_i - c_k\|^2$ is the squared Euclidean distance between m_i and c_k . After a random initialization, it optimizes the cluster centers c_k and the membership values u_{ik} with the same Picard iteration through Eqs. (4.6) and (4.7) as was shown in Chap. 4. Thus, it applies the same procedure, but instead directly on the reference vectors m_i .

Distance-Based Fuzzification

Instead of using FCM clustering on the units, Sarlin and Eklund (2013) compute the membership degrees directly using Euclidean distances between SOM units (or data) and the centroids of crisp clusters. For this, any crisp clustering method, as appropriate, is applicable. The crisp clustering is fuzzified by computing the inverse distance between reference vector m_i (or each data point x_j) and each cluster center c_k :

$$v_{ik} = \frac{1}{1 + \|m_i - c_k\|^{\frac{2}{\theta-1}}} \quad (6.2)$$

where $\theta \in (1, \infty)$ is again the fuzzy exponent (i.e., the fuzzifier) which controls the extent of overlap between the clusters. However, the similarity matrix v_{ik} is normalized to the following cluster membership matrix for each unit:

$$u_{ik} = \frac{v_{ik}}{\sum_{k=1}^C v_{ik}} \quad (6.3)$$

to fulfill the probabilistic constraint $\sum_{k=1}^C u_{ik} = 1$. The extent of overlap between the clusters is set by the fuzzy exponent θ . When $\theta \rightarrow 1$, the fuzzy clustering converges to a crisp clustering, while when $\theta \rightarrow \infty$ the cluster centers tend towards the center of the data set. $\theta = 2$ and $\theta = 3$ can be seen as benchmarks, since they give squared and simple Euclidean distances. The fuzzy exponents in the above and below approaches also follow these guidelines.

This approach resembles that in Cottrell and Letrémy (2005), but differs by being implemented on a second-level clustering instead of directly on the units, by not assuming inverse exponential distances, and by introducing a fuzzification parameter. More importantly, Cottrell and Letrémy (2005) use the derived memberships for imputing missing values. As the fuzzification is implemented on the units, it can be used for assessing the topological ordering and distance structure of the grid. While FCM clustering necessitates visualizing memberships of individual data according to those of their BMUs, this approach enables one to also compute them for individual data. This is particularly important as one-unit movements may be switches between clusters, but still changes in data may be minor. A plot of the memberships would capture this.

Class and Distance-Based Fuzzification

In cases when one possesses class information in data, it is not necessary to estimate clusters and their centroids. They can be derived from their distribution on the SOM. Thus, one might not only have class information, but also utilize a semi-supervised SOM with the classes in the ordering process. Following Sarlin (2013b), the above fuzzification of the SOM is adapted for computing class memberships. Let the input data consist of two parts: class vector $x_{j(cl)}$ and input vector $x_{j(in)}$. The SOM can be classified and fuzzified based upon the class vectors $x_{j(cl)}$ by assuming the following: the number of clusters C equals the number of classes K , i.e., $C = K$, and the cluster center c_k (where $k = 1, 2, \dots, C$) for each class is a perfect representative state vector, i.e.,

$$\forall k \quad c_k = \begin{cases} 1 & \text{if } k \text{ equals the state of } c_k \\ 0 & \text{otherwise} \end{cases} \quad (6.4)$$

While there exist other methods for class visualization on the SOM, such as Voronoi regions (Mayer et al. 2007), they fall short in dealing with imprecision in class memberships. Following the approach in Eqs. (6.2) and (6.3), we can compute a membership degree using Euclidean distances between units and state centers, but only use $x_{j(cl)}$ and $m_{i(cl)}$ for measuring these distances. The rationale for this is the focus on distances between mean profiles of classes $x_{j(cl)}$ rather than those between inputs $x_{j(in)}$. The SOM is fuzzified by computing the inverse distance between reference vector $m_{i(cl)}$ and each state center $c_{k(cl)}$, as in Eq. (6.2), and normalized as in Eq. (6.3) to fulfill the probabilistic constraint. In addition to computing membership degrees, one can also apply a defuzzification of the results using the maximum-membership method. This enables deriving crisp clusters of reference vectors such as in two-level clustering, which is also applicable for the FCM clustering.

Visualizing a Fuzzification

Similarly as feature planes for individual variables, membership degrees can be associated to each of the SOM units, and linked to the SOM grid, where one unique point represents the same unit on a SOM. Thereby, the structure of the clusters of a SOM model can be identified by studying so-called membership planes. They show the degree of membership in cluster k for each unit m_i on an own grid, such that the color code of each unit m_i represents its membership in cluster k . The temporal dimension of an individual entity can also be represented by computing for each data point an own membership degree in each cluster. This enables a line graph representation of the state switch probabilities over time for individual data, where cluster centers express representative states and variations of membership degrees represent their fluctuation over time.

An Illustrative Example

To illustrate the functioning and usefulness of fuzzifications, we turn to illustrative examples on simple real-world data. The SOM model created on data for European banks, when comparing methods in Chap. 5, can be used for illustrating the fuzzification. As the SOM is entirely unsupervised and already uses Ward's hierarchical method for a second-level clustering, the distance-based fuzzification is applied. Thus, the crisp clustering of the above SOM model is fuzzified using Euclidean distances. The fuzzifier $\theta \in (1, \infty)$ was tested for values in $(1, 10]$. Based upon these experiments, a benchmark θ -value of 2.0 provided an adequate fuzzification of the map. It introduces a fuzziness degree large enough to show relationships between clusters, but small enough not to completely eliminate cluster borders. The computation provides the membership of each unit in each cluster, as well as the memberships of each data point in each cluster.

Figure 6.1 shows membership planes, in each of which the membership degrees of all units in one cluster are visualized, and finally a membership plane to illustrate the crisp clusters and the trajectories of UniCredit Banca and ING Bank. The figure illustrates the crispness of the clusters, and locations in general. For instance, while the cluster center of cluster D is located in its lower part, where cluster memberships are somewhat crisp, one can observe the opposite for units on the borders between clusters A, D, E and G. In broad terms, as the memberships overall decrease over distances from cluster centers on the grid, the membership planes indicate no major concerns with topological ordering.

The line graphs in Fig. 6.2 illustrate an assessment of individual time-series points that have a partial membership in all identified, but overlapping, clusters. UniCredit Banca switches between clusters A, D, E and G, which is also illustrated by the low membership degrees. This functions as a particular motivation for using memberships, as only tiny differences in the underlying data may lead to switches between clusters. On the contrary, strong membership degrees can be observed in the case of ING, where the final movement to a unit that borders cluster D actually increases

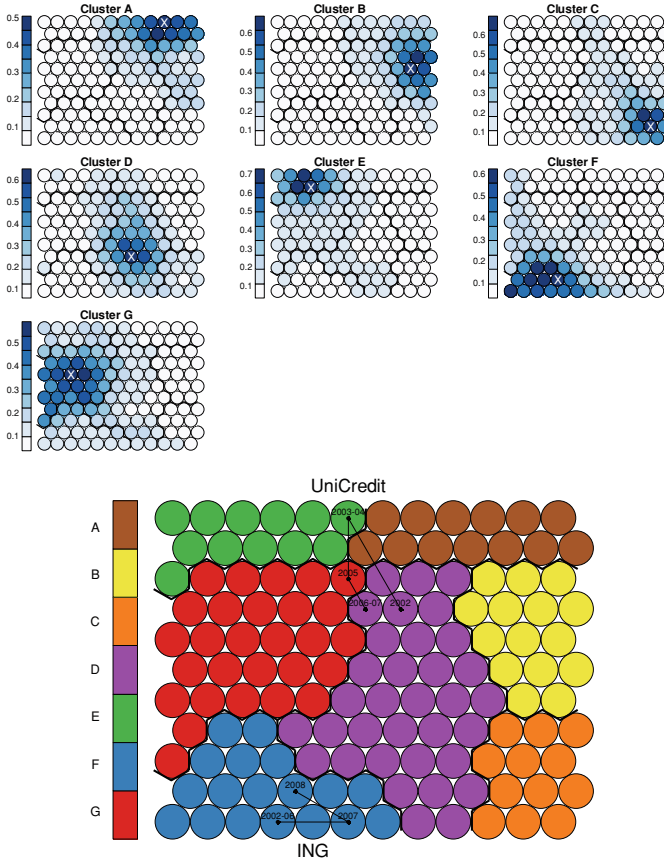


Fig. 6.1 An exemplification of a fuzzification of the SOM. *Notes* The figure links membership degrees to the SOM. The first seven *grid* representations are called membership planes, as each of the planes visualizes the membership degrees of all units to one cluster. The final *grid* is a crisp membership plane that shows the same *cluster* memberships as those displayed in Fig. 5.4. The crisp membership plane also overlays trajectories for UniCredit Banca and ING Bank

the membership in cluster F. Hence, additional aids are needed to understand the implications of movements on the SOM.

6.2.2 Transition Probabilities on the SOM

The SOM has been shown to be an ideal tool for building low-dimensional displays for the visualization of individual data. However, manually identifying the temporal patterns in a SOM model is not necessarily a simple process. As is already appar-




Fig. 6.2 Line graphs of a fuzzified SOM. *Notes* The figure visualizes times series of membership degrees for two banks, UniCredit Banca and ING Bank

ent from Chap. 3, as well as the above example, data are commonly drawn from a three-dimensional data cube, including the multivariate, cross-sectional and temporal dimensions. These types of data are not unique. Also in other fields than accounting, finance and economics, such as process monitoring [see Alhoniemi et al. (1999) and Fuertes et al. (2010)], it is more common than not for multivariate data to include both a temporal and cross-sectional dimension. For instance, when data are cyclical (or scarce), one may want to build a standard SOM model with data on several entities over time to include both the temporal and cross-sectional differences. Given a model with this type of data, an obvious interest would be the temporal properties of the model.

The standard SOM paradigm does not, however, explicitly address the issue of temporality. Variations of the SOM algorithm itself, as reviewed in Sect. 6.1, have been proposed for dealing with temporal data. Oftentimes, these extensions, however, turn their focus from the entities to the sequences, or otherwise enhance time-series prediction. The extensions, while holding promise for a wide range of other tasks, do not provide means for visualizing the temporal structure on a standard SOM. While trajectories have been a common means to illustrate temporal movements, as has already been illustrated in this book, small samples give no indication of overall patterns and large samples clutter the display. Thus, trajectories and the above presented fuzzification provide no overall information about trends in the dataset. For finding these patterns, be they cyclical or not, movements should be summarized from transition probabilities, something that is not apparent from only studying the elements of the SOM units. Transition probability matrices (TPMs) can be used to produce a probabilistic model of the temporal variation in a SOM model. This has been introduced through unit-to-unit transition probabilities (Sulkava and Hollmén 2003; Luysaert et al. 2004; Fuertes et al. 2010). These types of transition probabilities generalize the strengths and actual directions of the temporal patterns on the SOM.

Three Approaches to Transitions

This subsection presents the approach to transition probabilities put forward in Sarlin et al. (2012) by focusing on unit-to-cluster switches on SOM with a second-level clustering. Thus, the below presented framework provides better means

to compute, summarize and visualize strengths and actual directions of transition probability patterns.

Movements between units on the two-dimensional SOM are used to compute probabilities of switching from a unit to a specified region in a specified time period, where the location of data per time unit is their BMU (see Eq. 4.9). First, we compute for each unit m_i the probability of transition to any other unit m_u :

$$p_{iu}(t+s) = \frac{n_{iu}(t+s)}{\sum_{u=1}^M n_{iu}(t+s)} \quad (6.5)$$

where n_{iu} is the cardinality of data switching from m_i to m_u , t is a time coordinate and s is the time span for the switch. In other words, the transition probability $p_{iu}(t+s)$ equals the cardinality of transitions from unit m_i to unit m_u divided by the sum of transition from unit m_i to $m_{1,2,\dots,M}$. On a SOM grid with four units, this could in practice mean that for, say, unit m_1 the probability of being in period $t+1$ in $m_{1,2,\dots,4}$ could be 0.5, 0.2, 0.2 and 0.1, respectively. More formally, a TPM corresponds to a stationary first-order Markov model or maximum-likelihood estimates of the switches (Anderson and Goodman 1957). It can, however, be computed for different time spans, as appropriate, and summarized to switches between clusters or any other region on the map. For example, unit-to-cluster switches are computed using p_{il} , where the transition refers to movements from reference vector i to cluster l (where $l = 1, 2, \dots, C$), thus:

$$p_{il}(t+s) = \frac{n_{il}(t+s)}{\sum_{l=1}^C n_{il}(t+s)} \quad (6.6)$$

For larger samples, and thus more robust results, the TPMs p_{il} (as well as p_{iu}) can be computed as an average of several s values (where $s = 1, 2, \dots, S$):

$$p_{il}(t + \{1, 2, \dots, S\}) = \frac{\sum_{s=1}^S n_{il}(t+s)}{\sum_{s=1}^S \sum_{l=1}^C n_{il}(t+s)} \quad (6.7)$$

Thus, the following three computations are proposed:

- (i) TPMs for unit-to-cluster switches ($p_{il}(t+s)$) as in Eq. (6.6) for a specified set of s values.
- (ii) Summarize the TPMs from Step (i) by computing to which cluster l an observation in m_i is most likely to switch and with what likelihood, i.e., showing maximum transition probabilities ($\max_l(p_{il})$) conditional on switching. This combines the direction and strength of all probabilities into one vector.
- (iii) For summarizing the computations in Steps (i) and (ii) over time, compute average transition probabilities over a chosen set of s values ($p_{il}(t + \{1, 2, \dots, S\})$) as in Eq. 6.7).

Visualizing the Transitions

Similarly as membership planes, transition probabilities can be associated to the SOM units, and linked to the SOM grid. Thereby, the structure of the transitions on the SOM model can be directly identified by studying these so-called transition planes. The above computations are represented using the following three visualizations:

- (i) *Transition planes* show the probability to transit to cluster l for each unit m_i on an own grid, such that the color code of each unit m_i represents its probability of transition to cluster l .
- (ii) *Summarized transition planes* aggregate the transition planes for all C clusters to one grid by using a color code for m_i to represent the probability of the most likely switch and a label to represent that cluster.
- (iii) Create the same feature planes as in Steps (i) and (ii), but as an average over a chosen set of s values.

To normalize the color scales for different cluster sizes, but still show differences over time spans, the color scales of the feature planes for all s values and sets of s values can be specified as to that for the shortest time span $\min(s)$ (e.g., $s = 1$ and $t + 1$). The temporal dimension of an individual entity can as well be represented by associating each time-series point with the transition probability of its BMU (Eq. 4.9). This enables a line graph representation of the state switch probabilities over time for individual data, where clusters are representative states and the variation in transition probabilities represent changes in indications of future characteristics. The transition probabilities can also be used for profiling by presenting characteristics of low- and high-risk mean profiles based upon future transitions.

An Illustrative Example

The same model based upon European banks is also used to illustrate transition probabilities on the SOM. Thus, we follow the above three-step framework when computing the transition probabilities. First, TPMs are computed as switches from units to clusters ($p_{il}(t + s)$ as in Eq. 6.6). Second, the direction and strength of the switches are summarized by computing maximum transition probabilities $\max_l(p_{il})$ conditional on switching. Third, the above steps are computed for three different transition time spans ($t + 1$, $t + 2$ and $t + 3$) and an average for $S = 3$. However, for the sake of brevity, Fig. 6.3 only visualizes the average of all three time spans. The illustrated transitions on the SOM may be utilized for exploring patterns of interest. Clusters E, F and G can be seen as inherently stable, as there are few transitions from the units in these clusters. Clusters A, B and C, on the other hand, are less stable. Cluster D is an unstable transition cluster. Further, we can see that banks in clusters A, F and G are quite stable, while clusters B, C and E exhibit more transitions.

The differences in stability between cluster A and clusters B and C might be due to differing business activities, as an inspection of the feature planes (partially shown in Fig. 5.4) illustrate that cluster A differs from B and C primarily in capital ratios,

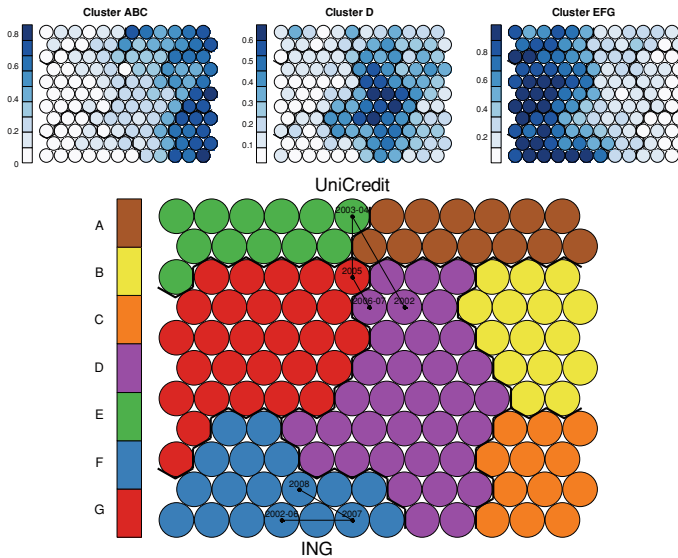


Fig. 6.3 An exemplification of transition probabilities on a SOM. *Notes* The figure links transition probabilities information to the SOM. The first three *grids* are so-called transition planes, as each plane visualizes the transition probability of all units to one *cluster*. The final *grid* is the same crisp membership plane shown in Fig. 5.4. The membership plane also overlays trajectories for UniCredit Banca and ING Bank

loan interest revenue and subordinated debt. This indicates that clusters B and C are higher risk clusters than cluster A, and thus probably more sensitive to changes in the business environment, such as interest rates and other macro-financial conditions. For cluster B, an interesting strong cluster-to-cluster pattern is the high probabilities of movements to cluster D. Another interesting pattern is the difference in stability between cluster E and clusters F and G. While E, F and G are quite similar clusters in terms of performance, a clear difference can be seen in the high ratio of non-operating items of cluster E. Non-operating items are items not related to ongoing, day-to-day operations, such as dividends, financial investments or significant write-downs, which might partially explain the unstable nature of positions in cluster E.

The line graphs in Fig. 6.4 show a practical bank-specific application of the transition-probability framework. The figure shows the state transition probabilities for UniCredit Banca and ING Bank for 2002–2008. If one is interested in likely future switches, the addresses of the switches and the probability trend of the most likely switch should be assessed, as the probability of staying in a cluster is mostly highest. The patterns for the two case banks resembles that of the fuzzification application. The transition probabilities for UniCredit, who also switches cluster frequently, are spread out in all three cluster groups, whereas those of ING are shown to be dominated by the only state it is a member of, cluster F.




Fig. 6.4 Line graphs of transition probabilities for SOM trajectories. *Notes* The figure visualizes times series of transition probabilities for two banks, UniCredit Banca and ING Bank

6.2.3 Shock Propagation on the SOM

The process of some event being transmitted to another entity goes by different names, such as contagion, shock propagation or spread of an event. On a SOM, this occurrence can be analyzed with two approaches: links between entities and similarities in inputs. This subsection presents the approaches put forward in Sarlin (2013b).

Networks on the SOM

While most thus far discussed tasks have utilized data from the cube representation, they have disregarded the fourth dimension of linkages. The first approach superimposes a cross-sectional network of bilateral links on the SOM. Network analysis, or link analysis, can be seen as the exploration of crucial relationships and associations between a large set of objects that may not be apparent from assessing isolated data. Networks of relationships are mostly expressed in matrix form, where the link between entities g and l in a matrix A is represented by element a_{gl} . The matrix is of size n^2 , where n is the number of entities. Matrices of directed graphs can be read in two directions: rows of A represent the relationship of g to l and columns of A represent the relationship of l to g . To combine SOMs and network analysis, network relations are superimposed on top of the standard SOM grid by visualizing relationships between entities. Labels of entities under analysis, say g and l , are projected to their BMUs on the SOM based upon their data x_g and x_l . After that, relations between entities g and l are visualized by edges between the locations of the BMUs of x_g and x_l on the SOM grid using elements a_{gl} and a_{lg} . This visualizes simultaneously the data topology of the SOM and a network topology of pure data relationships. While these two topologies have thus far been mainly assessed in isolation, they are oftentimes highly interrelated as changes in one of the topologies may have significant implications on the stability of the other or the combined topology.

Figure 6.5 exemplifies the visualization of a standard directed network on a dataset of bilateral financial exposures. As interbank exposures are not publicly available, the network is illustrated with country-level exposures. More specifically, the data represent banks' outstanding loans and holdings of securities, i.e. "claims", in other countries and are collected from the Bank for International Settlements (BIS) banking statistics. The figure shows a graph of financial relationships between countries where




Fig. 6.5 A network of financial linkages. *Notes* The figure illustrates an example of a network of financial linkages. Nodes of each economy are scaled as to the sum of exposures to other economies, whereas the *thickness* of the *edges* represent the size of external exposure to total exposures per economy and the color of the *edge* indicates the address of the exposure holder

objects are represented by nodes and bilateral relationships by edges. Number of objects n equals 16, giving us a matrix A of the form 16×16 , where each element A_{kl} represents the size of financial linkages between country k and l . Node size of each country is scaled based upon the sum of exposures to other countries and other countries' exposures to the base country. The thickness of each edge represents the size of exposure to total exposures of each country, where the color of the edge indicates the address of the exposure holder. Figure 6.5 illustrates, for instance, that the share of Ireland's exposure to UK is large, while UK only has a minor exposure to Ireland. Moving to a SOM grid from the representation in Fig. 6.5 involves only positioning the nodes as per their BMUs based upon the data x_j .

Neighborhoods on the SOM

The second approach follows that in Sarlin (2013b) to measure neighborhood effects on the SOM. First, we assign one of the classes C to be the event of interest. Then, we assign the locations of the event of interest in period t to be signals of similar events in that location (or some neighborhood) in period $t + s$, where s is the time span for transmission. One can choose to define the time span and neighborhood as suitable for a given task.

Contagion is exemplified on a SOM grid in Fig. 6.6, on which the labels of UniCredit Banca, ING Bank and ABN Amro are shown. Given the hypothetical failure of ABN Amro in 2002, it would also have been an indication of a failure of banks similar to it, such as UniCredit Banca in this case. It is worth noting that it is not conditional on location, such as the part of the grid that has experienced failures in the past, which implies that there is no dependence on historical data. This is a




Fig. 6.6 An exemplification of contagion on a SOM. *Notes* The figure shows the labels of UniCredit Banca, ING Bank and ABN Amro on the SOM. The final grid is the same crisp membership plane shown in Fig. 5.4

particularly important feature when dealing with events of changing nature, where data based upon history are from time to time of little use.

6.3 The Self-Organizing Time Map

Most often, the main concern of exploratory data analysis (EDA) is the analysis of either time-series or static cross sections. Given that data are drawn from a cube representation, a question of central importance is how to combine the tasks of cross-sectional and time-series analysis. That is, how to identify the occurrence and explore the properties of temporal structural changes in data, as well as their specific locations in the cross section. This can also be called exploratory temporal structure analysis.

For exploratory analysis on data from the data cube, it is critical to visualize, or present an abstraction across, all dimensions (i.e., multivariate, temporal and cross-sectional spaces). Using a standard two-dimensional SOM for exploratory temporal structure analysis, processing of the time dimension has thus far been proposed along two suboptimal directions: computing separate maps per time unit [e.g., Back et al. (1998), Denny and Squire (2005) and Denny et al. (2010)] or one map on pooled panel data [e.g., Back et al. (2001), Sarlin and Marghescu (2011) and Sarlin and Peltonen (2013)]. Owing to a possibly high number of time units and temporal differences in correlations and distributions, comparing separate maps per time unit is a laborious task while their structure may not in the least even be comparable. Denny and Squire (2005) and Denny et al. (2010) enhance temporal interpretability by applying specific initializations and visualizations. Nevertheless, the method has the drawback of an unstable orientation over time and complex comparisons of two-dimensional grids. SOMs trained with pooled data, for which time can be inferred as a type of latent

dimension that is definable but unordered, fail in describing the structure in each cross section.

While Sect. 6.1 presented several improvements to the SOM paradigm for temporal processing, the problem of visualizing changes in data structures over time has not been entirely addressed. The existing SOM literature can thus be said to have shortcomings in disentangling the temporal dimensions and cross-sectional structures for exploratory temporal structure analysis, which is the main focus of the SOTM. Thus, the SOTM can directly be related to the approach of evolutionary clustering (Chakrabarti et al. 2006), which concerns processing temporal data by producing a sequence of clustering solutions. An effective evolutionary clustering aims to achieve a balance between clustering results being faithful to current data and comparable with the previous clustering result. In this vein, Chakrabarti et al. (2006) illustrate that the usefulness of such an approach is fourfold: (i) consistency (i.e., familiarity with previous clustering), (ii) noise removal (i.e., a historically consistent clustering increases robustness), (iii) smoothing (i.e., a smooth view of transitions), and (iv) cluster correspondence (i.e., relation to historical context). The SOTM is a visual approach to evolutionary clustering by providing means to a low-dimensional representation of all three dimensions of data: (i) cross-sectional, (ii) temporal, and (iii) multivariate.

The SOTM, as proposed in Sarlin (2013d), uses the clustering and projection capabilities of the standard SOM for visualization and abstraction of temporal structural changes in data. However, here t (where $t = 1, 2, \dots, T$) is a time-coordinate in data, not in training iterations as is common for the standard SOM. To observe the cross-sectional structures of the dataset for each time unit t , the SOTM performs a mapping from the input data space $\Omega(t)$, with a probability density function $p(x, t)$, onto a one-dimensional array $A(t)$ of output units $m_i(t)$ (where $i = 1, 2, \dots, M$). After performing a mapping for all t , the timeline is created by arranging $A(t)$ in an ascending order of time t . The positions on the SOTM carry a different meaning than those on the standard SOM; the horizontal direction has a parametric interpretation of time t while the vertical direction represents positions in the data space $\Omega(t)$. Hence, the topology is rectangular rather than hexagonal and topology preservation is twofold, where the horizontal direction preserves time topology and the vertical preserves data topology.

The orientation preservation and gradual adjustment to temporal changes is performed as follows. The first principal component of Principal Component Analysis (PCA) is used for initializing $A(t_1)$ and setting the orientation of the SOTM. PCA on $\Omega(t_1)$ provides the eigenvector of the first principal component, which is used for initializing $A(t_1)$. For preserving the orientation between consecutive patterns in a time series, the model uses short-term memory to retain information about past patterns. Thus, the orientation of the map is preserved by initializing $A(t_2, 3, \dots, T)$ with the reference vectors of $A(t - 1)$. Adjustment to temporal changes is achieved by performing a batch update per time t . For $A(t_1, 2, \dots, T)$, each data point $x_j(t) \in \Omega(t)$ (where $j = 1, 2, \dots, N(t)$) is compared to reference vectors $m_i(t) \in A(t)$ and assigned to its BMU $m_b(t)$:

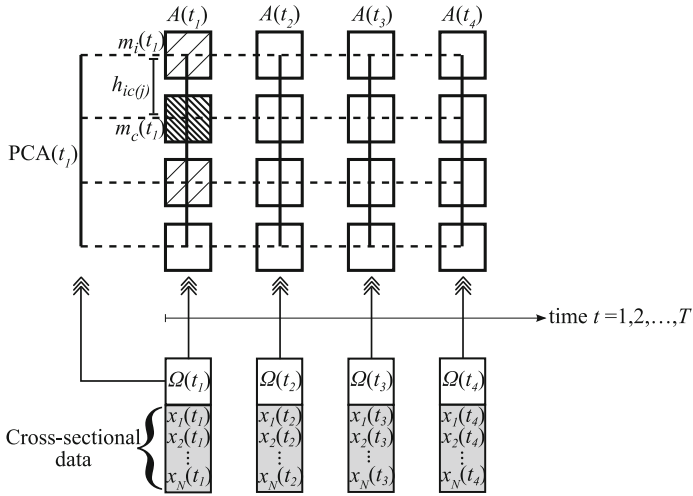


Fig. 6.7 The functioning principles of the SOTM. *Notes* This figure shows the functioning of the SOTM. The lower part of the figure represents the temporal data (where time increases from left to right) and the upper part represents the SOTM grid and its training

$$\|x_j(t) - m_b(t)\| = \min_i \|x_j(t) - m_i(t)\|. \tag{6.8}$$

Then each reference vector $m_i(t)$ is adjusted using the batch update formula:

$$m_i(t) = \frac{\sum_{j=1}^{N(t)} h_{ib(j)}(t)x_j(t)}{\sum_{j=1}^{N(t)} h_{ib(j)}(t)}, \tag{6.9}$$

where index j indicates the input data that belong to unit b and the neighborhood function $h_{ib(j)}(t) \in (0, 1]$ is defined as a Gaussian function

$$h_{ib(j)}(t) = \exp\left(-\frac{\|r_b(t) - r_i(t)\|^2}{2\sigma^2}\right), \tag{6.10}$$

where $\|r_b(t) - r_i(t)\|^2$ is the squared Euclidean distance between the coordinates of the reference vectors $m_b(t)$ and $m_i(t)$ on the one-dimensional array, and σ is the user-specified neighborhood parameter. From this follows obviously that neighborhood σ only includes vertical relationships. In contrast to what is common for the standard batch SOM, the neighborhood σ is constant over time for a comparable timeline, not a decreasing function of time as is common when time represents iterations.

To sum up, Fig. 6.7 presents the functioning principles of the SOTM. Yet, even though the figure illustrates the notion of a neighborhood function with a crisp $h_{ib(j)}$ above the BMU, it is worth noting that the function decreases gradually below the

BMU. Further, the algorithmic principles of the SOTM can be distinguished as follows:

```

 $t = 1$ 
initialize  $A(t)$  using PCA on  $\Omega(t)$ 
apply the batch update to  $A(t)$  using  $\Omega(t)$ 
while  $t < T$ 
     $t = t + 1$ 
    initialize  $A(t)$  using the reference vectors of  $A(t - 1)$ 
    apply the batch update to  $A(t)$  using  $\Omega(t)$ 
end
order  $A(t)$  in an ascending order of time  $t$ 

```

6.3.1 SOTM Properties

The above presented SOTM specification, while being flexible in nature, disposes some assumptions on distance metrics and grid shapes, as well as other computational details. Even though a SOTM mapping to one-dimensional arrays loses in granularity and detail to the two-dimensional one, the sole case of successful complete mathematical study of the SOM is in one dimension (though with one input dimension as well) [for a review see Cottrell et al. (1998)]. Further, a two-dimensional representation of the SOTM, while describing less detail, facilitates interpretation over the three-dimensional case. The SOTM is implemented using the Euclidean metric for the sake of simplicity and purpose herein as well as sticks to the standard batch SOM with exponential neighborhood functions. The batch SOM is preferred over the sequential SOM for its well-known properties of efficiency and precision [see, e.g., Kohonen (2001)]. Further, the disadvantage of all data points having to be available in batches is not a concern given that the entire cross section is accessible simultaneously at each time t . In this sense, the SOTM can be seen as a type of online batch SOM.

When compared in terms of computational cost, the SOTM is cheaper than a standard SOM of the same size since matching and learning is restricted by time t . Thus, the SOTM also has the asset of keeping the most important properties and the interpretation of the SOM as it has its basis in the very standard SOM algorithm. While the purpose of use of the SOTM is different, the functioning of it can also be linked to several other pieces of literature extending the SOM. For instance, the increase in number of units over time resembles the functioning of Growing SOMs (Fritzke 1994) and the short-term memory initialization resembles SOMs with feedback connections [e.g., Voegtlin (2002)].

While the SOTM herein uses specifications from the very standard SOM literature, such as batch training, Euclidean metric and exponential neighborhoods, its matching, learning and neighborhoods could be implemented in various modified fashions, such as those discussed in the related literature (Sect. 6.1). Parametrization of batch training can also be performed in a number of ways depending on the task and data at hand. For instance, the first array $A(t_1)$ may be trained until convergence if the

initialization is far from converged and the number of training iterations of each array $A(t)$ may be increased if quantization accuracy is relatively important. Idle units, i.e., units not attracting any data, while representing a discrepancy between array $A(t)$ and data $\Omega(t)$, may also be dealt with through increases in training iterations.

Finally, by simply interchanging the time dimension of a SOTM to a time-to-event dimension, Sarlin (2013e) shows that the SOTM provides means for illustrating patterns in time-to-event data. Time-to-event data are, in their most frequent definition, nothing more than the time that elapses until some specified event occurs. Yet, in addition to the time before an event, time-to-event data may have an afterlife as well as a life during the event. Hence, the time-to-event SOTM focuses on understanding dynamics in multivariate data before, during and after events, such as in the case of assessing the path to and afterlife of a failure of a financial institution or country and diagnosis of a disease in patient data. Yet, the x -axis of the SOTM need not be limited to the definitions of time and time to an event. It could, in fact, represent any variable, such as age in customer segmentation and states in process monitoring.

6.3.2 Qualities and Properties of the SOTM

Common quality measures for evaluating the goodness of a SOM are quantization error, distortion measure and topographic error. These, as well as other measures of the SOM, could be adapted to apply for quantifying the qualities and properties of SOTMs, where quality refers to the goodness of the mapping and property to characteristics of the data. Computations of quality and property measures can be distinguished as follows: quality measures of SOTMs are summed over T whereas property measures of a SOTMs depict the characteristics of data at each t . However, property measures obviously also illustrate time-specific qualities of SOTMs.

The fit of the SOTM to the data distribution can be measured with an adaptation of the standard quantization error and distortion measure. The time-restricted quantization errors ε_{qe} and $\varepsilon_{qe}(t)$ compute the average distance between $x_j(t) \in \Omega(t)$ and $m_b(t) \in A(t)$:

$$\varepsilon_{qe} = \frac{1}{T} \sum_{t=1}^T \frac{1}{N(t)} \sum_{j=1}^{N(t)} \|x_j(t) - m_{b(j)}(t)\|, \quad (6.11)$$

$$\varepsilon_{qe}(t) = \frac{1}{N(t)} \sum_{j=1}^{N(t)} \|x_j(t) - m_{b(j)}(t)\|. \quad (6.12)$$

The distortion measures ε_{dm} and $\varepsilon_{dm}(t)$ indicate similarly the fit of the map to the shape of the data distribution, but also account for the radius of the neighborhood:

$$\varepsilon_{dm} = \frac{1}{T} \sum_{t=1}^T \frac{1}{N(t)} \frac{1}{M(t)} \sum_{j=1}^{N(t)} \sum_{i=1}^{M(t)} h_{ib(j)}(t) \|x_j(t) - m_{b(j)}(t)\|, \quad (6.13)$$

$$\varepsilon_{dm}(t) = \frac{1}{N(t)} \frac{1}{M(t)} \sum_{j=1}^{N(t)} \sum_{i=1}^{M(t)} h_{ib(j)}(t) \|x_j(t) - m_{b(j)}(t)\|. \quad (6.14)$$

The topology preservation of the SOTM can also be measured using an adaptation of the standard topographic error. The time-restricted topographic errors ε_{te} and $\varepsilon_{te}(t)$ measure by $u(x_j(t))$ the average proportion of $x_j(t) \in \Omega(t)$ for which first and second BMUs (within $A(t)$) are non-adjacent units:

$$\varepsilon_{te} = \frac{1}{T} \sum_{t=1}^T \frac{1}{N(t)} \sum_{j=1}^{N(t)} u(x_j(t)), \quad (6.15)$$

$$\varepsilon_{te}(t) = \frac{1}{N(t)} \sum_{j=1}^{N(t)} u(x_j(t)). \quad (6.16)$$

While quantifying the degree of temporal changes in data is of central importance, it is oftentimes a difficult task. The SOTM enables approximating the structural change between time units $t - 1$ and t by an average Euclidean distance between $m_i(t - 1) \in A(t)$ and $m_i(t) \in A(t)$ for all pairs $i = 1, 2, \dots, M$. The distance is meaningful given that the ending point of $A(t - 1)$ is the starting point of $A(t)$ in training and given that the adjustment to temporal changes (i.e., σ) is constant over time. The structural changes ε_{sc} and $\varepsilon_{sc}(t)$ are computed as follows:

$$\varepsilon_{sc} = \frac{1}{T} \sum_{t=1}^T \frac{1}{M(t)} \sum_{i=1}^{M(t)} \|m_i(t - 1) - m_i(t)\|, \quad (6.17)$$

$$\varepsilon_{sc}(t) = \frac{1}{M(t)} \sum_{i=1}^{M(t)} \|m_i(t - 1) - m_i(t)\|. \quad (6.18)$$

When the quantization error $\varepsilon_{qe}(t)$, distortion measure $\varepsilon_{dm}(t)$ and topographic error $\varepsilon_{te}(t)$ are computed for $t = 1, 2, \dots, T$ and structural change $\varepsilon_{sc}(t)$ for $t = 2, 3, \dots, T$, they can be plotted over time. This is useful for identifying properties and qualities of data at each time unit, in particular the degree of temporal changes in data. Similarly, the ε_{qe} , ε_{dm} , ε_{te} and ε_{sc} can also be plotted over different free parameters, such as grid size and neighborhood radius.

6.3.3 Visualizations of the SOTM

The output of the SOTM is a two-dimensional array of units, with time on the horizontal direction and data structures on the vertical, which represents a multidimensional space. While there exist numerous visualizations for the SOM that could be applied to the SOTM framework, this subsection focuses on the standard ones that enhance the objectives of the SOTM. For each individual input, a feature plane represents the spread of its values. Thus, one can interpret vertical differences as cross-sectional properties and horizontal differences as temporal changes. As for the standard SOM, the feature planes are different views of the same map, where one unique point represents the same unit on all planes. The coloring of the feature planes is again performed using the ColorBrewer's (Harrower and Brewer 2003) scale, in which variation of a blue hue occurs in luminance and light to dark represent low to high values. As the scale is common for the entire SOTM (i.e., $A(t)$ for $t = 1, 2, \dots, T$) for each feature plane, the changes in the spread of values are shown by variations in shade.

While plots of $\varepsilon_{qe}(t)$, $\varepsilon_{dm}(t)$, $\varepsilon_{te}(t)$ and $\varepsilon_{sc}(t)$ show the changes in the measures over time, the assessment of structural differences on the horizontal and vertical dimensions of the SOTM can be enhanced by a Multidimensional Scaling (MDS) method, such as Sammon's non-linear mapping (Sammon 1969). The reason for preferring Sammon's mapping over other MDS methods is its focus on local distances. Time is disentangled by mapping all multidimensional SOTM units $m_i(t)$ (where $t = 1, 2, \dots, T$) to one dimension using Sammon's mapping and then plotting that dimension individually for each time t . Thus, this representation has Sammon's dimension on the y axis and time on the x axis. The detection of structural changes and topographic errors is facilitated by connecting adjacent units with solid (data topology) and dashed (time topology) lines for a net-like representation and showing topographic errors $u(x_j(t))$ through color coding. Moreover, a coloring method based upon that in Kaski et al. (2001) for revealing changes in cluster structures can be applied to the SOTM. The well-known uniform color space CIELab (1986) is used, where perceptual differences of colors represent distances in the data space, as approximated by the Sammon's mapping. However, as the SOMs of the SOTMs are one-dimensional, only one dimension (blue to yellow) of the color space is used.

As proposed in Sarlin and Yao (2013), the visualization of cluster structures on the SOTM may still be enhanced by pairing it with classical cluster analysis. This provides objective means for identification of *changing*, *emerging* and *lost* clusters over time. Hence, the three types of dynamics in cluster structures can be defined as follows: (i) a cluster is *lost* when one or more units are a member of it in time t and none is in $t + 1$, (ii) a cluster *emerges* when no unit is a member of it in time t and one or more are in $t + 1$, and (iii) a cluster *changes* when the (positive) number of units being a member of it in time t and $t + 1$ differ. Again, cluster memberships are visualized through a qualitative color scheme from ColorBrewer (Harrower and Brewer 2003), where groups are differentiated in hue contrast with nearly constant saturation and lightness.

6.3.4 Some Illustrative Examples

To illustrate the functioning, output and quality and property measures of the SOTM, toy data are generated and drawn from the three-dimensional data cube. This subsection motivates the choice of a SOTM over a naïve SOM model and validates the output of a SOTM by representing expected patterns, as well as provides a guide for interpreting patterns on a SOTM. The illustrative examples also cover second-level clustering of the SOTM and a time-to-event SOTM.

Toy Data

The data need to come from a three-dimensional cube, where one dimension represents time, one the cross-sectional entities and one the input variables, such as the data cube in Fig. 3.1. The toy data are generated by setting five weights w_{1-5} that adjust a mixture of randomized shocks on four different levels: group-specific (g), time-specific (t), variable-specific (r) and common (j) properties. For each variable, group-level differences are included to have artificial clusters, time-level properties to introduce temporal trends, and group-specific and common shocks to introduce general noise. Data $x(r, g, j, t)$ are generated by combining group-specific trends E with common shocks across data and over time,

$$x(r, g, j, t) = E(r, g, t) + w_4(r, g) e_4(r, t) + w_5(r, g) e_5(r, j, t), \quad (6.19)$$

and

$$E(r, g, t) = w_1(r) e_1(g) + w_2(r) e_2(g) t + w_3(r) e_3(g, t), \quad (6.20)$$

where $e_{1,3-5} \sim N(0, 1)$, $e_2 \sim U(0, 1)$, r stands for variables, g for groups, t for time and j for entities, and E computes group-specific trends. The rationale for drawing e_2 from a uniform rather than a normal distribution is to have larger variation in the group-specific slopes. Finally, each variable x_j is transformed into $[0,1]$ through a logistic sigmoidal function.

Weights specify the following properties of data: w_1 sets the group-specific intercepts, w_2 the group-specific slopes over time, w_3 the magnitude of group-specific random shocks, w_4 the magnitude of time-specific common shocks, and w_5 the magnitude of common shocks. Figure 6.8 plots four variables and reports the used weights for generating 100 entities over 10 periods, where the color coding illustrates five groups of entities. Particular characteristics of the below four variables are as follows: x_1 has small differences in intercepts and a positive slope; x_2 has large differences in intercepts, a negative slope and minor group-level and common shocks over time and across entities; x_3 has large differences in intercepts, and a constant trend with minor common shocks across entities and over time; and x_4 has large differences in intercepts and large common shocks over time.




Fig. 6.8 Four toy variables for the standard SOTM. *Notes* Data consist of 5 groups of 20 cross-sectional entities over 10 periods, where the color coding illustrates the groups, the x axis represents time and y axis the values

A Naïve SOM Model

Comparing the applicability of methods for EDA in general and exploratory temporal structure analysis in particular is not an entirely straightforward task. The absence of a quantitative evaluation, such as common prediction or classification comparisons, is due to the lack of a comparable evaluation function. Instead, the focus herein is on illustrating the advantages of the SOTM by comparing it to a naïve one-dimensional SOM model on the entire dataset Ω . Although a fair comparison would make use of a two-dimensional SOM, the exercise is still feasible for illustrating how time, when being embedded, cannot be fully represented on a standard SOM, not even when utilizing post-processing techniques. In this SOM, the pooled toy dataset is used as an input to a SOM with 5 units as per the number of groups in data. Figure 6.9 shows the SOM, its feature planes and a post-processed trajectory for the toy dataset. Figure 6.9a shows the SOM where differences in units are represented by perceptual differences in colors. Its feature planes in Fig. 6.9b depict characteristics of the data in Fig. 6.8, but obviously disregard the time dimension. For instance, neither time trends of x_1 and x_2 nor time shocks of x_4 are depicted. Variable x_3 is, however, correctly depicted as it is close to constant over time. A trajectory of an entity can be used for describing its evolution on the SOM over time. In Fig. 6.9c, a trajectory of an arbitrary data point over the 10 periods exemplifies that, while temporal movements of individual data exist, changes in cluster (or unit) structures are not represented. In particular, this illustrates that the evolution of data structures in Fig. 6.8 is not represented by a static SOM.

Illustrating a Standard SOTM

A natural next step is to apply the SOTM on the toy data. Although we make use of the above presented standard SOTM specification, the free parameters still have to be specified. The SOTM is chosen to have 5×10 units, where 5 units represent data topology at time t on the vertical direction and 10 units the time topology on the horizontal direction. The number of units on the horizontal axis is set by the




Fig. 6.9 A naïve one-dimensional SOM. *Notes* The figure shows **a** a naïve one-dimensional SOM, **b** feature planes of the SOM, and **c** an exemplification of temporal movements of an arbitrary data point on the SOM

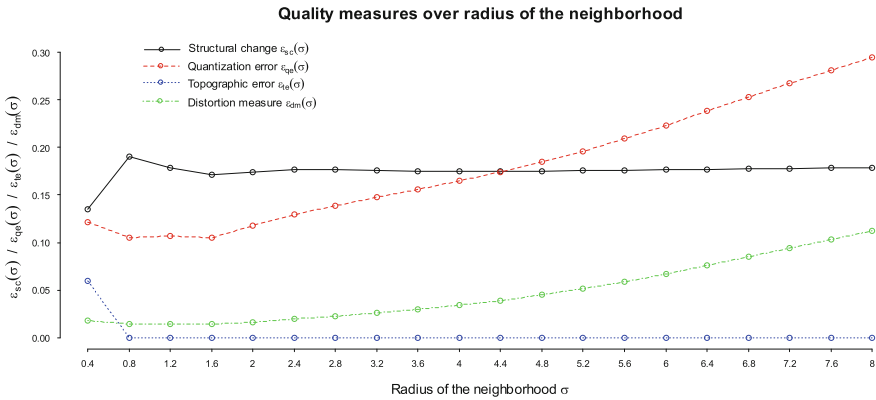


Fig. 6.10 An example of quality measures of the SOTM. *Notes* For models with a 5×10 array of units, the errors (ε_{qe} , ε_{dm} and ε_{te}) are computed as aggregates of all time units $t = 1, 2, \dots, T$ and structural changes ε_{sc} of time units $t = 2, 3, \dots, T$ over neighborhood radii $\sigma = \{0.4, 0.8, \dots, 8\}$

number of time units T in data, while the number of units (or clusters) on the vertical axis equals the number of groups in the data. The quality measures presented in Sect. 6.3 are used for evaluating performance over different parameters. For all time units t , the distortion measure ε_{dm} and quantization error ε_{qe} measure the fit to data Ω , while topographic error ε_{te} measures the aggregated topology preservation. The structural change ε_{sc} , on the other hand, shows the distance between horizontal units. Figure 6.10 shows the quality measures over radius of the neighborhood σ ranging from 0.4 to 8. The figure illustrates aspects of not only these data in particular, but also SOTM training in general. It shows the strength of the topology preservation in the SOTM; a topology error ε_{te} is only found for experiments with $\sigma = 0.4$. Though the magnitude of quantization error ε_{qe} and distortion measure ε_{dm} differs due to simple and squared distances, an obvious effect is the increase of the measures when σ increases. The structural change starts to decrease when $\sigma = 0.6$, and decreases until it stabilizes for $\sigma \geq 1.6$. When aiming at data abstraction and exploratory analysis, choosing optimal parameter values for a SOTM, likewise for a SOM, is a difficult task;

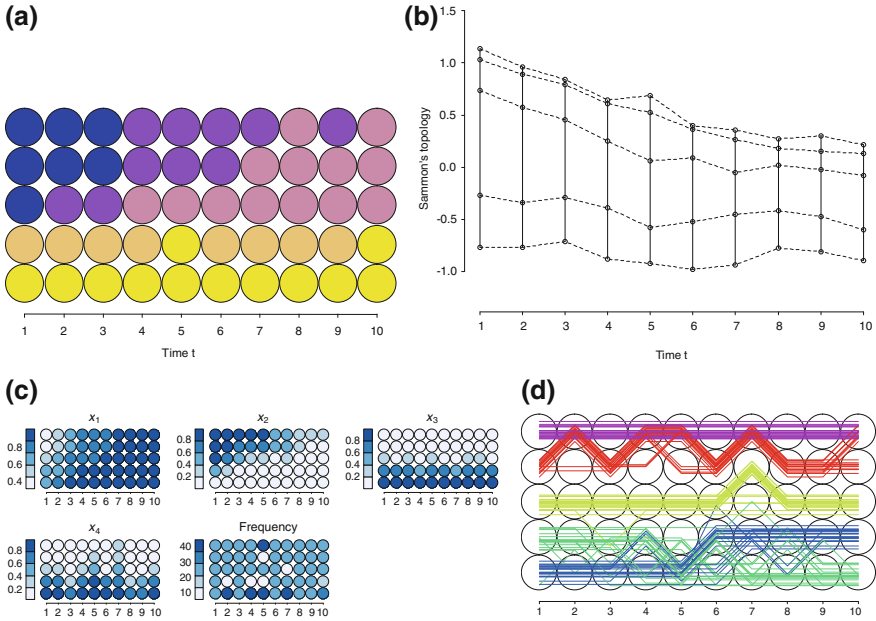


Fig. 6.11 An illustrative example of the SOTM. *Notes* The figure shows **a** a SOTM grid with perceptual differences in color representing distances between units, **b** a plot of the SOTM units according to Sammon’s topology on the vertical axis and time on the horizontal axis where neighboring units are connected with lines, **c** feature planes and a frequency plot on the SOTM grid, and **d** the data overlaid as time series, or trajectories, on top of the SOTM grid with coloring that corresponds to that in Fig. 6.8

the choice can be said to depend on the relative preferences of the analyst between topographic and quantization errors. However, as the interpretation of a SOTM relies heavily on topology preservation, not the least the time dimension, topographic errors ought to be of higher importance. As we here only have topographic errors for $\sigma = 0.4$, we can choose a SOTM with minimum quantization error and distortion. The chosen SOTM has thus a radius of the neighborhood $\sigma = 1.6$.

The final SOTM is found in Fig. 6.11a and a Sammon’s mapping of it in Fig. 6.11b. The coloring of the SOTM uses the CIELab unified color space, where perceptual differences in colors represent differences between units as approximated by Sammon’s mapping. Feature planes in Fig. 6.11c represent layers of the SOTM, while Fig. 6.11d reports trajectories of all data on the SOTM. Figure 6.12 illustrates a plot of property measures $\varepsilon_{qe}(t)$, $\varepsilon_{dm}(t)$, $\varepsilon_{te}(t)$ and $\varepsilon_{sc}(t)$ over time.

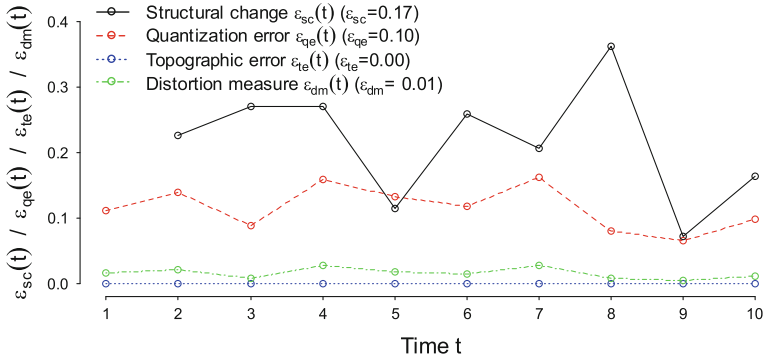


Fig. 6.12 An example of property measures of the SOTM. *Notes* The errors ($\varepsilon_{qe}(t)$, $\varepsilon_{dm}(t)$ and $\varepsilon_{te}(t)$) are computed for time units $t = 1, 2, \dots, T$ and $\varepsilon_{sc}(t)$ for time units $t = 2, 3, \dots, T$ for the final SOTM with an 5×10 array of units and $\sigma = 1.6$

A Guide for Interpreting the SOTM

This part gives a brief guide for interpreting the SOTM and its visualizations. A key to interpreting the SOTM is to understand the grid structure and the following representation of data along two directions. The vertical direction (or columns of units) has a similar interpretation as a standard SOM (cf. Fig. 6.9), but each one refers to a specific time unit. Thus, it represents the cross-sectional data structure, or data topology, at time t , where similar units are located close together. The horizontal direction (or rows of units), while being conceptually different from a standard SOM, has a similar interpretation. It represents the time structure, or time topology, where similar units are again located close together, but refers instead to resembling units at different points in time. Hence, differences along both directions represent differences between respective topologies when interpreting properties of high-dimensional structures, values of individual inputs or any other linked information. Below, we use the above toy example for discussing the interpretation of the SOTM visualizations.

Figures 6.11a and b give information on the distance structure of the SOTM. Perceptual differences in colors (blue to yellow) in Fig. 6.11a represent differences between units as per distances in the Sammon's mapping in Fig. 6.11b. In Fig. 6.11b, differences between units on both vertical and horizontal directions should, however, be interpreted by values of Sammon's topology (color in Fig. 6.11a and y axis in Fig. 6.11b). The differences in values of units on the vertical direction represent distances in cross-sectional data structures at a specific time t and differences in values of units on the horizontal direction represent distances over time. In the Sammon's mapping, solid connections between units represent data topology and dashed connections time topology. The figures show that the data are clustered into two distinct groups: the three uppermost horizontal rows (yellow, green and blue, cf. Fig. 6.8) and two lowest rows (red and purple, cf. Fig. 6.8). The structure of the SOTM illustrates two types of temporal changes: common trends of the entire structures and

movements of individual units. The former type preserves distances between units at each point in time, but moves the entire structure to some direction, while the latter type illustrates changes in distances to neighboring units. Figures 6.11a and b show that the two distinct groups converge over time, in particular that the uppermost groups of data move towards the rest of the data, as the raw data in Fig. 6.8 confirm. Convergence is mostly a result of inputs x_1 and x_2 moving towards maximum and minimum values over time, in particular the large changes of x_2 .

Figure 6.11c illustrates the spread of values for each of the four inputs and should similarly be interpreted along the two directions. One type of validation of the SOTM is that the four feature planes correspond to the description of differences in group-level intercepts and slopes, as well as time-specific shocks, for the inputs (cf. Fig. 6.8). That is, x_1 has small differences in intercepts and a positive slope, x_2 has large differences in intercepts and a negative slope, x_3 has large differences in intercepts and a constant trend, and x_4 has large differences in intercepts and large common shocks over time. The frequency plane in Fig. 6.11c represents density of data on the SOTM grid and is particularly useful for two purposes. Since the SOTM attempts to update cluster structures in $A(t-1)$ to $A(t)$ by a batch update, while structures in data $\Omega(t-1)$ and $\Omega(t)$ may be of different nature, one purpose of use is locating idle units. While idle units represent a change in cluster structures, the reference vectors are still transmitted to $A(t)$ through the short-term memory.¹ The frequency plots also enable observing evolution of densities over time. While changes in the spread of values indeed indicate changes in data, frequencies are an equally important property of structures. In this toy example, the main interpretation is the absence of idle units. Another validation of the SOTM is the plot of all individual data on the SOTM in Fig. 6.11d. The coloring of the trajectories corresponds to that in Fig. 6.8 and illustrates the evolution of the groups on the SOTM. While the groups are separated during most of the periods, some overlap and interchange of positions occurs over time. The one-period overlaps of red and purple groups accurately correspond to the time-specific shocks of x_4 . The occurrence of position interchanges of blue and green groups at periods 3–5 are likely due to change in input x_1 and finally in period 7 due to substantial changes in input x_2 .

Plots of property measures over time in Fig. 6.12 illustrate the variation of $\varepsilon_{qe}(t)$, $\varepsilon_{dm}(t)$, $\varepsilon_{te}(t)$ and $\varepsilon_{sc}(t)$ over time. When assessing properties for each time unit t , the structural change $\varepsilon_{sc}(t)$ measures divergence of $m_i(t)$ from the units $m_i(t-1)$, whereas the rest mainly visualize quantization and topographic qualities across a SOTM. While increases in quantization error $\varepsilon_{qe}(t)$ and distortion $\varepsilon_{dm}(t)$ represent the fit of data $\Omega(t)$ to units $m_i(t)$, increases in topographic error $\varepsilon_{te}(t)$ represents the topology preservation for each array $A(t)$. For the toy data, the large variation in $\varepsilon_{sc}(t)$ depicts the existence of large differences between data structures. In particular, we can see that highest values of $\varepsilon_{sc}(t)$ in periods 3–4 and 6–8 co-occur with large common temporal shocks in x_4 . Small or none variation of $\varepsilon_{qe}(t)$, $\varepsilon_{dm}(t)$ and $\varepsilon_{te}(t)$ confirms that quantization and topographic errors are low over time, while

¹ I suggest to illustrate with idle units through some color coding. While idle units have implemented to be colored in gray, these specific cases are not encountered in the experiments performed here.




Fig. 6.13 Four toy variables for the clustering of the SOTM. *Notes* Data consist of 5 groups of 20 cross-sectional entities over 10 periods, where the color coding illustrates the groups, the x axis represents time and the y axis the values. The figure reports the used weights w_{1-5} for generating the x_{1-4}

the difference in the magnitude of the quantization accuracies is a result of them being measured with simple and squared distances, respectively.

Illustrating the Clustering of the SOTM

This part introduces the second-level clustering of SOTMs with experiments on toy data. The experiments on data with expected patterns illustrate the usefulness of combining clustering techniques with the SOTM. For the experiments, data are created using the process for generating toy data in Eqs. (6.19) and (6.20). Figure 6.13 reports the used weights for generating the four variables with five groups of 20 cross-sectional entities over 10 periods, where the color coding illustrates the groups. Particular characteristics of the below four variables are obviously that the time series of x_{1-2} are quasi-stationary and those of x_{3-4} are non-stationary. While we may not fulfill all conditions of stationarity, particularly not constant variance and autocorrelation, the aim of these data is to have two variables with a somewhat constant data structure and two with a time-varying and converging structure. That is, the time-varying data are generated such that parts of the groups converge whereas others diverge. Evident changes of x_{3-4} in Fig. 6.13 are, for instance, that the red group diverges from the purple group and the blue converges to the green group.

This part presents two toy examples of the SOTM: one with quasi-stationary and one with non-stationary data. The quasi-stationary data x_{1-2} and the non-stationary data x_{3-4} are separately used as inputs for the SOTM. The SOTM is specified to have 5×10 units, where 5 vertical units represent data topology and 10 horizontal units time topology. Again, the number of units (or clusters) on the vertical axis is set to equal the number of generated groups in data, while the number of units on the horizontal axis is fixed by the number of time units T . When choosing the final specification of the SOTM, the above presented quality measures are used, but not shown, for the sake of brevity. As the main focus ought to be on topographic accuracy, the neighborhood is chosen as to minimize quantization error given no topographic errors.




Fig. 6.14 A SOTM applied to data with and without temporal variation. *Notes* For both rows, the two first grids are feature planes and represent the spread of values of x_{1-2} and x_{3-4} . The final grid shows cluster memberships using color coding. Vertical color scales on the left of feature planes link a cluster number to each color

To the trained SOTM, a second-level clustering is applied using Ward’s method. As the number of generated groups is predefined, the performance of different K need not be compared using clustering validation measures. On the quasi-stationary variables x_{1-2} , K is set to equal to the number of created groups. In the non-stationary case, the same K is used to better illustrate the difference to the quasi-stationary case. However, exploring different K would be useful for illustrating properties of the data. The first row of Fig. 6.14 shows feature planes and cluster memberships for a SOTM on the quasi-stationary data x_{1-2} . It is worth noting that the coloring of the clusters follows the coloring of the groups in Fig. 6.13. Indeed, one can observe that both the two feature planes and the cluster memberships are constant over time. This illustrates that quasi-stationary data may be labeled by the rows of the SOTM. Likewise, the second row of Fig. 6.14 shows feature planes for a SOTM on the non-stationary data x_{3-4} . The feature planes clearly depict the increasing and decreasing trends in data. This is also reflected in the cluster memberships. The feature of approximating the probability density functions of data $p(x, t)$ lead to a direct interpretation of memberships; the denser a part of the data space $\Omega(t)$, the higher is the number of units in that location. The convergence of the green, blue and orange groups and divergence of the red from the purple group (as also shown in Fig. 6.13) are shown as increases and decreases of vertical units in a second-level cluster. That is, in Fig. 6.14, we can observe that cluster 2 (blue) *disappears* in period 7 and that cluster 1 (red) *emerges* from cluster 4 (purple) in period 3, as well as cluster 3 (green) and 4 (purple) *change* from one vertical unit to two and from two units to one, respectively. This motivates the need for a second-level clustering to aid in interpreting changes in cluster structures.

Illustrating a Time-to-Event SOTM

This part aims at illustrating and validating the performance of the SOTM on time-to-event data. The SOTM for time-to-event data has a different interpretation for time t . Rather than representing the time span in data, it represents the time to a specific event. Hence, it takes, for instance, the following form: $t = -T, -T + 1, \dots, T - 1, T$, where T sets the range of time units before and after the event. For illustrative purposes, the experiments use an equal number of periods before and after the events, but the SOTM obviously sets no such restriction. The process is a modified version of that in Eqs. (6.19) and (6.20) and is steered with similar parameters for setting the properties of data. Whereas data need to come from a three-dimensional space, where one dimension represents time, one the cross-sectional entities and one the input variables, the time dimension needs to be transformed to represent the distance of each data point to some event v , in order to represent time-to-event data. The process starts by first drawing random events for a specified number of entities over a number of periods. The number of periods is set to equal the number of time-to-event states: $2T + 1$. Hence, for each entity j , the events v are drawn from a discrete uniform distribution as follows: $v \sim U(1, 2T + 1)$. Then, the data are generated by setting the five weights w_{1-5} , where only w_2 differs by representing time-to-event specific properties.

After generating the multivariate time-series, as well as time-stamped events, they are turned into time-to-event data. The used data consists of 5 groups with 20 entities each over 19 periods (where $T = 9$). The time stamps for the events are drawn from the above presented discrete uniform distribution, and are hence equally likely for all 19 periods. Figure 6.15a plots the four generated variables and reports the used weights w_{1-5} . The figure illustrates the group memberships through ColorBrewer's qualitative color scheme. While only showing minor randomized shocks over time, x_4 better illustrates the five distinguished groups in data. Although the data are generated with group-specific effects, one key message of Fig. 6.15a is that the groups cannot clearly be distinguished. This is due to the patterns being related to the events and the events being randomly distributed. Thus, the data are transformed to time-to-event data by ordering them according to time to the events. Figure 6.15b shows a plot of the same above used data, but in time-to-event format, where the color coding again illustrates 5 groups of 20 entities each. This figure better illustrates the generated patterns in data. It shows increases towards the events for x_1 and x_3 , decreases for x_2 and weak decreasing patterns related to the events for x_4 .

For illustrative purposes, two experiments are conducted on these data: a standard SOTM and a time-to-event SOTM. In both experiments, the quality measures are used for finding an adequate SOTM, but are not reported for brevity. First, the standard SOTM is applied on data ordered and grouped as per the periods ($t = 1, 2, \dots, 19$). Figure 6.16a shows for the standard SOTM the feature planes for variables x_{1-4} and the events v , i.e., the spread of variable values on the SOTM grid. The event feature plane shows averages of data located in each unit for the binary variable indicating whether or not an event occurs, and illustrates that the events mostly occur in the upper part of the SOTM. The information illustrated by the figure complements the

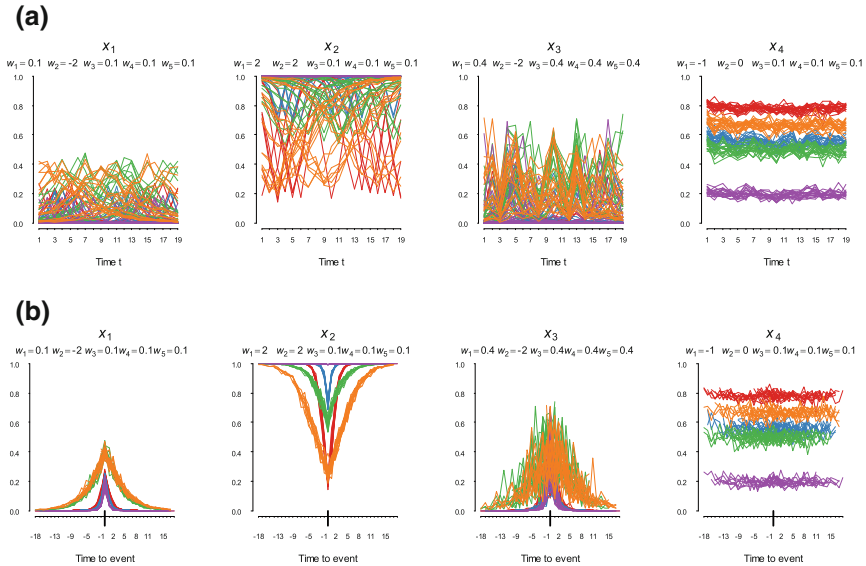


Fig. 6.15 Four toy variables and their time-to-event counterparts. *Notes* The figure represents **a** four toy variables, and **b** their time-to-event counterparts. Data consist of 5 groups of 20 cross-sectional entities over 19 periods, where the color coding illustrates the groups, the x axis represents **(a)** time or **(b)** time to an event, and the y axis the values. The figure reports the used weights w_{1-5} for generating the x_{1-4} above each plot

patterns in Fig. 6.15a (and confirms those in Fig. 6.15b): high values for x_1 and x_3 and low for x_2 are positively related to the events, whereas x_4 shows no clear patterns related to the events. Yet, this says little about the dynamics before, during and after the events, such as how early prior to events do the changes start. The time-to-event SOTM in Fig. 6.16b addresses these patterns. The figure includes the feature planes for the time-to-event SOTM, and thus illustrates the dynamics before, during and after the events. The focus is on $T = 9$ (i.e., $t - 9$ to $t + 9$) to keep the time dimension comprehensible. Again, the patterns in Fig. 6.16b follow those in Fig. 6.15b. The events are obviously all in the column of units at $t = 0$. The values for x_1 and x_3 increase towards the events, values for x_2 decrease towards the events and values for x_4 do not vary over time-to-event dates. Likewise, values for x_1 and x_3 decrease and values for x_2 increase after the events. While the patterns are symmetrical and quite well-behaving, one can observe for x_1 and x_3 that towards the events positive slopes increase and away from the events negative slopes decrease, and *vice versa* for x_2 . Further, while Fig. 6.15b illustrates differences in x_1 and x_3 , Fig. 6.16b shows that the general time-to-event patterns are close to similar. A detailed look at Fig. 6.16b does, however, illustrate that the reaction of x_1 to the events is more peaked than that of x_3 . The former starts increasing at a later stage with a peak at the event higher than the rest, whereas the latter increases at an earlier stage and can be seen as having a peak with broader shoulders.

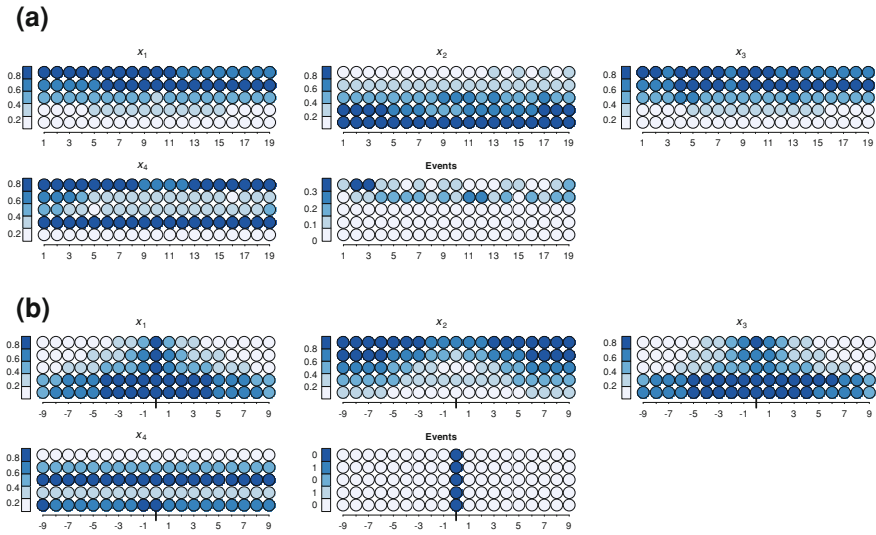


Fig. 6.16 A standard and a time-to-event SOTM on toy data. *Notes* The figure represents (a) a standard SOTM and (b) a time-to-event SOTM on the four generated toy variables. For both figures (a) and (b), the four first *grids* are feature planes and represent the spread of values of x_{1-4} , whereas the final *grid* shows the spread of events on the map

6.4 Concluding Summary

This chapter has discussed a number of extensions to the standard SOM paradigm. After a broad, yet brief, literature review of time in SOMs, the chapter introduces a number of extensions mainly focusing on improving temporal processing. Besides a focus on time, the general aim has been to enhance the SOM paradigm for processing data from the data cube in Fig. 3.1, i.e., along multivariate, temporal and cross-sectional dimensions. The suggested extensions are four. First, a fuzzification of the SOM provides means for visualizing temporal belongingness of individual data to second-level clusters and the cluster structures on the SOM. Second, transition probabilities enable visualizing probabilities of transition of individual data to second-level clusters and for assessing the overall cyclical and temporal structure on the SOM. Third, neighborhoods on the SOM and superimposed portfolio network visualizations provide means for assessing links between entities and potential for the spread of events. Finally, the fourth extension, the SOTM, enables visualizing and assessing changes in cluster structures over time. Further, the proposed second-level clustering of the SOTM enables objective identification of the temporal changes, whereas a time-to-event SOTM enables assessing patterns in multivariate data before, during and after user-specified events.

At this point, we have compared and chosen the most suitable method in Chap. 5 and suggested a range of extensions for the task at hand. The sequel of this book focuses on how these methods are applied in macroprudential oversight for the two tasks of risk identification and assessment, not to forget the third task of risk communication.

References

- Alhoniemi, E., Hollmén, J., Simula, O., & Vesanto, J. (1999). Process monitoring and modeling using the self-organizing map. *Integrated Computer Aided Engineering*, 6(1), 3–14.
- Anderson, T., & Goodman, L. (1957). Statistical inference about markov chains. *Annals of Mathematical Statistics*, 28(1), 89–110.
- Andrienko, N., Andrienko, V., Bremm, S., Schreck, T., von Landesberger, P., & Keim, D. (2010). Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns. *Computer Graphics Forum*, 29(3), 913–922.
- Back, B., Sere, K., & Vanharanta, H. (1998). Managing complexity in large data bases using self-organizing maps. *Accounting, Management and Information Technologies*, 8(4), 191–210.
- Back, B., Toivonen, J., Vanharanta, H., & Visa, A. (2001). Comparing numerical data and text information from annual reports using self-organizing maps. *International Journal of Accounting Information Systems*, 2, 249–259.
- Barreto, G. (2007). Time series prediction with the self-organizing map: A review. In P. Hitzler & B. Hammer B (Eds.), *Perspectives on neural-symbolic integration*. Heidelberg: Springer-Verlag.
- Barreto, G., & Araújo, A. (2001). Time in self-organizing maps: An overview of models. *International Journal of Computer Research*, 10(2), 139–149.
- Barreto, G., Araujo, F., & Kremer, S. (2003). A taxonomy for spatiotemporal connectionist networks revisited: The unsupervised case. *Neural Computation*, 15(6), 1255–1320.
- Bezdek, J. (1981). *Pattern recognition with fuzzy objective function algorithms*. New York: Plenum Press.
- Chakrabarti, D., Kumar, R., & Tomkins, A. (2006). Evolutionary clustering. *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 06)* (pp. 554–560). Philadelphia, PA: ACM.
- Chappell, G., & Taylor, J. (1993). The temporal Kohonen map. *Neural Networks*, 6, 441–445.
- Chen, C. (2005). Top 10 unsolved information visualization problems. *IEEE Computer Graphics and Applications*, 25(4), 12–26.
- CIE Lab. (1986). *Colorimetry*. (CIE Publication No. 15:2).
- Cottrell, M., Fort, J., & Pagès, G. (1998). Theoretical aspects of the SOM algorithm. *Neurocomputing*, 21(1–3), 119–138.
- Cottrell, M., & Letrémy, P. (2005). Missing values: Processing with the Kohonen algorithm. *Proceedings of Applied Stochastic Models and Data Analysis (ASMDA 05)* (pp. 489–496). France: Brest.
- Denny, G. W., & Christen, P. (2010). Visualizing temporal cluster changes using relative density self-organizing maps. *Knowledge and Information Systems*, 25(2), 281–302.
- Denny., & Squire, D. M. (2005). Visualization of cluster changes by comparing self-organizing maps. In *Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD 05)*. (pp. 410–419). Hanoi, Vietnam.
- Dunn, J. (1973). A fuzzy relative of the isodata process and its use in detecting compact, well-separated clusters. *Cybernetics and Systems*, 3, 32–57.
- Fritzke, B. (1994). A growing neural gas network learns topologies. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.), *Advances in Neural Information Processing Systems* (Vol. 7, pp. 625–632). MA, Cambridge: MIT Press.
- Fuertes, J., Domínguez, M., Reguera, P., Prada, M., Díaz, I., & Cuadrado, A. (2010). Visual dynamic model based on self-organizing maps for supervision and fault detection in industrial processes. *Engineering Applications of Artificial Intelligence*, 23(1), 8–17.
- Guimarães, G. (2000). Temporal knowledge discovery with self-organizing neural networks. *International Journal of Computers, Systems and Signals*, 1(1), 5–16.
- Guimarães, G., Lobo, V., & Moura-Pires, F. (2003). A taxonomy of self-organizing maps for temporal sequence processing. *Intelligent Data Analysis*, 7(4), 269–290.

- Guimarães, G., & Utsch, A. (1999). A method for temporal knowledge conversion. In *Proceedings of the International Symposium on Intelligent Data Analysis (IDA 99)* (pp. 369–382). Amsterdam, The Netherlands.
- Guo, D., Chen, J., MacEachren, A., & Liao, K. (2006). A visualization system for space-time and multivariate patterns (VIS-STAMP). *IEEE Transactions on Visualization and Computer Graphics*, 12(6), 1461–1474.
- Hagenbuchner, M., Sperduti, A., & Tsoi, A. (2003). Self-organizing map for adaptive processing of structured data. *IEEE Transactions on Neural Networks*, 14, 191–505.
- Hammer, B., Micheli, A., Neubauer, N., Sperduti, A., & Strickert, M. (2005). Self organizing maps for time series. *Proceedings of the Workshop on Self-Organizing Maps (WSOM 05)* (pp. 115–122). Paris: France.
- Hammer, B., Micheli, A., Sperduti, A., & Strickert, M. (2004). A general framework for unsupervised processing of structured data. *Neurocomputing*, 57, 3–35.
- Harrower, M., & Brewer, C. (2003). ColorBrewer.org: An online tool for selecting color schemes for maps. *The Cartographic Journal*, 40(1), 27–37.
- Horio, K., & Yamakawa, T. (2001). Feedback self-organizing map and its application to spatio-temporal pattern classification. *International Journal of Computational Intelligence and Applications*, 1(1), 1–18.
- Kangas, J. (1990). Time-delayed self-organizing maps. *Proceedings of the International Joint Conference on Neural Networks (IJCNN 90)* (pp. 331–336). San Diego, USA: IEEE Press.
- Kangas, J. (1992). Temporal knowledge in locations of activations in a self-organizing map. *Proceedings of the International Conference on Artificial Neural Networks (ICANN 92)* (pp. 117–120). Brighton, England: IEEE Press.
- Kaski, S., Venna, J., & Kohonen, T. (2001). Coloring that reveals cluster structures in multivariate data. *Australian Journal of Intelligent Information Processing Systems*, 60, 2–88.
- Kohonen, T. (1988). The “neural” phonetic typewriter. *Computer*, 21(3), 11–22.
- Kohonen, T. (1991). The hypermap architecture. In T. Kohonen., K. Mäkisara., O. Simula., & J. Kangas (Eds.), *Artificial Neural Networks* (Vol. II, pp. 1357–1360). Amsterdam, The Netherlands: Elsevier.
- Kohonen, T. (2001). *Self-organizing maps* (3rd ed.). Berlin: Springer-Verlag.
- Koskela, T., Varsta, M., Heikkonen, J., & Kaski, K. (1998). Time series prediction using RSOM with local linear models. *International Journal of Knowledge-Based Intelligent Engineering Systems*, 2(1), 60–68.
- Luyssaert, S., Sulkava, M., Raitio, H., & Hollmén, J. (2004). Evaluation of forest nutrition based on large-scale foliar surveys: Are nutrition profiles the way of the future? *Journal of Environmental Monitoring*, 6(2), 160–167.
- Martín-del Brío, B., & Serrano-Cinca, C. (1993). Self-organizing neural networks for the analysis and representation of data: Some financial cases. *Neural Computing and Applications*, 1(2), 193–206.
- Mayer, R., Abdel Aziz, T., & Rauber, A. (2007). Visualising class distribution on self-organising maps. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN 07)*. Porto, Portugal.
- Principe, J., & Wang, L. (1995). Non-linear time series modeling with self-organizing feature maps. *Proceedings of the IEEE Workshop on Neural Networks for Signal Processing* (pp. 11–20). Piscataway, NJ: IEEE Computer Society.
- Sammon, J. (1969). A non-linear mapping for data structure analysis. *IEEE Transactions on Computers*, 18(5), 401–409.
- Sarlin, P. (2013b). Exploiting the self-organizing financial stability map. *Engineering Applications of Artificial Intelligence*, 26(5–6), 1532–1539.
- Sarlin, P. (2013d). Self-organizing time map: An abstraction of temporal multivariate patterns. *Neurocomputing*, 99(1), 496–508.

- Sarlin, P. (2013e). A self-organizing time map for time-to-event data. In *Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining (CIDM'13)* (pp. 230–237). Singapore: IEEE Press.
- Sarlin, P., & Eklund, T. (2011). Fuzzy clustering of the self-organizing map: Some applications on financial time series. *Proceedings of the Workshop on Self-Organizing Maps (WSOM 11)* (pp. 40–50). Helsinki, Finland: Springer-Verlag.
- Sarlin, P., & Eklund, T. (2013). Financial performance analysis of European banks using a fuzzified self-organizing map. *International Journal of Knowledge-Based and Intelligent Engineering Systems*, 17(3), 223–234.
- Sarlin, P., & Marghescu, D. (2011). Visual predictions of currency crises using self-organizing maps. *Intelligent Systems in Accounting, Finance and Management*, 18(1), 15–38.
- Sarlin, P., & Peltonen, T. (2013). Mapping the state of financial stability. *Journal of International Financial Markets, Institutions and Money*, 26, 46–56.
- Sarlin, P., & Yao, Z. (2013). Clustering of the self-organizing time map. *Neurocomputing*, 121, 317–327.
- Sarlin, P., Yao, Z., & Eklund, T. (2012). A framework for state transitions on the self-organizing map: Some temporal financial applications. *Intelligent Systems in Accounting, Finance and Management*, 19(1), 189–203.
- Strickert, M., & Hammer, B. (2005). Merge SOM for temporal data. *Neurocomputing*, 64, 39–72.
- Strickert, M., Hammer, B., & Blohm, S. (2005). Unsupervised recursive sequences processing. *Neurocomputing*, 63, 69–98.
- Sulkava, M., & Hollmén, J. (2003). Finding profiles of forest nutrition by clustering of the self-organizing map. *Proceedings of the Workshop on Self-Organizing Maps (WSOM 03)* (pp. 243–248). Kitakyushu, Japan: Hibikino.
- Ultsch, A., Guimaraes, G., & Schmidt, W. (1996). Classification and prediction of hail using self-organizing neural networks. In *Proceedings of the International Conference on Neural Networks (ICNN 96)* (pp. 1622–1627). Washington, USA.
- Ultsch, A., & Siemon, H. (1990). Kohonen's self organizing feature maps for exploratory data analysis. In *Proceedings of the International Conference on Neural Networks (ICNN 90)* (pp. 305–308). Dordrecht, The Netherlands.
- Varsta, M., Heikkonen, J., & Millán, J. (1997). Context learning with the self organizing map. *Proceedings of the Workshop on Self-Organizing Maps (WSOM 97)* (pp. 197–202). Finland: Espoo.
- Vesanto, J., & Alhoniemi, E. (2000). Clustering of the self-organizing map. *IEEE Transactions on Neural Networks*, 11(3), 586–600.
- Voegtlin, T. (2002). Recursive self-organizing maps. *Neural Networks*, 15(8–9), 979–992.
- Wismüller, A., Verleysen, M., Aupetit, M., & Lee, J. (2010). Recent advances in nonlinear dimensionality reduction, manifold and topological learning. *Proceedings of the European Symposium on Artificial Neural Networks (ESANN 10)* (pp. 71–80). Belgium: Bruges.
- Wong, P., Shen, H., Johnson, C., Chen, C., & Ross, R. (2012). The top 10 challenges in extreme-scale visual analytics. *IEEE Computer Graphics and Applications*, 32(4), 63–77.

Chapter 7

Self-Organizing Financial Stability Map

I would very much welcome inspiration from other disciplines: physics, engineering, psychology, biology. Bringing experts from these fields together with economists and central bankers is potentially very creative and valuable. Scientists have developed sophisticated tools for analysing complex dynamic systems in a rigorous way. These models have proved helpful in understanding many important but complex phenomena: epidemics, weather patterns, crowd psychology, magnetic fields. [...] I am hopeful that central banks can also benefit from these insights in developing tools to analyse financial markets and monetary policy transmission.

– Jean-Claude Trichet, President of the ECB, Frankfurt am Main,
18 November 2010

This chapter ties together most of the previous parts of this book. Macroprudential oversight and data alike not only motivate, but also provide guidelines for building tools with visual capabilities. Data and dimension reductions, as well as their combinations, provide means for creating visual displays for a wide range of tasks, whereas a qualitative comparison shows that the Self-Organizing Map (SOM) is suitable for the task we have at hand. This chapter unifies the above discussed topics by creating a SOM-based financial stability map, coined the Self-Organizing Financial Stability Map (SOFSM). The task involves five key building blocks: the SOM, crisis dates, vulnerability indicators, a model training framework and a model evaluation framework.

The aim of this chapter is to put forward a framework for creating, as well as to build, a two-dimensional display that represents a high-dimensional financial stability space. The map represents a financial stability cycle consisting of pre-crisis, crisis, post-crisis and tranquil states. Whereas the key aim of the SOFSM is to function as a display for visualizing the state of financial stability, the evaluation of it is performed

This chapter is partly based upon previous research. Please see the following works for further information: Sarlin and Peltonen (2013) and Sarlin (2013c)

with a focus on predictive performance. In particular, it is evaluated in terms of an early-warning model and according to policymakers' preferences between missing a crisis (type I errors) and issuing a false alarm (type II errors). Yet, the creation of the SOFSM in this chapter merely sets a starting point for visualizing threats to financial stability, which is the focus of Chap. 8.

Most parts of this chapter is based upon material in Sarlin and Peltonen (2013). The creation of the SOFSM follows the process of knowledge discovery in databases (KDD) described in Sect. 4.1. The six steps are performed as follows.

- (i) **Domain understanding:** This mainly relates to discussions in previous chapters. Whereas Chap. 2 discussed broadly the domain, Chap. 5 defined the task at hand: to represent high-dimensional data concerning financial entities, be they countries, markets or institutions, on low-dimensional displays to facilitate the identification, assessment and communication of vulnerabilities and risks.
- (ii) **Data understanding:** Whereas Chap. 3 discussed macroprudential data from a broad viewpoint, Sect. 7.1 in this chapter presents the process of collecting data for the task.
- (iii) **Data preparation:** Relating to the previous step, Sect. 7.1 in this chapter also presents the process of transforming and preprocessing the collected data so that they lend to analysis.
- (iv) **Data mining:** The performed data mining makes use of data and dimension reduction methods and follows the basis put forward in Chaps. 4 and 5, as well as Chap. 6. This chapter focuses on using the standard SOM to create the SOFSM, whereas extensions are applied in Chap. 8. In this chapter, Sect. 7.3 introduces a model training framework, which is to be applied in Sect. 7.4.
- (v) **Performance evaluation:** An essential part of the KDD process is to evaluate the performance of models. The model evaluation framework discussed in Sect. 7.2, in addition to other internal quality measures discussed in Sect. 4.4, provides direct means for evaluating the models. The evaluation of the models is discussed in Sects. 7.4 and 7.5.
- (vi) **Knowledge consolidation and deployment:** The final step involves tasks partly outside the scope of this chapter. The SOFSM is exploited in Chap. 8, which also functions as knowledge consolidation and a type of deployment.

7.1 Data

The data used for creating the SOFSM have been chosen with the ultimate goal of representing financial stability as broadly and globally as possible. This obviously imposes challenges in the retrieval of data, as small emerging market economies (EMEs) differ in data provision in comparison to larger advanced economies (AEs). The necessary data for the task consist of vulnerability measures commonly used in the macroprudential literature and binary class information representing pre-crisis, crisis, post-crisis and tranquil periods. In the work in this book, large parts of the main

dataset used follow that in Lo Duca and Peltonen (2013). Quarterly data are collected for 28 countries, 10 AE and 18 EMEs, spanning from 1990Q1–2011Q2. The AEs are Australia, Denmark, the euro area, Japan, New Zealand, Norway, Sweden, Switzerland, the United Kingdom (UK), and the United States (US), while the EMEs are Argentina, Brazil, China, the Czech Republic, Hong Kong, Hungary, India, Indonesia, Malaysia, Mexico, the Philippines, Poland, Russia, Singapore, South Africa, Taiwan, Thailand and Turkey. That is, a multivariate panel dataset, consisting of both a cross-sectional and a temporal dimension, such as the macroprudential data cube in Fig. 3.1. The rationale for using cross-sectional data, rather than creating country-specific models, is threefold: the relatively small number of crisis events in individual countries, the strive to capture a wide variety of crisis types, and the requirement of a global policy approach. Further, results indicate that accounting for country and time-specific effects in early-warning models lead to an improved in-sample fit, while it decreases predictive performance on out-of-sample data (e.g., Fuertes and Kalotychou 2006). Hence, a data vector $x_j \in \mathbb{R}^{18}$ is formed of a class vector $x_{j(cl)} \in \mathbb{R}^4$ and an indicator vector $x_{j(in)} \in \mathbb{R}^{14}$ for each quarter and country in the sample. The data are retrieved from Haver Analytics, Bloomberg and Datastream.

To assess linkages among economies, the events and indicators are complemented with exposures between economies. The focus herein is on the real transmission channel and balance-sheet exposures among economies. The network of financial linkages is based upon external assets (equities and bonds), i.e., holdings of one economy in another, as reported in the Coordinated Portfolio Investment Survey by the International Monetary Fund (IMF). It is worth noting that exposures of central banks are not included due to the different nature of their holdings. Following Sect. 6.2.3, the relationships are expressed in matrix form. Hence, the link between object k and l is represented as element A_{kl} in an $n \times n$ sized matrix, where n is the number of economies. However, the linkages are only used in Chap. 8.

In the following, we focus on the two key types of data: crisis events and vulnerability indicators.

7.1.1 Identifying Systemic Financial Crises

This section explains how systemic financial crises are identified and how all four class variables are defined for enabling assessment of the entire financial stability cycle. The identification of systemic financial crises is done using the Financial Distress Index (FDI) (Lo Duca and Peltonen 2013). This approach provides an objective criterion for the definition of the starting date of a systemic event. While there are several composite indices for measuring financial stress, the FDI differs from most indices by focusing on systemic events. More importantly, the general specification, including the three key market segments, enables applying the FDI to economies of different nature, such as advanced and emerging economies. Financial stress indices for advanced economies, such as the Composite Indicator of Systemic Stress

(CISS) for the euro area (Holló et al. 2012), include a substantially larger number of indicators, and covers also other financial market segments, which may not always be available for emerging markets.

The rationale behind the FDI is that the larger and broader the shock is (i.e., the more systemic the shock), the higher the co-movement among variables reflecting tensions in different market segments. By aggregating five variables to an index that measures stress across market segments, the FDI captures the starting and ending points of a systemic financial crisis. The FDI is a country-specific composite index that covers the money market, equity market and foreign exchange market segments of the domestic financial market:

- (i) the spread of the 3-month interbank rate over the 3-month government bill rate (Ind_1);
- (ii) negative quarterly equity returns (Ind_2);
- (iii) the realized volatility of the main equity index (Ind_3);
- (iv) the realized volatility of the nominal effective exchange rate (Ind_4); and
- (v) the realized volatility of the yield on the 3-month government bill (Ind_5).¹

Each indicator Ind_j for country i at quarter t is transformed into an integer from 0 to 3 according to the quartile of the country-specific distribution, after which the transformed variable is denoted $q_{j,i,t}(\text{Ind}_{j,i,t})$. For example, a value for indicator j falling into the third quartile of the distribution would be transformed to a “2”. The FDI is computed for country i at time t as a simple average of the transformed variables as follows:

$$\text{FDI}_{i,t} = \frac{\sum_{j=1}^5 q_{j,i,t}(\text{Ind}_{j,i,t})}{5} \quad (7.1)$$

To define systemic financial crises, the FDI is first transformed into a binary variable. While discretization leads to some loss of information, it provides a useful means to define the starting and ending points of crises. Moreover, the fundamental idea of predicting vulnerabilities prior to financial crisis, i.e., pre-crisis periods, does not allow modeling a continuous, coinciding index of financial stress.

Hence, a systemic financial crisis is defined as a period of extreme financial stress that has in the past on average been followed by negative consequences for the real economy (i.e., output loss in relation to potential output). One motivation for calibrating the models by choosing average real consequences is to not have a selection bias. If one would only select events with strictly negative real consequences, we would have a selection bias for modeling only events which the policymaker had either failed to predict or the potential policy action that she had taken had not been successful in preventing the negative impact on the real economy. Given that controlling for policy actions are beyond the scope of this book, the level of financial stress

¹ When the 3-month government bill rate is not available, the spread between interbank and T-bill rates of the closest maturity is used. The equity returns are multiplied by minus one, so that negative returns increase stress, while positive returns are set to 0. When computing realized volatilities for components Ind_{3-5} , average daily absolute changes over a quarter are used.

is calibrated to average negative real consequences similar to Lo Duca and Peltonen (2013).

In practice, a binary “crisis” variable, denoted $C0$, is created by assigning it a value 1 in the quarter when the FDI is above the threshold of the 90th percentile of its country-specific distribution $\Omega_i^{90th}(FDI_{i,t})$ and 0 otherwise:

$$C0_{i,t} = \begin{cases} 1 & \text{if } FDI_{i,t} > \Omega_i^{90th} \\ 0 & \text{otherwise} \end{cases} \quad (7.2)$$

This approach identifies a set of 94 systemic events over 1990–2011 for the 28 countries in the sample. To describe the financial stability cycle, the $C0$ variable is turned into a set of other class variables. First, a “pre-crisis” class variable $C18$ is created by setting the binary variable to 1 in the 18 months preceding the systemic financial crisis, and to 0 in all other periods. The pre-crisis variable mimics an ideal leading indicator that perfectly signals a systemic financial crisis in the 18 months before the event. In order to evaluate robustness for different horizons, $C18$ is turned into other pre-crisis class variables, by setting the binary variables $C24$, $C12$ and $C6$ to 1 in the 24, 12 and 6 months before the systemic event and zero otherwise. Similarly, the creation of “post-crisis” class variables $P6$, $P12$, $P18$ and $P24$ take the value 1 in the 6, 12, 18 and 24 months after the systemic event. Finally, when none of the benchmark horizons $C18$, $C0$ and $P18$ take the value 1, then a period is called “tranquil”, denoted as $T0$. Thereby, the class vector $x_{j(ct)} \in \mathbb{R}^4$ consists of the benchmark horizons $C18$, $C0$, $P18$ and $T0$.

7.1.2 Macro-financial Indicators of Vulnerabilities and Risks

The set of indicators consists of commonly used measures in the macroprudential literature for capturing the build-up of vulnerabilities and imbalances in the domestic and global economy (see, e.g., Alessi and Detken 2011; Borio and Lowe 2002, 2004). The key included variables measure asset price developments and valuations, and proxy for credit developments and leverage. In addition, traditional variables (e.g., government budget deficit and current account deficit) are used to control for vulnerabilities stemming from macroeconomic imbalances. Following the indicators used in Lo Duca and Peltonen (2013), this work uses only two of the indicator groups for macroprudential oversight identified in Sect. 3.2: macroeconomic and market-based indicators. With the aim of a global dataset, banking system data are not used due to poor availability for emerging markets.

Following the literature, several transformations of the indicators are constructed to proxy for imbalances, misalignments and a build-up of vulnerabilities. The transformations are levels, annual changes, deviations from short (8 quarters) and long (20 quarters) moving averages, deviations from short ($\lambda = 1600$ in Hodrick-Prescott detrending) and long ($\lambda = 400000$) trends, which results in total into more than 200 indicators. Further, to proxy for global macro-financial imbalances and

vulnerabilities, a set of global indicators are calculated by averaging the transformed variables for the US, the euro area, Japan and the UK. The indicator vector $x_{j(in)} \in \mathbb{R}^{14}$ consists of the best-performing transformation per indicator in terms of their univariate performance in predicting systemic events. The performance is tested with the univariate signaling approach (see Sect. 2.3). The indicators and their summary statistics and transformations are shown in Table 7.1.

Statistical properties of the chosen indicators (Table 7.1) reveal that the data are significantly skewed and non-mesokurtic, and thus do not exhibit normal distributions. To take into account cross-country differences and country-specific fixed effects, this work follows Kaminsky et al. (1998) by measuring indicators in terms of country-specific percentiles. While such outlier trimming is unnecessary for the clustering of the SOM, the even distribution of percentile scales still facilitates judgment and interpretation of the visualization.

Finally, the analysis is conducted in a real-time fashion to the extent possible. Thus, publication lags are taken into account by using lagged variables. For gross domestic product (GDP), money and credit related indicators, the lag ranges from 1 to 2 quarters depending on the country. The variables are also detrended and measured in terms of country-specific percentiles using the latest available information, such that data at time t are only related to data prior to t . Hence, it is worth remembering in the subsequent analyses and visualizations that data refer to the date they are available, rather than the reference period. To test the predictability of the 2007–2008 financial crisis, the sample is split into two sub-samples: the training set spans 1990Q4–2005Q1, while the test set spans 2005Q2–2009Q2.

7.2 Model Evaluation Framework

Crisis data require evaluation criteria that account for their complex nature. Crises are oftentimes outlier events in three aspects:

- (i) they differ significantly from tranquil times,
- (ii) they are commonly more costly, and
- (iii) they occur more rarely.

Given these properties, especially the two latter ones, the evaluation framework in Sarlin (2013c) better resembles the decision problem faced by a policymaker. After briefly reviewing the literature on evaluating early-warning models, we discuss a general framework for deriving a policymaker's loss function and the Usefulness of a model.

While an own strand of literature has focused on the evaluation of early-warning models, the utilized measures seldom cover the wide spectrum of factors that may concern a policymaker. The seminal study by Kaminsky et al. (1998) utilized the

Table 7.1 Statistical properties of the dataset

Type	Variable	Abbreviation	Mean	SD	Min.	Max.	Skew.	Kurt.	KSL	AD
Domestic	Inflation ^a	Inflation	0.89	5.17	-10.15	42.53	4.80	26.72	0.29*	263.90*
Domestic	Real GDP ^b	Real GDP growth	3.73	3.76	-17.54	14.13	-0.86	3.16	0.06*	11.34*
Domestic	Real credit to private sector to GDP ^b	Real credit growth	234.07	4724.00	-69.42	101870.34	20.76	429.59	0.51*	Inf*
Domestic	Real equity prices ^b	Real equity growth	5.93	33.01	-84.40	257.04	0.99	4.31	0.05*	7.28*
Domestic	Credit to private sector to GDP ^a	Leverage	3.48	51.64	-62.78	1673.04	22.76	673.35	0.29*	Inf*
Domestic	Stock market capitalisation to GDP ^a	Equity valuation	3.90	28.32	-62.79	201.55	0.77	2.41	0.03*	3.86*
Domestic	Current account deficit to GDP ^c	CA deficit	-0.02	0.07	-0.27	0.10	-0.98	0.73	0.09*	33.12*
Domestic	Government deficit to GDP ^c	Government deficit	0.01	0.05	-0.19	0.22	-1.09	3.46	0.09*	35.90*
Global	Inflation ^a	Global inflation	0.03	0.64	-1.33	2.29	0.71	1.28	0.08*	12.12*
Global	Real GDP ^b	Global real GDP growth	1.84	1.59	-6.34	4.09	-3.02	11.74	0.20*	122.16*
Global	Real credit to private sector to GDP ^b	Global real credit growth	3.87	1.68	-0.23	7.20	-0.21	-0.31	0.07*	8.82*
Global	Real equity prices ^b	Global real equity growth	2.31	19.08	-40.62	37.77	-0.57	-0.68	0.15*	41.90*
Global	Credit to private sector to GDP ^a	Global leverage	1.15	2.79	-2.79	11.21	1.84	3.40	0.22*	105.26*
Global	Stock market capitalisation to GDP ^a	Global equity valuation	0.89	17.41	-40.54	27.46	-0.50	-0.43	0.09*	19.11*

Notes Transformations: ^a deviation from trend; ^b annual change; ^c level
 KSL Lilliefors' adaption of the Kolmogorov-Smirnov normality test
 AD the standard Anderson-Darling normality test
 Significance levels 1%, *

simple noise-to-signal ratio to set an optimal threshold value.² Based upon Receiver Operating Characteristics (ROC) curves and the area below them, measures applied by Sarlin and Marghescu (2011a) to early-warning model evaluations, Jordá and Taylor (2011) formulated a Correct Classification Frontier (CCF) with advantages like providing visual means and summarizations of results for all possible thresholds. Yet, the measures do not properly pay regard to varying misclassification costs and imbalanced data, and suffer from the fact that some thresholds may be far from policy relevant (e.g., both ends of the CCF). Likewise, while the comprehensive toolbox for evaluating early-warning models by Candelon et al. (2012) provides significant contributions to statistical inference for testing the superiority of one early-warning model over another, they lack an explicit focus on variations in misclassification costs and imbalanced data. A crucial characteristic of measures attempting to grasp a problem of this order of complexity is to explicitly tailor forecasting objectives and validations to the preferences of a decision-maker and the properties of the underlying data.

The literature on the derivation of a policymaker's loss-function has attempted to deal with these so-called low-probability, high-impact events. Demirgüç-Kunt and Detragiache (2000) introduced the notion of a policymaker's loss-function in a banking crisis context, where the policymaker has a cost for preventive actions and type I and II errors (i.e., probability of not receiving a warning conditional on a crisis occurring and of receiving a warning conditional on no crisis occurring). Later, adaptations of this type of loss functions have been introduced to early-warning models for other types of crises, e.g., debt crises (Fuentes and Kalotychou 2007), currency crises (Bussière and Fratzscher 2008), and asset price boom/bust cycles (Alessi and Detken 2011). While Bussière and Fratzscher (2008) still focused on costs of preventive actions, the later literature has mainly focused on the trade-off between type I and II errors. There are two key motivations for focusing on relative preferences between the errors:

- (i) the costs of actions and no actions can be incorporated in preferences between type I and II errors as unrealized benefits can be "rolled up" into error costs (Elkan 2001; Fawcett 2006), and
- (ii) the uncertainty of exact costs associated with preventive actions, false alarms and missing crises.

In addition to a loss function, Alessi and Detken (2011) also propose a Usefulness measure that indicates whether the loss of the prediction is smaller than the loss

² The noise-to-signal ratio is a ratio of the probability of receiving a signal conditional on no crisis occurring to the probability of receiving a signal conditional on a crisis occurring. Demirgüç-Kunt and Detragiache (2000) and El-Shagi et al. (2012) showed that minimizing the noise-to-signal ratio could lead to a relatively high share of missed crisis episodes (i.e., only noise minimization) if crises are rare and the cost of missing a crisis is high. This type of a common corner solution to the optimization problem is mainly due to the fact that the marginal rate of substitution between type I and II errors is unrestricted. Lund-Jensen (2012) concludes the same, and chooses not to use the measure, while Drehmann et al. (2011) choose to minimize the noise-to-signal-ratio subject to at least two thirds of the crises being correctly called. Likewise, Sarlin (2013c) also illustrates such a corner solution.

Table 7.2 A contingency matrix

		Actual class I_j	
		Crisis	No crisis
Predicted class P_j	Signal	A <i>True positive (TP)</i>	B <i>False positive (FP)</i>
	No Signal	C <i>False negative (FN)</i>	D <i>True negative (TN)</i>

of disregarding the model. However, while the above evaluation frameworks have become state-of-the-art, they fail to account for characteristics of imbalanced data.³ In the following, we discuss the use of a loss function and Usefulness measure to better account for the complex nature of crises.

A loss function and usefulness measure. The occurrence of crisis can be represented with a binary state variable $I_j(0) \in \{0, 1\}$ (where observation $j = 1, 2, \dots, N$). Predicting the exact timing of distress does not, however, provide enough reaction time for a policymaker. The wide variety of triggers may also complicate the task of identifying exact timings. To enable policy actions for preventing or decreasing further build-up of vulnerabilities and strengthening the financial system, the focus should rather be on identifying pre-crisis periods $I_j(h) \in \{0, 1\}$ with a specified forecast horizon h . Let $I_j(h)$ be a binary indicator that equals one during pre-crisis periods and zero otherwise. Using univariate or multivariate data, various methods can be used for turning indicators into estimated probabilities of an impending crisis $p_j \in [0, 1]$ (i.e., probability forecasts). To mimic the ideal leading indicator $I_j(h)$, the probability p_j is transformed into a binary point forecast P_j that equals one if p_j exceeds a specified threshold λ and zero otherwise. The correspondence between P_j and I_j can be summarized into a so-called contingency matrix (i.e., frequencies of prediction-realization combinations), as shown in Table 7.2.

From the elements of the above matrix, one can then define various goodness-of-fit measures. The problem is herein approached from the viewpoint of a policymaker, and specific traits related to policymaking.⁴ In a two-class prediction problem, policymakers can be assumed to have relative preferences of conducting two types of errors: issuing false alarms and missing crises. Type I errors represent the probability of not receiving a warning conditional on a crisis occurring $P(p \leq \lambda \mid I_j(h) = 1)$ and type II errors the probability of receiving a warning conditional on no crisis occurring $P(p > \lambda \mid I_j(h) = 0)$. The loss of a policymaker consists of T_1 and T_2

³ While the seminal loss function by Demirgüç-Kunt and Detragiache (2000) accounts for unconditional probabilities, they do not propose a Usefulness measure for the function. Given their complex definition of loss, deriving the Usefulness would not be an entirely straightforward exercise. Further, the version applied in Bussière and Fratzscher (2008) neither accounts for unconditional probabilities nor distinguishes between losses from correct and wrong calls of crisis.

⁴ A further discussion on shaping decision-makers' problems through loss functions, as well as on the relation between statistical and economic value of predictions, can be found in Granger and Pesaran (2000) and Abhyankar et al. (2005).

weighted according to her relative preferences between missing crises ($\mu \in [0, 1]$) and giving false alarms ($1 - \mu$). However, when only using T_1 and T_2 weighted according to relative preferences, we fail to account for imbalances in class size.⁵ Finally, given probabilities p_j of a model, the policymaker should aim at choosing a threshold λ such that her loss is minimized.

The preference parameters may also be derived from a benefit/cost matrix that matches the contingency matrix. A standard 2×2 benefit/cost matrix may easily be manipulated to only include error costs by scaling and shifting entries of columns without affecting the decisions (Elkan 2001; Fawcett 2006). A benefit may be treated as a negative error cost and hence unrealized benefits can be “rolled up” into error costs. For instance, the costs c for the elements of the matrix with two degrees of freedom can be derived to a simpler matrix of class-specific costs c_1 and c_2 with one degree of freedom: $c_1 = c_C - c_A$ and $c_2 = c_B - c_D$ (the subscripts refer to Table 7.2). Most likely, c_B and c_C have a non-negative cost, while c_A and c_D have a non-positive cost. From this, we can derive the relative preferences $\mu = c_1/(c_1 + c_2)$ and $1 - \mu = c_2/(c_1 + c_2)$.

By accounting for unconditional probabilities of crises $P(I_j(h) = 1)$ and tranquil periods $P(I_j(h) = 0) = 1 - P_1$, a loss function is as follows:

$$L(\mu) = \mu T_1 P_1 + (1 - \mu) T_2 P_2 \quad (7.3)$$

As the parameters are unknown *ex ante*, we can use in-sample frequencies to estimate them. Given a threshold λ and forecast horizon h , P_1 and P_2 are estimated with the frequency of the classes ($P_1 = (A + C) / (A + B + C + D)$ and $P_2 = (B + D) / (A + B + C + D)$) and T_1 and T_2 with the error rates ($T_1 = C / (A + C)$ and $T_2 = B / (B + D)$). Using the loss function $L(\mu)$, we can then define the Usefulness of a model. A policymaker could achieve a loss of $\min(P_1, P_2)$ by always issuing a signal of a crisis if $P_1 > 0.5$ or never issuing a signal if $P_2 > 0.5$. However, by weighting with policymakers’ preferences, as she may be more concerned of one of the classes, we achieve the loss $\min(\mu P_1, (1 - \mu) P_2)$ when ignoring the model. First, we derive the absolute Usefulness $U_a(\mu)$ of a model by computing the loss generated by the model subtracted from the loss of ignoring it:

$$U_a(\mu) = \min(\mu P_1, (1 - \mu) P_2) - L(\mu). \quad (7.4)$$

This measure highlights the fact that achieving well-performing, useful models on highly imbalanced data is a difficult task. Hence, already an attempt to build an

⁵ The loss function used by Alessi and Detken (2011) differs from the one introduced here as it assumes equal class size. Their Usefulness measure does, similarly, not account for imbalanced classes, as the loss of disregarding a model depends solely on the preferences. Usefulness measures close to that in Alessi and Detken (2011) have been applied in a large number of works, such as Lo Duca and Peltonen (2013), Sarlin and Marghescu (2011a), El-Shagi et al. (2012), Bisiyas et al. (2012). Similar loss functions have been applied in Fuertes and Kalotychou (2007), Candelon et al. (2012), Lund-Jensen (2012), Knedlik and Schweinitz (2012).

early-warning model with imbalanced data implicitly necessitates a policymaker to be more concerned of the rare class. With a non-perfectly performing model, it would otherwise easily pay-off for the policymaker to always signal the high-frequency class. Second, we compute the share of $U_a(\mu)$ to the maximum possible Usefulness of the model with a measure that is coined relative Usefulness:

$$U_r(\mu) = \frac{U_a(\mu)}{\min(\mu P_1, (1 - \mu) P_2)}. \quad (7.5)$$

That is, $U_r(\mu)$ reports $U_a(\mu)$ as a percentage of the Usefulness that a policymaker would gain with a perfectly performing model. This derives from the fact that if $L(\mu) = 0$ then $U_a(\mu) = \min(\mu P_1, (1 - \mu) P_2)$. The $U_r(\mu)$ provides means for representing the Usefulness as a ratio rather than only reporting a number difficult to judge. In particular, it facilitates comparisons of models for policymakers with different preferences. Within the above framework, we can derive $U_a(\mu)$ and $U_r(\mu)$ for policymakers of different kinds depending on their preferences, which is essentially a parameter to be specified *ad hoc*.

This derives to a cost matrix with costs μ for type I errors and $1 - \mu$ for type II errors. While constants could be added to these entries and their scaling may be modified, this approach favors simplicity. Hence, the rationale for preferring this framework is that it enables setting relative preferences of the errors. Setting specific costs for each entry of the cost matrix is a difficult task in a real-world setting not only because the problem with two degrees of freedom may be difficult to untangle, but also because most often exact values of cost matrix entries are unknown.

In addition to the above framework, the use of pooled panel data motivates including observation-specific costs into the loss function, as the importance of a single country in the evaluation phase may vary depending on the objectives of the policymaker. In an evaluation framework, this leads to a need for weighting entities in terms of their importance, such as systemic relevance or size. The entity-level importance is, however, also a time-varying parameter, and should thus more preferably be defined on the observation level. Although a policymaker's loss function and Usefulness measure that depend on observation-varying costs are shown in Sarlin (2013c), the work in this book focuses only on class-specific costs (that is, does not discriminate between the importance of countries).

Other goodness-of-fit measures. The literature has provided and applied a wide range of goodness-of-fit measures. A large number of them can be defined from the elements of the contingency matrix in Table 7.2. Thus, the following goodness-of-fit measures are used to support the evaluation of models: recall and precision rates, False Positive (FP), True Positive (TP), False Negative (FN) and True Negative (TN) rates, and overall accuracy.⁶ In addition, the global performance of models can be measured using ROC curves and the area under the curve (AUC), i.e., under the ROC

⁶ Recall positives = $TP/(TP + FN)$, Recall negatives = $TN/(TN + FP)$, Precision positives = $TP/(TP + FP)$, Precision negatives = $TN/(TN + FN)$, Accuracy = $(TP + TN)/(TP + TN + FP + FN)$, TP rate = $TP/(TP + FN)$, FP rate = $FP/(FP + TN)$, FN rate = $FN/(FN + TP)$ and TN rate = $TN/(FP + TN)$.

curve. The ROC curve shows the trade-off between the benefits and costs of choosing a certain threshold. When two models are compared, the better model has a higher benefit (expressed in terms of TP rate on the vertical axis) at the same cost (expressed in terms of FP rate on the horizontal axis). In general, the ROC curve plots, for the whole range of measures, the conditional probability of positives to the conditional probability of negatives: $ROC = P(P = 1 | C = 1) / (1 - P(P = 0 | C = 0))$. In this sense, as each FP rate can be associated with a threshold for classifying crisis and tranquil events, the measure shows performance over all thresholds. The size of the AUC is estimated using trapezoidal approximations. It measures the probability that a randomly chosen crisis observation is ranked higher than a randomly chosen tranquil one. A random ranking has an expected AUC of 0.5, while a perfect ranking has an AUC equal to 1.

7.3 Model Training Framework

A key part in modeling in general and dimension reduction in particular is how to train and parametrize the models. In the analysis, a semi-supervised SOM is employed by using data vector $x_j \in \mathbb{R}^{18}$, including class variables (C18, C0, P18 and T0), in training. In contrast to Sarlin and Marghescu (2011a), where only the indicator vector $x_{j(in)} \in \mathbb{R}^{14}$ is used in determining the best-matching units (BMUs), the class vector $x_{j(cl)} \in \mathbb{R}^4$ also has an impact when determining the BMUs in training. By including the class variables in the topology preservation, the projection better separates the classes, which yields the benefit of easier interpretation of the stages of the financial stability cycle. As discussed in Sect. 4.4, this follows the semi-supervised SOMs in general and multi-class supervision of SOMs in particular.

The predictive feature of the model is obtained by assigning to each data point $x_{j(in)} \in \mathbb{R}^{14}$ the C18 (as well as C6, C12 and C24 when testing robustness) value of its BMU.⁷ The performance of a model is then evaluated using the $U_a(\mu)$ and $U_r(\mu)$ for a policymaker. The performance is computed using static and pooled models, i.e., the coefficients or reference vectors m_i are not re-estimated recursively over time and across countries. Following Fuertes and Kalotychou (2006), it can be assumed that by not deriving new models per time unit and country, the parsimonious pooled models better generalize in-sample data and predict out-of-sample data. Although static models have the drawback of ignoring the latest available information, they are a necessity for visualizations of long time series (see Chap. 8). Yet, it is worth noting that recursive re-estimations would computation-wise be feasible when using the model in real-time fashion. Moreover, to account for a possible adjustment process that economic variables go through in between crisis and tranquil periods, i.e., a

⁷ The BMU is the unit that has the shortest Euclidean distance to a data point. When evaluating an already trained SOM model, all data are projected onto the map using only the indicator vector $x_{j(in)} \in \mathbb{R}^{14}$. For each data point, probabilities of a crisis in 6, 12, 18 and 24 months are obtained by retrieving the values of C6, C12, C18 and C24 of its BMU ($m_{b(cl)}$).

crisis and post-crisis bias (Bussière and Fratzscher 2006), the crisis and post-crisis class variables (C0 and P18) are included in SOM training.

The training framework and choice of the SOM specification is implemented with respect to three aspects:

- (i) the model does not overfit the in-sample data (parsimonious);
- (ii) the framework does not include out-of-sample performance (objective); and
- (iii) visualization is taken into account (interpretable).

For a parsimonious benchmark model that avoids overfitting, a logit model similar to the one in Lo Duca and Peltonen (2013) is estimated.⁸ The SOM is parametrized as follows. Whereas the number of units M and radius σ are varied, the map format (75:100) and training length are kept constant. As is recommended by Kohonen (2001) for a stable orientation, this particular map format approximates the ratio of the two largest eigenvalues. Generally, the varied parameters, M and radius σ , have the following effect on performance: an increase in the M value increases the in-sample Usefulness, where $U_r(\mu) \rightarrow \max(U_r(\mu)) = 1$ when $M \rightarrow \infty$, but decreases out-of-sample Usefulness. In fact, if M equals the cardinality of x_j , then perfect in-sample performance may be obtained by each m_i attracting one data point. This would, however, be an overfitted model for out-of-sample prediction. Increases in radius decrease quantization accuracy, and thus in-sample Usefulness, whereas experiments do not show a direct effect on out-of-sample performance. The following training framework assures a focus on a parsimonious, objective and interpretable model:

- (i) Train and evaluate in terms of in-sample $U_a(\mu)$ models for $\sigma = \{\sim 0, 0.3, 0.5, 0.75, 1.0, 1.5, 2.0\}$ and $M = \{50, 100, 150, 200, 250, 300, 400, 500, 600, 1000\}$. For each model, set the threshold on the probability of a crisis such that the $U_a(\mu)$ is maximized. For each M -value, order the models in a descending order.
- (ii) Find for each M -value the first model with in-sample $U_a(\mu)$ equal to or better than that of the benchmark logit model. Choose none of the models if for an M -value all or none of the models' $U_a(\mu)$ exceed that of the logit model.
- (iii) Evaluate the interpretability of the models chosen in Step (ii). Choose the one that is easiest to interpret and has the best topological ordering.

Due to a lack of consensus on a single topology-preservation metric of the SOM projection, it is evaluated following an approach discussed in Kaski et al. (2001). The units m_i are projected into two- and three-dimensional spaces using Sammon's mapping, a distance-preserving mapping from a high-dimensional input space to a

⁸ The logistic regression proceeds as follows. First, it forms a predictor variable which is a linear combination of the explanatory variables. The values of this predictor variable are transformed into probabilities by a logistic function. This logistic function operates through $f(z) = \frac{1}{1+e^{-z}}$, where $z = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \dots + \beta_k x_k$, β_0 is the intercept and $\beta_1 + \beta_2 + \beta_3 + \dots + \beta_k$ are the regression coefficients of $x_1 + x_2 + x_3 + \dots + x_k$, respectively. The value of z measures the total contribution of all the predictor variables used in the model. It is worth to note that $f(z) \rightarrow 0$ when $z \rightarrow -\infty$, and $f(z) \rightarrow 1$ when $z \rightarrow \infty$. Moreover, when $z = 0$, then $f(z) = 0.5$. Thereby, a response curve for a logistic regression is S-shaped.

lower dimension. Topology preservation is defined to be adequate if the map is not twisted at any point and has only adjacent units as neighbors in the Euclidean space. Interpretability is a subjective measure of the SOM visualization defined by the user. The above evaluation framework results in a performance matrix with positions for each M - σ combination, highlights first models per M to outperform the logit model and uses information on topological ordering and interpretability for choosing the final model.

To partition the map into a reduced number of clusters, the units are grouped using Ward's clustering. By performing the clustering on the class variables (C18, C0, P18 and T0), the map is partitioned according to the four stages in the financial stability cycle. This creates four crisp so-called class clusters or financial stability states. The clustering given by lines on a map is, however, oftentimes overlapping, and should thus only be interpreted as an aid in finding the four stages of the financial stability cycle rather than four distinct clusters.

7.4 Training and Evaluation of the SOFSM

This section creates the SOFSM with the help of the five building blocks: the SOM, crisis dates, vulnerability indicators, and model training and model evaluation frameworks. The model training phase starts by estimating a pooled logit model as a benchmark. The logit model is estimated using the quarterly in-sample panel data for 28 countries from 1990Q4–2005Q1. The estimates are reported in Table 7.3 and are later used for predicting out-of-sample data from 2005Q2–2009Q2. In this work, a policymaker is assumed to be more concerned of calling crises, and thus $\mu = 0.8$. A preference parameter of 0.8 belongs to a policymaker who is substantially more concerned about missing a crisis than issuing a false alarm. The rationale for this is that the model is targeted for use in risk identification, leading to further risk assessments, rather than direct policy recommendations. This also follows the historically large costs of financial crises (see Chap. 1) relative to the costs of an internal in-depth investigation of risks and vulnerabilities. On the in-sample data, the pooled logit model has $U_a(\mu) = 0.08$. The training of the SOFSM is performed on the same panel data and the evaluation results are shown in Table 7.4. For $M = 50, 400, 500, 600, 1000$ no model is chosen for analysis, as they never or always exceed the $U_a(\mu)$ of the logit model ($U_a(\mu) = 0.08$). Finally, of the five highlighted models, the one with $M = 150$ and $\sigma = 0.5$ (shown in bold) is selected for its interpretability and topological ordering. The Sammon's mapping used as an aid in judging topological ordering is shown in Fig. 7.1.

The chosen model has 132 units on an 11×12 grid. Figure 7.2 presents the two-dimensional SOFSM that represents the high-dimensional data. By performing Ward's clustering on the class variables, four class clusters are created according to the stages of the financial stability cycle. The upper left cluster represents the pre-crisis cluster (Pre-crisis), the lower left represents the crisis cluster (Crisis), the center and lower-right cluster represents the post-crisis cluster (Post-crisis) and

Table 7.3 The estimates of the logit model

Variable	Estimate	Error	Z	Sig.
Intercept	-6.744	0.612	-11.024	0.000 ^a
Inflation	-0.100	0.300	-0.334	0.738
Real GDP growth	0.076	0.334	0.229	0.819
Real credit growth	-0.001	0.001	-0.613	0.540
Real equity growth	1.791	0.382	4.685	0.000 ^a
Leverage	0.003	0.001	3.204	0.001 ^a
Equity valuation	0.002	0.001	2.689	0.007 ^a
CA deficit	1.151	0.308	3.741	0.000 ^a
Government deficit	0.076	0.342	0.223	0.823
Global inflation	0.207	0.341	0.608	0.543
Global real GDP growth	1.156	0.419	2.761	0.006 ^a
Global real credit growth	0.685	0.381	1.799	0.072 ^c
Global real equity growth	0.832	0.419	1.985	0.047 ^b
Global leverage	0.712	0.427	1.668	0.095 ^c
Global equity valuation	0.959	0.472	2.029	0.042 ^b

Notes Significance levels: 1%, ^a; 5 %, ^b; 10 %, ^c. The model has benchmark specifications of $\mu = 0.8$ and $h = 18$ months

Table 7.4 The evaluation of the SOFSM over M and σ values

M (#Units)	σ (Tension)	0.001	0.3	0.5	0.75	1	1.5	2
50 (52)		0.07	0.06	0.06	0.06	0.06	0.06	0.06
100 (85)		0.08	0.07	0.07	0.06	0.06	0.06	0.06
150 (132)		0.09	0.07	0.08	0.07	0.06	0.06	0.06
200 (188)		0.09	0.09	0.09	0.07	0.07	0.06	0.06
250 (247)		0.09	0.09	0.09	0.07	0.07	0.07	0.06
300 (331)		0.09	0.09	0.09	0.08	0.07	0.07	0.06
400 (408)		0.10	0.10	0.10	0.09	0.09	0.09	0.09
500 (493)		0.11	0.10	0.10	0.10	0.09	0.09	0.09
600 (609)		0.11	0.11	0.10	0.10	0.09	0.09	0.09
1000 (942)		0.11	0.11	0.11	0.10	0.10	0.10	0.10

Notes The table evaluates the SOFSM over M and σ values for the benchmark specifications $\mu = 0.8$ and $h = 18$. Over the neighborhood radii σ , first models to outperform the logit model ($U_a(\mu) = 0.08$) per M value are highlighted in gray and the chosen map is shown in bold. The real number of units is shown in parenthesis since fulfilling the map ratio (75:100) affects the number of units.

the upper right represents the tranquil cluster (Tranquil). Yet, as already noted, the four clusters are overlapping and hence the lines should only be used as an aid in interpreting the map. When maximizing the $U_a(\mu)$ for policymakers with different preferences, Fig. 7.3 shows how the map is classified into two parts, where the shaded area represents early-warning units and the rest tranquil units.




Fig. 7.1 A Sammon's mapping of the SOFSM. *Notes* The figure displays a Sammon's mapping of the SOFSM. The *color coding* corresponds to that in Fig. 7.2

7.5 Performance and Robustness of the SOFSM

Even though the aim of the SOFSM is a two-dimensional display for visualizing threats to financial stability, the evaluation of it is performed with a focus on predictive performance. This section focuses on two types of performance evaluation: comparisons to benchmark models and robustness tests of the SOFSM.

First, this section compares the performance of the semi-supervised SOFSM with an unsupervised counterpart and a logit model. An unsupervised model with the same specifications as the SOFSM is trained to compare their performance. In Table 7.5, the in-sample and out-of-sample performance with the benchmark specifications ($\mu = 0.8$ and C18) are shown for the semi-supervised SOFSM, unsupervised SOFSM (denoted only by the SOM) and the logit model. As anticipated, the unsupervised SOFSM performs to some extent better than the SOFSM along all measures, but also lacks the separation of classes, which is necessary for interpreting the stages of the financial stability cycle. Hence, as this is a key feature of the SOFSM, henceforth the focus is only on comparing the semi-supervised SOFSM and the logit model.

For the benchmark models, the overall performance is similar between the SOFSM and the logit model. On the train set, the SOFSM performs slightly better than the logit model in terms of recall positives, precision negatives and the AUC measure, while the

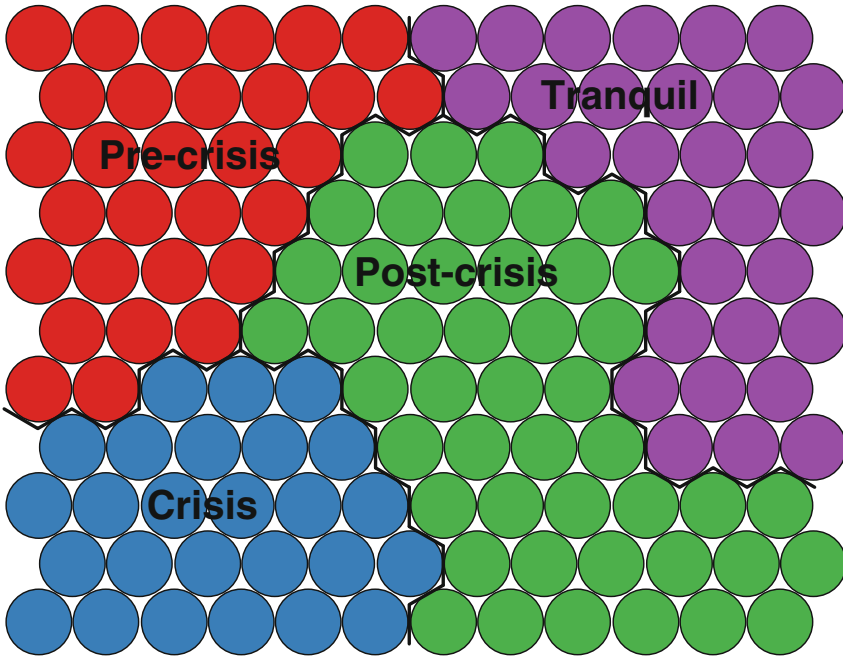


Fig. 7.2 The two-dimensional grid of the SOFSM. *Notes* The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space. The *four clusters* representing financial stability states, distinguished by *lines* and *colors*, are derived using the values of the class variables (C18, C0, P18, T0). Hence, the location on the SOFSM represents the state of financial stability. Distributions of the individual indicators and class variables are shown in Figs. 8.2 and 8.3




Fig. 7.3 Early-warning units for different policymakers’ preferences. *Notes* In the figure, the *shaded area* on the SOFSM (same map as in Fig. 7.2) represents the part of the map that is classified as early-warning units when maximizing the policymakers’ preferences with three different parameter values ($\mu = 0.7$, $\mu = 0.8$ and $\mu = 0.9$) and a horizon of 18 months according to the evaluation framework

logit model outperforms on the other measures, and Usefulness is by definition equal. The classification of the models are of opposite nature, as the SOFSM issues a larger share of false alarms (FP rate = 31 %) than it misses crises (FN rate = 19 %), whereas the logit model misses a larger share of crises (31 %) than it issues false alarms (19 %).

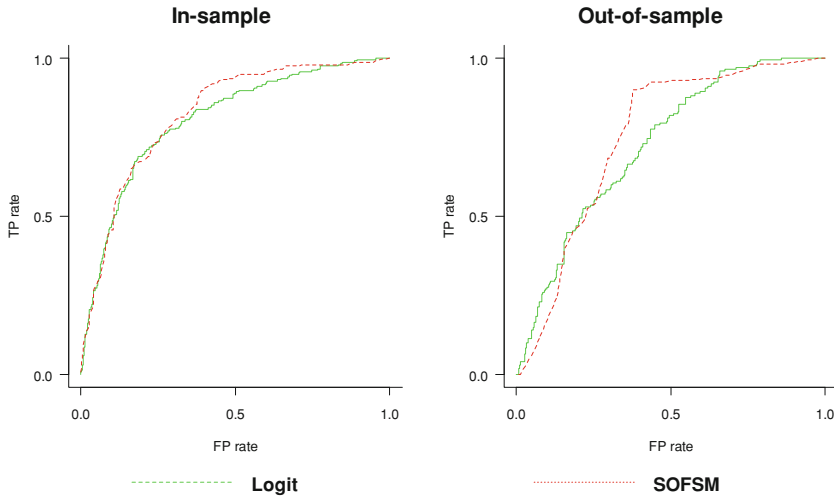


Fig. 7.4 ROC curves for the SOFSM and the logit model. *Notes* The figure shows on in-sample and out-of-sample data ROC curves given policymakers' preferences $\mu = 0.8$ and forecast horizons $h = 18$ months. The *vertical* and *horizontal axes* represent True Positives (TP) rate ($TP/(TP + FN)$) and False Positives (FP) rate ($FP/(FP + TN)$). The AUC, given in Tables 7.5, 7.6 and 7.7, measures the *area below these curves*

That explains also the difference in the overall accuracy, since the class sizes are imbalanced (around 20% pre-crisis and 80% tranquil periods). The performance of the models on the test set differs, in general, similarly as the performance on the train set, except for the SOFSM having slightly higher overall accuracy and Usefulness. This may, in general, be due to the higher share of crisis episodes in the out-of-sample dataset. In terms of out-of-sample $U_r(\mu)$, the SOFSM outperforms the logit model by 10 percentage points and underperforms the unsupervised counterpart by 14% points.

Second, the robustness of the SOFSM is tested with respect to policymakers' preferences ($\mu = 0.7$ and $\mu = 0.9$), forecast horizon (6, 12 and 24 months before a crisis) and thresholds ($\lambda \in [0, 1]$ with the AUC measure). The results of the robustness tests are shown in Tables 7.6, 7.7 and Fig. 7.4. Table 7.6 shows the performance over different policymakers' preferences, Table 7.7 over different forecast horizons and Fig. 7.4 and the second last column of Tables 7.6 and 7.7 over all possible thresholds.

For a policymaker, who is less concerned about issuing false alarms ($\mu = 0.9$), the performance of the models are similar, except for slightly higher Usefulness of the SOFSM compared to the logit model. This confirms that the SOFSM better detects the rare crisis occurrences. For a policymaker, who is less concerned about missing crises ($\mu = 0.7$), the Usefulness of the models is similar, but the nature of the prediction is reversed; the SOFSM issues less false alarms than it misses crises, whereas the logit model issues more false alarms than misses crises.

Table 7.5 In-sample and out-of-sample performance of the benchmark models

Model	Data set	Threshold	TP	FP	TN	FN	Positives		Negatives		Accuracy	$U_a(\mu)$	$U_r(\mu)$	AUC
							Precision	Recall	Precision	Recall				
Logit	Train	0.72	162	190	830	73	0.46	0.69	0.92	0.81	0.79	0.08	0.50	0.81
SOFSM	Train	0.60	190	314	706	45	0.38	0.81	0.94	0.69	0.71	0.08	0.50	0.83
SOM	Train	0.58	215	319	701	20	0.40	0.91	0.97	0.69	0.73	0.09	0.60	0.88
LOGIT	Test	0.72	77	57	249	93	0.57	0.45	0.73	0.81	0.68	0.04	0.27	0.72
SOFSM	Test	0.60	112	89	217	58	0.56	0.66	0.79	0.71	0.69	0.06	0.37	0.75
SOM	Test	0.58	139	95	211	31	0.59	0.82	0.87	0.69	0.74	0.08	0.51	0.76

Notes The table reports results for the logit, semi-supervised SOFSM and unsupervised SOM on the train and test datasets, as well as the optimal threshold, for the benchmark specifications $\mu = 0.8$ and $h = 18$. To assess the performance of the models, the table also reports in columns the following measures: $TP = Truepositives$, $FP = Falsepositives$, $TN = Truenegatives$, $FN = Falsenegatives$, $Precisionpositives = TP/(TP + FP)$, $Recallpositives = TP/(TP + FN)$, $Precisionnegatives = TN/(TN + FP)$, $Accuracynegatives = (TP + TN)/(TP + TN + FP + FN)$, $U = Usefulness$ (see Sect. 7.2), $AUC =$ area under the ROC curve (TP rate to FP rate, see Sect. 7.2). For further information on the evaluation measures, see Sect. 7.2. The best accuracy measure, as per data set and evaluation measure, is shown in bold.

Table 7.6 Robustness tests for different preference values μ

Model	Data set	μ	Threshold	TP	FP	TN	FN	Positives		Negatives		Accuracy	$U_a(\mu)$	$U_r(\mu)$	AUC
								Precision	Recall	Precision	Recall				
Logit	Train	0.7	0.72	162	190	830	73	0.46	0.69	0.92	0.81	0.79	0.06	0.39	0.81
SOFSM	Train	0.7	0.75	153	166	854	82	0.48	0.65	0.91	0.84	0.80	0.05	0.37	0.83
Logit	Train	0.8	0.72	162	190	830	73	0.46	0.69	0.92	0.81	0.79	0.08	0.50	0.81
SOFSM	Train	0.8	0.60	190	314	706	45	0.38	0.81	0.94	0.69	0.71	0.08	0.50	0.83
Logit	Train	0.9	0.54	197	381	639	38	0.34	0.84	0.94	0.63	0.67	0.03	0.30	0.81
SOFSM	Train	0.9	0.50	214	419	601	21	0.34	0.91	0.97	0.59	0.65	0.04	0.32	0.83
Logit	Test	0.7	0.72	77	57	249	93	0.57	0.45	0.73	0.81	0.68	0.01	0.08	0.72
SOFSM	Test	0.7	0.75	76	56	250	94	0.58	0.45	0.73	0.82	0.68	0.01	0.07	0.75
Logit	Test	0.8	0.72	77	57	249	93	0.57	0.45	0.73	0.81	0.68	0.04	0.27	0.72
SOFSM	Test	0.8	0.60	112	89	217	58	0.56	0.66	0.79	0.71	0.69	0.06	0.37	0.75
Logit	Test	0.9	0.54	110	109	197	60	0.50	0.65	0.77	0.64	0.64	0.01	0.13	0.72
SOFSM	Test	0.9	0.50	134	109	197	36	0.55	0.79	0.85	0.64	0.70	0.03	0.27	0.75

Notes The table tests robustness on in-sample and out-of-sample data for different μ values given $h = 18$. See the notes for Table 7.5. The best accuracy measure, as per method, for each forecast horizon and evaluation measure, is shown in bold.

Table 7.7 Robustness tests for different horizons h

Model	Data set	Horizon	Threshold	TP	FP	TN	FN	Positives		Negatives		Accuracy	$U_a(\mu)$	$U_r(\mu)$	AUC
								Precision	Recall	Precision	Recall				
Logit	Train	C6	0.72	70	282	882	21	0.20	0.77	0.98	0.76	0.76	0.08	0.53	0.81
SOFSM	Train	C6	0.51	88	530	634	3	0.14	0.97	1.00	0.54	0.58	0.08	0.51	0.83
Logit	Train	C12	0.72	117	235	855	48	0.33	0.71	0.95	0.78	0.77	0.08	0.49	0.80
SOFSM	Train	C12	0.69	123	267	823	42	0.32	0.75	0.95	0.76	0.75	0.08	0.50	0.84
Logit	Train	C18	0.72	162	190	830	73	0.46	0.69	0.92	0.81	0.79	0.08	0.50	0.81
SOFSM	Train	C18	0.60	190	314	706	45	0.38	0.81	0.94	0.69	0.71	0.08	0.50	0.83
Logit	Train	C24	0.58	242	286	673	54	0.46	0.82	0.93	0.70	0.73	0.08	0.52	0.81
SOFSM	Train	C24	0.63	233	241	718	63	0.49	0.79	0.92	0.75	0.76	0.08	0.54	0.85
Logit	Test	C6	0.72	18	116	302	40	0.13	0.31	0.88	0.72	0.67	0.00	0.03	0.57
SOFSM	Test	C6	0.51	47	205	213	11	0.19	0.81	0.95	0.51	0.55	0.05	0.32	0.65
Logit	Test	C12	0.72	49	85	275	67	0.37	0.42	0.80	0.76	0.68	0.03	0.19	0.64
SOFSM	Test	C12	0.69	51	102	258	65	0.33	0.44	0.80	0.72	0.65	0.02	0.16	0.68
Logit	Test	C18	0.72	77	57	249	93	0.57	0.45	0.73	0.81	0.68	0.04	0.27	0.72
SOFSM	Test	C18	0.60	112	89	217	58	0.56	0.66	0.79	0.71	0.69	0.06	0.37	0.75
Logit	Test	C24	0.58	132	68	185	91	0.66	0.59	0.67	0.73	0.67	0.05	0.32	0.76
SOFSM	Test	C24	0.63	150	51	202	73	0.75	0.67	0.73	0.80	0.74	0.07	0.47	0.80

Notes The table tests robustness on in-sample and out-of-sample data for different forecast horizons h given the policymakers' preferences $\mu = 0.8$. See the notes for Table 7.5
 The best accuracy measure, as per method, for each preference value and evaluation measure, is shown in bold.

Over different forecast horizons, the in-sample performance is generally similar. However, the out-of-sample Usefulness, with the exception of forecast horizon of 12 months (C12), is better for the SOFSM than for the logit model. Interestingly, the logit model fails to yield any Usefulness ($U_r(\mu) = 0.03$) at a forecast horizon of 6 months. Finally, the AUC measure, which summarizes the performance of a model over all thresholds, can be computed for all models by calculating the areas under the ROC curves, such as those shown in Fig. 7.4 for the benchmark models ($\mu = 0.8$ and C18). It is the only measure to consistently show superior performance for the SOFSM. A caution regarding the AUC measure is, however, that parts of the ROC curve that are not policy relevant are included in the computed area. When comparing Usefulness for each pair of models, the SOFSM shows consistently equal or superior performance except for a single out-of-sample evaluation with a forecast horizon of 12 months. To sum up, we can conclude that the SOM performs, in general, as well as or better than a logit model in both classifying the in-sample data and in predicting out-of-sample the global financial crisis that started in 2007.

7.6 Concluding Summary

The essence of this chapter was to describe how the SOFSM is created. The general framework used for creating the SOFSM consists of five building blocks: the SOM, crisis dates, vulnerability indicators, a model training framework and a model evaluation framework. This chapter has discussed the identification of systemic financial crises, the use of macro-financial vulnerabilities, risks and imbalances, and model evaluation and training frameworks. However, the general framework should not be restricted to precise definitions of the building blocks used herein. For an application with another focus, the components should obviously be defined differently. For instance, the choice of explanatory variables and the dating of financial crises should be designed according to the task at hand, such as the events being banking, debt or currency crises and vulnerabilities being indicators measuring banking systems, solvency or exchange-rate pressure. Likewise, the framework could be applied to firm-level data, where the events could be bank failures and indicators financial ratios based upon balance-sheet and income-statement data. Ironically, the view of a “financial stability cycle” could still apply, as banks tend not to disappear due to a failure.

The outcome of this chapter is a two-dimensional display for visualizing the high-dimensional state of financial stability. Hence, this chapter only provides a basis for monitoring threats to financial stability, whereas this display can be used as a groundwork for a wide range of tasks. The following chapter focuses on exploiting the SOFSM for assessing and identifying the three key systemic risks of macroprudential oversight: (i) endogenous build-up of widespread imbalances (*early-warning models*); (ii) exogenous aggregate shocks (*macro stress-testing models*); and (iii) contagion and spillover (*contagion and spillover models*).

References

- Abhyankar, A., Sarno, L., & Valente, G. (2005). Exchange rates and fundamentals: Evidence on the economic value of predictability. *Journal of International Economics*, 66, 325–348.
- Alessi, L., & Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity. *European Journal of Political Economy*, 27(3), 520–533.
- Bisias, D., Flood, M., Lo, A., & Valavanis, S. (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics*, 4, 255–296.
- Borio, C., Lowe, P., 2002. Asset prices, financial and monetary stability: Exploring the nexus, BIS Working Papers No. 114.
- Borio, C., Lowe, P., 2004. Securing sustainable price stability: Should credit come back from the wilderness?, BIS Working Papers No. 157.
- Bussière, M., & Fratzscher, M. (2006). Towards a new early warning system of financial crises. *Journal of International Money and Finance*, 25(6), 953–953.
- Bussière, M., & Fratzscher, M. (2008). Low probability, high impact: Policy making and extreme events. *Journal of Policy Modeling*, 30, 111–111.
- Candelon, B., Dumitrescu, E., & Hurlin, C. (2012). How to evaluate an early-warning system: Toward a unified statistical framework for assessing financial crises forecasting methods. *IMF Economic Review*, 60(1), 75–113.
- Demirgüç-Kunt, A., & Detragiache, E. (2000). Monitoring banking sector fragility: A multivariate logit. *World Bank Economic Review*, 14(2), 287–307.
- Drehmann, M., Borio, C., & Tsatsaronis, K. (2011). Anchoring countercyclical capital buffers: The role of credit aggregates. *International Journal of Central Banking*, 7(4), 189–240.
- El-Shagi, M., Knedlik, T., von Schweinitz, G., 2012. Predicting financial crises: The (statistical) significance of the signals approach, IWH Discussion Papers No. 3.
- Elkan, C. (2001). The foundations of cost-sensitive learning. Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 01) (pp. 973–978). USA: Seattle.
- Fawcett, T. (2006). ROC graphs with instance-varying costs. *Pattern Recognition Letters*, 27(8), 882–891.
- Fuertes, A.-M., & Kalotychou, E. (2006). Early warning system for sovereign debt crisis: The role of heterogeneity. *Computational Statistics and Data Analysis*, 5, 1420–1441.
- Fuertes, A.-M., & Kalotychou, E. (2007). Towards the optimal design of an early warning system for sovereign debt crises. *International Journal of Forecasting*, 23(1), 85–100.
- Granger, C., & Pesaran, M. (2000). Economic and statistical measures of forecast accuracy. *Journal of Forecasting*, 19, 537–560.
- Holló, D., Kremer, M., & Lo Duca, M. (2012). *CISS - a composite indicator of systemic stress in the financial system*. ECB working paper, No. 1426.
- Jordá, O., & Taylor, A. (2011). *Performance evaluation of zero net-investment strategies*. NBER working paper, No. 17150.
- Kaminsky, G., Lizondo, S., & Reinhart, C. (1998). Leading indicators of currency crises. *IMF Staff Papers*, 45(1), 1–48.
- Kaski, S., Venna, J., & Kohonen, T. (2001). Coloring that reveals cluster structures in multivariate data. *Australian Journal of Intelligent Information Processing Systems*, 60, 2–88.
- Knedlik, T., & von Schweinitz, G. (2012). Macroeconomic imbalances as indicators for debt crises in Europe. *Journal of Common Market Studies*, 50(5), 726–745.
- Kohonen, T. (2001). *Self-organizing maps* (3rd ed.). Berlin: Springer.
- Lo Duca, M., & Peltonen, T. (2013). Assessing systemic risks and predicting systemic events. *Journal of Banking & Finance*, 37(7), 2183–2195.
- Lund-Jensen, K. (2012). *Monitoring systemic risk based on dynamic thresholds*. IMF working paper, No. 12/159.
- Sarlin, P. (2013). On policymakers' loss functions and the evaluation of early warning systems. *Economics Letters*, 119(1), 1–7.

- Sarlin, P., & Marghescu, D. (2011). Neuro-genetic predictions of currency crises. *Intelligent Systems in Accounting, Finance and Management*, 18(4), 145–160.
- Sarlin, P., & Peltonen, T. (2013). Mapping the state of financial stability. *Journal of International Financial Markets, Institutions & Money*, 26, 46–76.

Chapter 8

Exploiting the SOFSM

Lastly, novel methods such as self-organising financial stability maps provide an alternative means of gauging systemic stress through visual means—thereby providing a useful complement to numerical signalling methodologies.

– Vítor Constâncio, Vice-President of the ECB,
Frankfurt am Main, 18 November 2010

This chapter exploits the Self-Organizing Financial Stability Map (SOFSM) for tasks in macroprudential oversight. The SOFSM was created in Chap. 7, whereas the Self-Organizing Map (SOM) extensions used for exploiting it were introduced in Chap. 6. The tasks performed with the SOFSM are two, risk identification and assessment, of which the former is supported by early-warning models and the latter by macro stress-testing and contagion or spillover models. The three models target the three respective forms of systemic risk: widespread imbalances, aggregate shocks and contagion and spillover risk. Drawing upon Sarlin and Peltonen (2013) and Sarlin (2013), the SOFSM is exploited by the means of the following eight approaches (where the numbering refers to sections and the *parenthesis* represents the addressed systemic risk).¹

- 9.1 Assessing distributions of the macro-financial indicators and all class variables with the help of the feature planes of the SOFSM (*imbalances*).
- 9.2 Mapping the state of financial stability for individual data and aggregates by the means of labels and trajectories on the SOFSM (*imbalances*).

This chapter is partly based upon previous research. Please see the following works for further information: Sarlin and Peltonen (2013), Sarlin (2013, 2014a)

¹ Beyond the static representations herein, the implementation developed by infolytika provides an interactive, web-based interface to the SOFSM (<http://risklab.fi/demo/macropru/fsm/>). For a description, see Sarlin (2014a).

- 9.3 Fuzzification of the SOFSM for visualizing temporal belongingness to financial stability states of individual data and class distance structures on the map (*imbalances*).
- 9.4 Probabilistic modeling of state transitions on the SOFSM for visualizing probabilities of transition to financial stability states of individual data and for assessing the cyclical and temporal structure of the financial stability cycle (*imbalances*).
- 9.5 Scenario analysis for economies on the SOFSM by assessing the effects of positive and negative shocks, both domestically and globally (*aggregate shocks*).
- 9.6 Using superimposed portfolio network topologies and neighborhoods on the SOFSM to assess the spread of financial distress and shock propagation (*contagion and spillover*).
- 9.7 Computing distances between data and their mean profiles on the SOFSM to find extreme events and imbalances in economies' macro-financial conditions (*imbalances*).
- 9.8 Complementing the SOFSM with a solely predictive model that uses genetically optimized neural networks for the identification of risks (*imbalances*).

Figure 8.1 relates the eight means for exploiting the SOFSM, as well as the Self-Organizing Time Map (SOTM) in the subsequent chapter, to risk identification, risk assessment and risk communication. The red components represent risks and vulnerabilities, the green components represent the need for risk identification, assessment or communication, and the blue frame marks the contributions of this book. The final ingredient of the process highlights the need for visualization tools not only for internal communication, but also for external risk warnings, policy recommendations and Financial Stability Reports in general. The figure illustrates separate feedback loops of internal and external risk communication, where the **solid** black line shows that internal communication interacts with risk identification and assessment (green components), and the **dashed** black line shows that external communication has effects on potential sources of systemic risk, vulnerabilities, and material risks (red components). The lack of focus on aids for risk communication directly follows from the literature review in Sect. 8.1 and the conventional macroprudential oversight process presented in Sect. 2.4 (see Fig. 2.4). Hence, the figure highlights the importance of visual means for external communication of the results of risk identification and assessment tools, in addition to the visuals' inherent properties of amplifying cognition and understanding of policymakers in the internal monitoring process.

The creation of the SOFSM in the previous chapter was related to the process of knowledge discovery in databases (KDD), in which this chapter was mainly positioned as the final step of knowledge consolidation and deployment. Yet, this chapter can also be related to the visual analytics process introduced in Chap. 4. There is a direct link to the fields of information visualization (Card et al. 1999) and visual analytics (Thomas and Cook 2005) in that the tools illustrated herein provide means to amplify cognition through visual representations, as well as a combination with analytical reasoning and methods. The connection can be illustrated by Keim's (2006) visual analytics mantra: "*Analyze first, show the important, zoom, filter and analyze further, details on demand*". The SOFSM created in the previous chapter provides

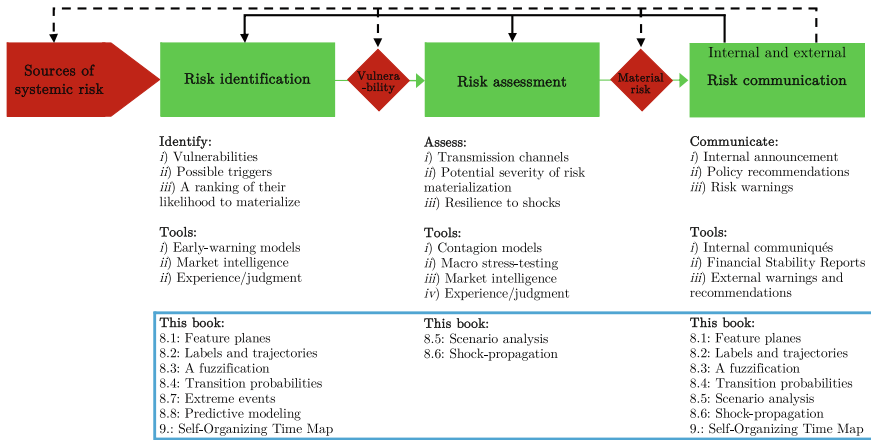


Fig. 8.1 The SOFSM for risk identification, assessment and communication. *Notes* The figure represents the role of the tools in this book in the process of risk identification, assessment and communication. The *red components* represent risks and vulnerabilities, the *green components* represent the need for risk identification, assessment or communication, and the *blue frame* marks the location of contributions of the work in this book. The feedback loops of internal and external risk communication are illustrated with *solid black lines* and the *dashed black lines*, respectively

an analytical solution for the first step of *analyze first*. In this chapter, many of the visualizations on the SOFSM provide means for *showing the important, zooming, and filtering* (e.g., mappings of individual data). Moreover, the analytical approaches put forward in this chapter provide means for *analyzing further*, after which *details on demand* can be viewed. Thus, the mantra involves automated analytical analysis before and after the use of visual representations.

From the viewpoint of macroprudential oversight, each of the following sections discusses how that particular approach aids in either risk identification or assessment, given an ultimate aim of risk communication.

8.1 The SOFSM: Its Output and Interpretation

This section presents the output of the SOFSM and an interpretation of it. In particular, the SOFSM is used for describing the four states of the financial stability cycle. For this purpose, we can make use of Figs. 7.2, 8.2 and 8.3, in which the SOFSM, feature planes for the 14 macro-financial indicators and the main classes, and feature planes for all the class variables are shown, respectively. Figure 7.2 displays the two-dimensional SOFSM that represents a high-dimensional financial stability space. The feature planes in Figs. 8.2 and 8.3 are layers of the SOFSM. Figure 8.2 shows the distribution of the indicators and the four main class variables (Pre-crisis, Crisis,

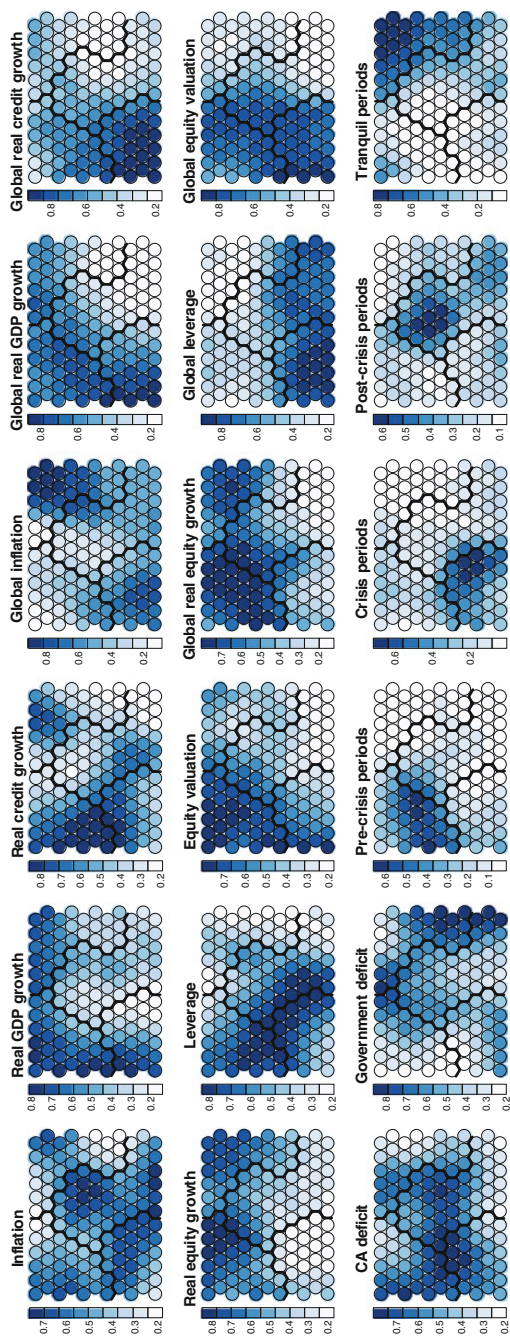


Fig. 8.2 Feature planes for the SOFSM. *Notes:* The figure shows feature planes for the 14 indicators and the benchmark class variables. The feature planes are layers of the SOFSM in Fig. 7.2. While the indicators are defined in Table 7.1, the four main class variables are Pre-crisis (C18), Crisis (C0), Post-crisis (P18) and Tranquil periods (T0). As each data vector consists of 14 indicators and 4 main class variables, these feature planes show the distribution of each data column on the SOFSM grid. In the case of binary class variables that take values 1 and 0, high values represent a high proportion of data in different periods (pre-crisis, crisis, post-crisis or tranquil periods). These views highlight the fact that location on the SOFSM represents the state of financial stability, where each location can be associated with variable values




Fig. 8.3 Feature planes for all classes. *Notes* The figure shows the distributions of different pre- and post-crisis horizons. As in Fig. 8.2, these are layers of the SOFSM in Fig. 7.2. The feature planes C24, C18, C12, C6, P24, P18, P12 and P6 shows the map distribution of class variables that represent 24, 18, 12 and 6 months before and after a crisis, respectively. While C0 and T0 shows the distribution of crisis and tranquil periods, PPC0 represents the co-occurrence of pre- and post-crisis periods

Post-crisis and Tranquil periods), whereas the feature planes in Fig. 8.3 shows the distribution of the classes on the SOFSM. An assessment of these figures may aid in understanding relations among the variables, including all macro-financial indicators and all classes, which is a key ingredient of risk identification.

In contrast to early-warning models using binary classification methods, such as discrete choice techniques, the SOFSM enables simultaneous assessment of the associations with all four stages of the financial stability cycle, i.e., class clusters. Thus, new models need not be derived for different forecast horizons or definitions of the dependent variable. The feature planes in Figs. 8.2 and 8.3, which disentangle the individual vulnerabilities and risks of the SOFSM in Fig. 7.2, enable one to directly detect signals of a crisis (or any of the four states). For instance, the following strong associations are found. First, we can differentiate between “early” and “late” signs of a crisis by assessing differences within the pre-crisis cluster. The strongest early signs of a crisis (upper right part of the cluster) are high domestic and global real equity growth and equity valuation, while most important late signs of a crisis (lower left part of the cluster) are domestic and global real gross domestic product (GDP) growth, and domestic real credit growth, leverage, budget surplus, and current account deficit. Second, the highest values of global leverage and real credit growth in the crisis cluster exemplify the fact that increases in some indicators may reflect a rise in financial stress only up to a specific threshold. Increases beyond that level are, in these cases, more concurrent than preceding signals of a crisis. Similarly, budget deficits characterize the late post-crisis and early tranquil periods. The characteristics of the financial stability states are summarized in Table 8.1 through summary statistics.

Table 8.1 Summary statistics of the financial stability states

Variable	Pre crisis		Crisis		Post crisis		Tranquil	
	Centre	Range	Centre	Range	Centre	Range	Centre	Range
Inflation	0.49	[0.22, 0.66]	0.55	[0.30, 0.69]	0.59	[0.26, 0.76]	0.37	[0.17, 0.68]
Real GDP growth	0.67	[0.40, 0.80]	0.48	[0.14, 0.83]	0.34	[0.25, 0.50]	0.53	[0.30, 0.72]
Real credit growth	0.66	[0.28, 0.85]	0.55	[0.35, 0.82]	0.39	[0.18, 0.68]	0.43	[0.21, 0.75]
Real equity growth	0.68	[0.41, 0.85]	0.28	[0.16, 0.58]	0.39	[0.23, 0.80]	0.61	[0.40, 0.74]
Leverage	0.63	[0.31, 0.80]	0.59	[0.37, 0.81]	0.52	[0.23, 0.83]	0.29	[0.18, 0.51]
Equity valuation	0.73	[0.62, 0.80]	0.55	[0.27, 0.81]	0.33	[0.17, 0.66]	0.45	[0.30, 0.63]
CA deficit	0.58	[0.30, 0.78]	0.54	[0.26, 0.80]	0.48	[0.25, 0.77]	0.41	[0.19, 0.66]
Government deficit	0.38	[0.19, 0.74]	0.45	[0.22, 0.62]	0.53	[0.32, 0.85]	0.61	[0.26, 0.85]
Global inflation	0.33	[0.08, 0.61]	0.61	[0.34, 0.76]	0.46	[0.20, 0.79]	0.63	[0.11, 0.90]
Global real GDP growth	0.67	[0.54, 0.74]	0.67	[0.30, 0.86]	0.29	[0.13, 0.69]	0.45	[0.13, 0.71]
Global real credit growth	0.55	[0.28, 0.77]	0.86	[0.61, 0.92]	0.37	[0.16, 0.67]	0.33	[0.15, 0.52]
Global real equity growth	0.72	[0.47, 0.80]	0.4	[0.23, 0.63]	0.34	[0.11, 0.79]	0.54	[0.20, 0.73]
Global leverage	0.35	[0.18, 0.60]	0.79	[0.57, 0.91]	0.58	[0.17, 0.77]	0.33	[0.16, 0.73]
Global equity valuation	0.67	[0.48, 0.82]	0.81	[0.54, 0.91]	0.36	[0.14, 0.76]	0.27	[0.19, 0.55]

Notes Columns represent characteristics (cluster center and range) of the financial stability states on the SOFSM and rows represent indicators. Since data are transformed to country-specific percentiles, the summary statistics are comparable across indicators and clusters

In the remainder of this chapter, the SOFSM is mostly used for mapping the state of financial stability by combining the SOFSM display with data concerning economies' macro-financial conditions

8.2 Visualizing the State of Financial Stability on the SOFSM

In this section, cross-sectional and temporal samples of the panel dataset are mapped on the two-dimensional SOFSM. Aggregates for groups of countries are also computed in order to explore the state of financial stability globally, in advanced economies and in emerging economies. Data points are mapped onto the grid by projecting them to their best-matching units BMUs using only the indicator vector $x_{j(in)} \in \mathbb{R}^{14}$. Trajectories of consecutive time-series data are shown with arrows. These mappings provide means for visualizing the state of financial stability on the SOFSM, which is clearly linked to not only the task of risk identification, but also provides direct means for risk communication.




Fig. 8.4 A mapping of the US and the euro area. *Notes* The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Fig. 7.2). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. Data points are mapped onto the grid by projecting them to their BMUs using only macro-financial indicators. Consecutive time-series data are linked with arrows. The data for both the US and the euro area represent the first quarters of 2002–2011 as well as the second quarter of 2011

For a simultaneous temporal and comparative analysis, the state of financial stability is mapped based upon the evolution of macro-financial conditions for the United States (US) and the euro area in Fig. 8.4. The data for both economies represent the first quarters of 2002–2010 and the final point of the sample, 2011Q2. Without a precise empirical treatment for accuracy, the map well recognizes for both economies the pre-crisis, crisis and post-crisis stages of the financial stability cycle by circulating around the map during the analyzed period. The early-warning units in Fig. 7.3 confirm that even a policymaker with $\mu \leq 0.7$ would have correctly predicted crises in both economies. Interestingly, the euro area is located in the tranquil cluster in 2010Q1. This indicates that the aggregated macro-financial measures for the euro area as a whole did not reflect the elevated risks in the euro area periphery at that point in time. However, it also coincides with a relatively low Financial Distress Index (FDI) for the aggregate euro area. This can be explained by the weaknesses and financial stress in smaller economies being averaged out by improved macro-financial conditions in larger euro area economies, highlighting the importance of

country-level analysis. As the SOFSM is flexible with respect to input data, it is of central importance that the included set of vulnerability indicators capture the particular events of interest. The macro-financial vulnerabilities currently used are best suited for capturing the build-up of vulnerabilities in the form of boom-bust cycles. However, they are less useful in identifying situations, where, for example, bank funding constraints or counterparty risks in a post-crisis recovery phase cause elevated financial stress that feeds back to the real economy, increasing the probability of a financial crisis. Furthermore, by using the traditional macro-financial vulnerabilities, it is rather difficult to capture situations where, as in the ongoing debt crisis, self-fulfilling expectations drive the equilibrium outcomes. Nevertheless, the euro area has moved to the border of the pre-crisis cluster in 2010Q4, and to an adjacent unit in 2011Q1 and Q2. This reflects the ongoing sovereign and banking crises as with $\mu \leq 0.7$ this particular location is an early-warning unit (see Fig. 7.3). The (US) is located in the post-crisis cluster in 2010Q1 and in the tranquil cluster in 2011Q2. Figure 8.5 represents a cross-sectional mapping of the state of financial stability for all countries in 2010Q3 and in 2011Q2, which is the latest data point in the analysis. In 2010Q3, the countries are divided into three groups of financial stability states. The map indicates elevated risks in several emerging market economies (Mexico, Turkey, Argentina, Brazil, Taiwan, Malaysia and the Philippines), while most of the advanced economies are in the lower right corner of the map (post-crisis and tranquil cluster). Three countries (Singapore, South Africa and India) are located on the border of the tranquil and pre-crisis clusters, which is an indication of a possible future transition to the pre-crisis cluster. Interestingly, in 2011Q2, most economies are located in the tranquil cluster, while the euro area has the highest financial stress by being located close to the pre-crisis cluster.

Further, the state of financial stability is mapped for three aggregates: the world, emerging market economies and advanced economies. The state of financial stability for the aggregates is computed by weighting the indicators for the countries in our sample using stock market capitalization to proxy their financial importance. Hence, an aggregated data vector is computed as follows: $x_{agg(i,t)} = \sum_{i=1}^I (w_{i,t}/W_t)x_{i,t}$, where $x_{i,t}$ is a data vector for country i at time t , $w_{i,t}$ is stock market capitalization, W is aggregated stock market capitalization and I represents all countries. These aggregates can, like any data point, be projected onto the map to their BMU.

The upper map in Fig. 8.6 shows the evolution of global macro-financial conditions in the first quarters of 2002–2011. The global state of financial stability enters the pre-crisis cluster in 2006Q1 and the crisis cluster in 2007Q1. It moves via the post-crisis and tranquil cluster back to the post-crisis cluster in 2011Q1. This coincides with the global evolution of the FDI. More interestingly, the model signals out of sample a global financial crisis as early as in 2006Q1. The separation of the global aggregate into emerging market and advanced economies is shown in the lower map in Fig. 8.6. The mapping of the advanced economy aggregate is very similar to the one of the world aggregate, which is mainly a result of the high share of stock market capitalization of the advanced economies. Notably, the movements of the financial stability states of the emerging markets are also similar to those in the

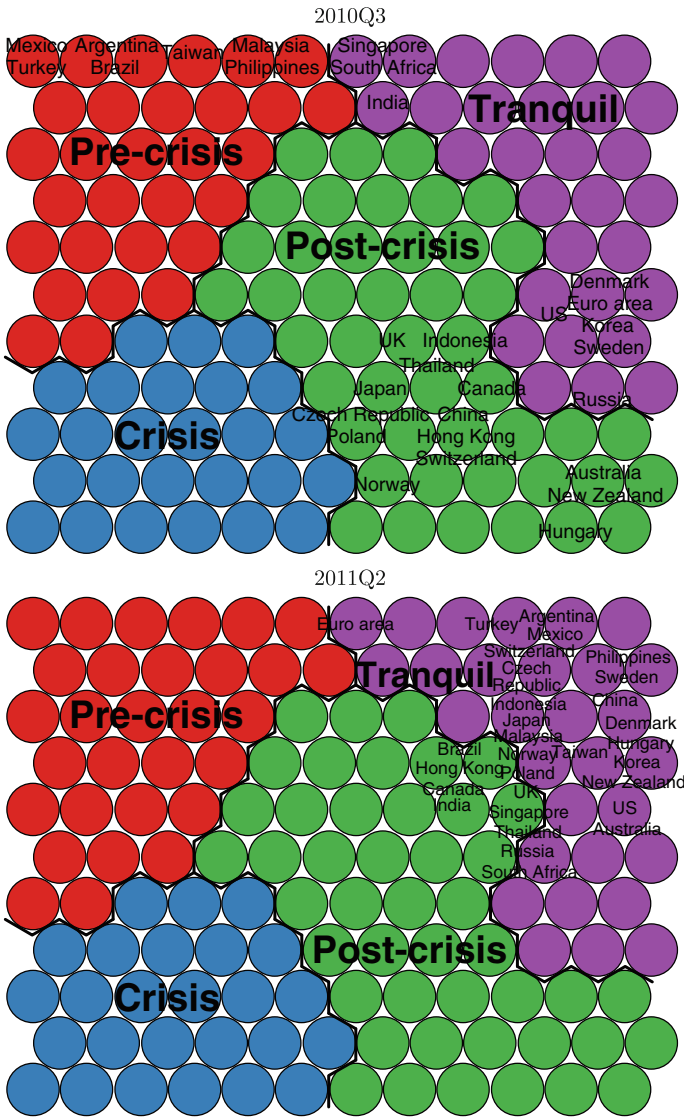


Fig. 8.5 A mapping of all countries in 2010Q3 and 2011Q2. *Notes* The figure presents a cross-sectional mapping of financial stability states for all countries in the sample in 2010Q3 and 2011Q2. The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Fig. 7.2). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. Data points are mapped onto the grid by projecting them to their BMUs using only macro-financial indicators, but positions are approximate to fit all labels. The data for all economies represent the third quarter of 2010 and the second quarter of 2011

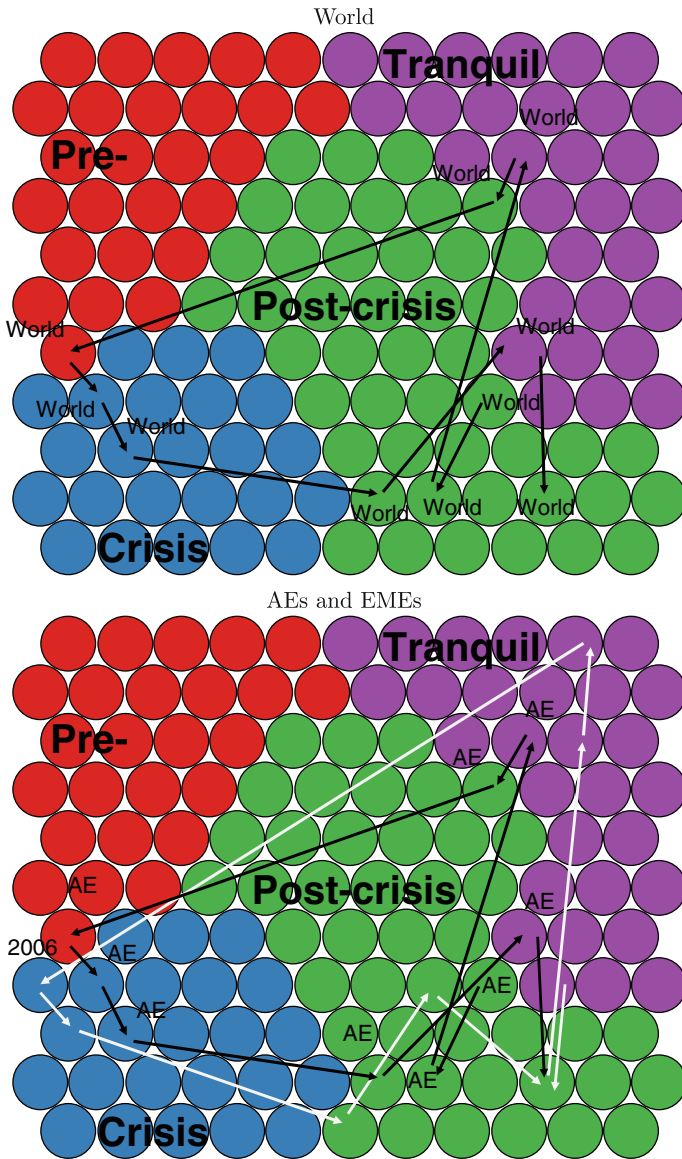


Fig. 8.6 A mapping of aggregates. *Notes* The figures displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Fig. 7.2). The *lines* that separate the map into four parts are based on the distribution of the four underlying financial stability states. Data points are mapped onto the grid by projecting them to their BMU using only macro-financial indicators. Consecutive time-series data are linked with *arrows*. On the first figure, the data for the aggregated world economy represent the first quarters of 2002–2011. On the second figure, the data for both advanced economies (AEs) and emerging market economies (EMEs) represent the first quarters of 2002–2011

advanced economies, illustrating the global dimension of the current crisis. While the emerging market cycle moves around that of the advanced economies, it does not indicate significant differences in the timeline or strength of financial stress.

8.3 Fuzzification of the SOFSM

Judging the degree of membership in a cluster on the SOFSM is not an entirely straightforward task. This section interchanges the current clustering of the SOFSM by fuzzifying and classifying it with a distance-based metric. As we not only have class information, but also utilize a semi-supervised SOM with the classes in the ordering process, there is no need to estimate clusters and their centroids. Following Sarlin (2013) (and Sect. 6.2.1), membership degrees are computed using inverse Euclidean distances by only using the class vector $x_{j(cl)} \in \mathbb{R}^4$. The rationale for this is the focus on distances between mean profiles of classes rather than those between indicators. The SOFSM is fuzzified by computing the inverse distance between reference vector $m_{i(cl)}$ and each perfect representative state center $c_{k(cl)}$ (as in Eq. 6.2), and normalized to fulfill the probabilistic constraint (as in Eq. 6.3). Yet, a defuzzification of the results using the maximum-membership method provides a crisp clustering. This enhances the visualization capability by enabling assessment of temporal belongingness to the financial stability states, where the states are expressed by representative cluster centers and fluctuations in macro-financial conditions are represented by the temporal variation of belongingness. In addition, visualizations of the memberships of units on a SOM grid enable assessment of the class structures. In the macroprudential oversight process, this supports risk identification, not the least the communication of individual data on the SOFSM.

Thereby, class information is accounted for by setting the number of states equal to the number of classes, i.e., four, and their centers as perfect states of the financial stability cycle: pre-crisis, crisis, post-crisis and tranquil states. To test different specifications, the fuzzifier is varied over a wide range ($\theta = 1.0, 1.2, \dots, 5.0$). Finally, squared Euclidean distances ($\theta = 2$) are chosen since that allows for overlapping, yet neither entirely crisp nor erased, state borders. While the differences between the most extreme choices of θ are significant, the results are stable for values close to $\theta = 2$. In Fig. 8.7, memberships to each state are shown on membership planes, where also the defuzzified crisp states are shown by contour lines (as well as on all other following grids, e.g., Fig. 8.7). The crispest part is the upper right corner of the tranquil state, whereas the rest have more overlap. The location of the cluster center (i.e., units closest to the perfect representative state center $c_{k(cl)}$), as shown by white X-marks, also depict the location of largest memberships (Fig. 8.8). For instance, in the pre-crisis cluster, units closest to the crisis cluster have the largest memberships. Figure 8.9 shows how the fuzzification can be turned into line graphs for the trajectories in Fig. 8.7, where membership in states and their variation over time represent fluctuation in the current state of financial stability, and vertical lines represent occurred crises. In addition to the crude trajectories on the SOFSM, this

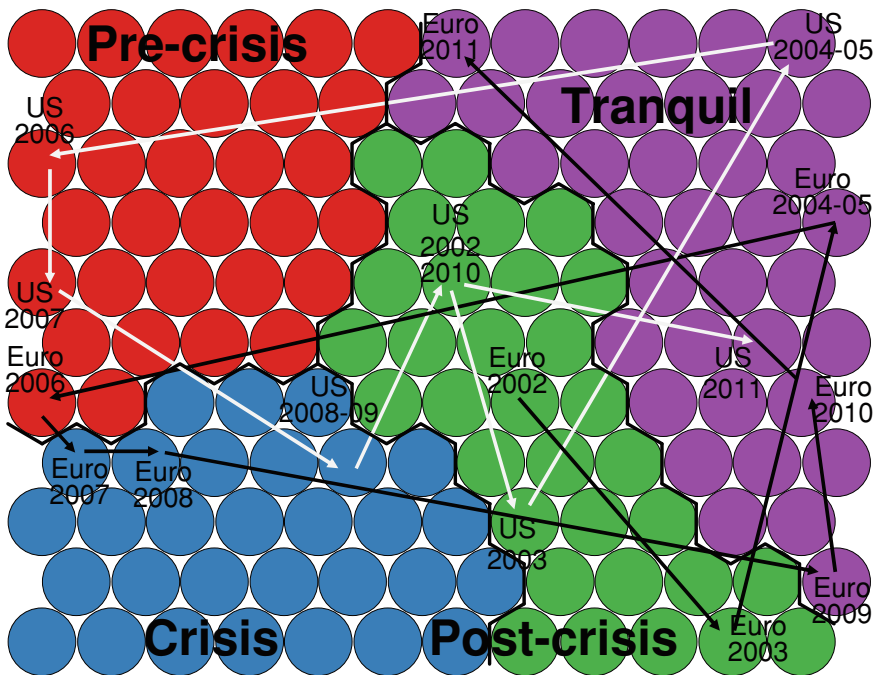


Fig. 8.7 The US and the euro area on a fuzzified SOFSM. *Notes* The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Fig. 7.2), but differs with respect to the partitioning of the map. A defuzzification of the financial stability states derives the lines that separate the map into four clusters. Data points are mapped onto the grid by projecting them to their BMUs using only macro-financial indicators. Consecutive time-series data are linked with *arrows*. The data for both US and the euro area represent the first quarters of 2002–2011

enables one to assess how the degree of membership in the financial stability states vary over time. The line graphs clearly depict increases in membership degrees in the pre-crisis states prior to crises. The figures depict, for instance, that the pre-crisis memberships in the US were of a larger magnitude than in the euro area.

8.4 Transitions on the SOFSM

Probabilities of transition provide means to support the judgment of the temporal structure on the SOFSM in general and the cyclical nature of the financial stability cycle in particular. The temporal patterns are approached by the means of computing, summarizing and visualizing probabilities of future state transitions. From the viewpoint of macroprudential oversight, this enables not only (risk) identification of the




Fig. 8.8 Membership planes for the SOFSM. *Notes* The figure represents memberships to the financial stability states, where *white X*-marks show the location of perfect representative state centers. The *contour lines* that separate the map into four clusters are derived using a defuzzification of the memberships

most likely future state transitions for each unit, and thus also pairing to individual data, but also country profiling of low- and high-risk financial stability states.

As shown in Fig. 8.10, the transition probabilities are computed for unit-to-state switches and visualized on own transition planes, and summarized as maximum transition probabilities conditional on switching, where labels show location and color probability. In this work, a wide range of time spans were tested ($s = 6, 12, 18, 24, 48$). Yet, for analysis was used a time span of 18 months ($s = 18$) that corresponds to that of the benchmark forecast horizon. Hence, transition probabilities represent the likelihood of switching to a state within 18 months. The rationale behind choosing $s = 18$ is that the SOFSM is also calibrated for optimal performance in terms of predicting vulnerable states 18 months prior to a crisis. Moreover, the transition patterns are considerably robust to changes in s . The length of movements increase with increases in s , as expected, while the directions of movements are stable. Most notably, while the transition patterns validate the assumed financial stability cycle, the cycle is shown not to be entirely well-behaving or continuous. For instance, the SOFSM shows high probability of transition to the crisis state on the border between the tranquil and pre-crisis states, as well as during extreme tranquil times. One can perform a similar line graph representation as that for the fuzzification, but instead with indications of future states, where vertical lines again represent occurred crises. Figure 8.11, while depicting probabilities according to the financial stability cycle, illustrates the crisis indications in the extreme part of the tranquil state and a rise of new instabilities in 2010.

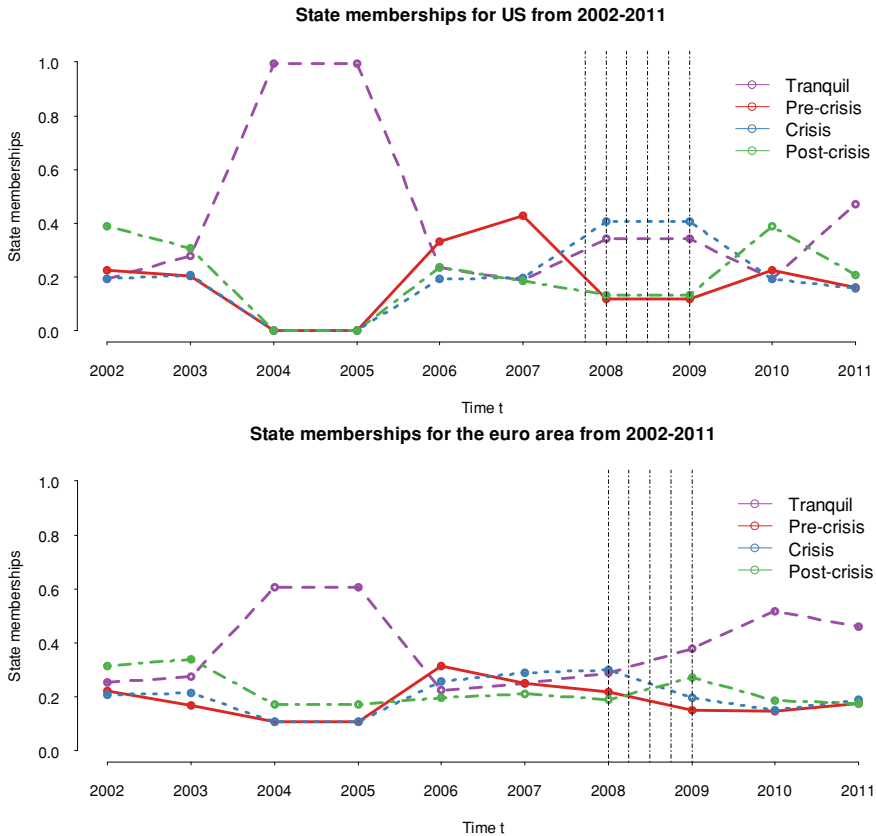


Fig. 8.9 Line graphs of US and euro area membership degrees. *Notes* The figure represents memberships of the US and the euro area trajectories in the financial stability states (see trajectory in Fig. 8.7). The vertical lines represent occurred crises

8.5 Scenario Analysis on the SOFSM

This section applies the SOFSM to scenario analysis. While the approach herein is an extremely simple version of what-if or scenario analysis, as it excludes all ingredients of more advanced macro stress-testing, the focus is on illustrating how the SOFSM suits for visualizing potential scenarios. We have previously used the SOFSM as a low-dimensional display onto which we have projected realized observations of macro-financial conditions, i.e., history. Scenario analysis differs only in the sense that the projected data are various scenarios of future conditions rather than historical patterns. However, more advanced macro stress-testing approaches should be used to derive effects of various scenarios. This relates to the task of risk assessment, in addition to the simultaneous means for risk identification.

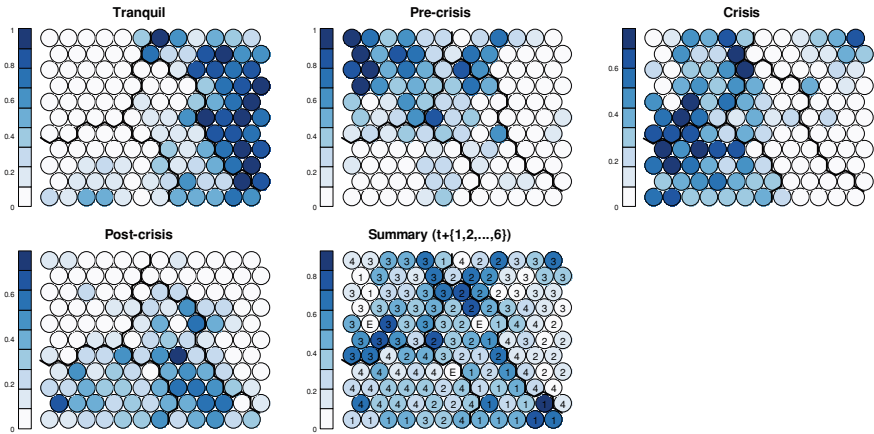


Fig. 8.10 Transition planes for the SOFSM. *Notes* The figure represents transition planes and summarized maximum-state transitions for the SOFSM. Each transition plane has its own color scale and the labels on the final grid correspond to the location of maximum-state transitions conditional on switching (where *E* represents empty units) and the *color* is the corresponding probability

Figure 8.12 presents transitions of the euro area given five different scenarios. In order to facilitate the visual representation of the scenarios, cluster memberships are illustrated with texture (i.e., a Bertin’s selective variable) and the scenarios in hue (i.e., a Bertin’s associative variable). The three types of introduced shocks are as follows.

- (1) Univariate shocks: ± 20 percentile variation in any variable.
- (2) Internal shocks: (2a) a positive (+30 percentiles) and (2b) negative (−30 percentiles) shock to domestic variables.
- (3) External shocks: (3a) a positive (+30 percentiles) and (3b) negative (−30 percentiles) shock to global variables.

The internal shocks represent changes in domestic macro-financial conditions that involve changes in real GDP, credit and equity growth, as well as leverage and equity valuation. Likewise, the external shocks involve the same changes in macro-financial conditions, but on a global level. The aim of positive and negative shocks is to represent increases and decreases in boom-like conditions.

The results of the scenario analysis in Fig. 8.12 are as follows. First, the introduced univariate ± 20 percentile variation in any variable shows that the euro area is not substantially sensitive to minor changes. The black arrows illustrate only a one-unit transition. Second, the introduction of positive and negative internal shocks in Fig. 8.12 shows different behavior with *solid* green and red arrows. A positive shock, involving booms in macro-financial conditions, would move the euro area to a substantially more vulnerable position. On the contrary, a negative internal shock would only involve a one-unit transition towards a less vulnerable state. Third, Fig. 8.12 also tests the resilience of the euro area conditions to external positive and negative shocks




Fig. 8.11 Line graphs of US and euro area transition probabilities. *Notes* The figure represents transition probabilities of the US and the euro area trajectories to the financial stability states (see trajectory in Fig. 8.7). The vertical lines represent occurred crises

(shown with *dashed green and red arrows*), involving global increases and decreases in macro-financial vulnerabilities. Opposite to internal shocks, the results show that the euro area is more sensitive to negative external shocks than positive ones. That is, a positive shock to macro-financial conditions only illustrates a two-unit transition towards a more vulnerable state, whereas a negative shock, or decrease in boom-like conditions, would involve a substantial transition towards a less vulnerable state. In a policy context, this could be related to the extent that one should be concerned about worsened macro-financial conditions globally, or some other specified shock.




Fig. 8.12 Scenario analysis on the SOFSM. *Notes* The figure displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Fig. 7.2). The lines that separate the map into four parts are based on the distribution of the four underlying financial stability states. The euro area is mapped onto the grid by projecting it to its BMU using only the macro-financial indicators. The three types of introduced shocks are as follows: (i) univariate shocks: ±20 percentile variation in any variable (black solid line); (ii) internal shocks: (2a) a positive (+30 percentiles, green solid line) and (2b) negative (−30 percentiles, red solid line) shock to domestic variables; and (iii) external shocks: (3a) a positive (+30 percentiles, green dashed line) and (3b) negative (−30 percentiles, red dashed line) shock to global variables

8.6 Shock propagation on the SOFSM

Transmission of financial shocks is often defined by a wide variety of measures, such as financial or trade linkages, proxies of financial shock propagation, equity market co-movement or geographical relations [see, e.g., Dornbusch et al. (2000), Pericoli and Sbracia (2003)]. Transmission of shocks on the SOFSM can be assessed with two methods: a superimposed portfolio network topology and neighborhoods on the SOFSM. This enables analyzing the spread of financial instabilities from two points of views: the portfolio network topology indicates propagation of financial stress through asset-based real linkages, while the financial stability topology indicates propagation to similar macro-financial conditions. The latter type of spread of events could propagate through both real (e.g., common exposures) and information (e.g., similar risks as judged by the markets) channels.

The SOFSM grid in Fig. 8.13 superimposes a network of financial links in 2010Q1 with the US as its center and a network in 2011Q2 with the euro area as its center. The networks are based upon external assets (equities and bonds) as reported in the

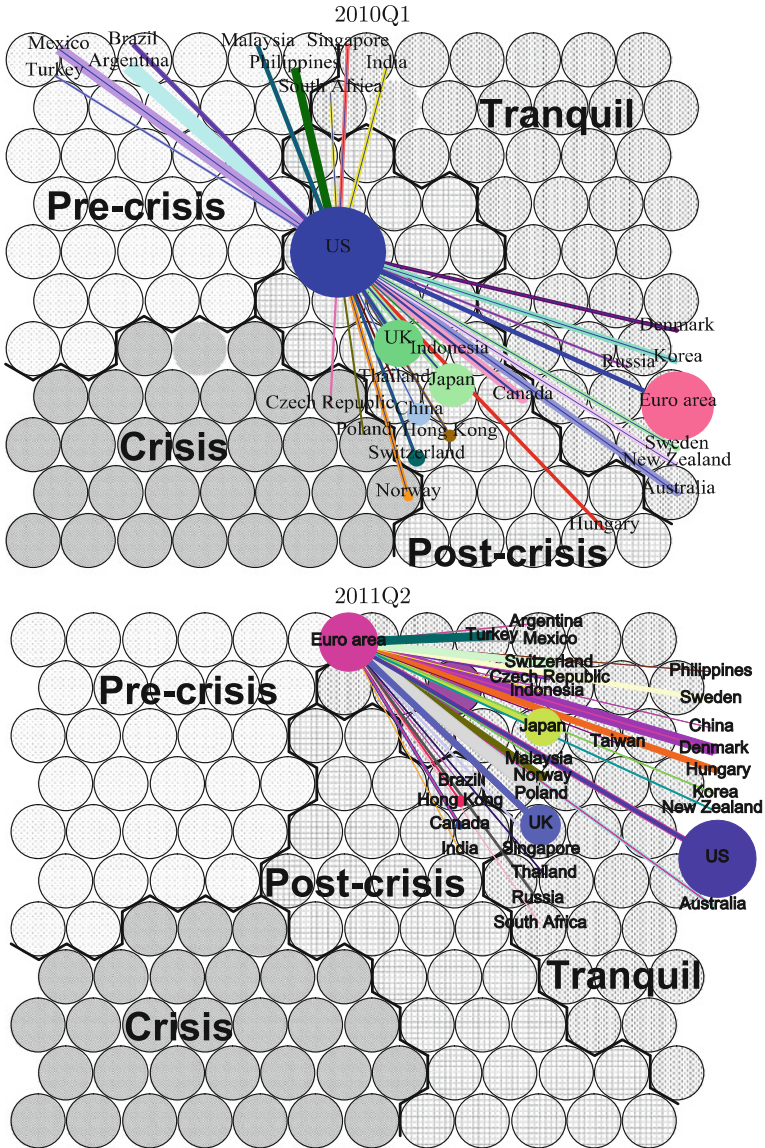


Fig. 8.13 A financial network topology on the SOFSM. *Notes* The figures displays the two-dimensional SOFSM that represents a high-dimensional financial stability space (same as in Fig. 7.2). The financial stability states are differentiated with *texture*, rather than *color*. The figure superimposes a financial network on the SOFSM, of which the US and the euro area are in the *center*, respectively. The network of financial linkages is based upon external assets (equities and bonds). Nodes of each economy are located in their BMUs m_b , but positions are approximate to fit all nodes and edges. The size of the nodes is scaled as to the sum of exposures to other economies. The width of the edges represents the size of external exposure to total exposures per economy, where the color of the edge indicates the address of the exposure holder. The data for all economies represent the first quarter of 2010 and the second quarter of 2011

Coordinated Portfolio Investment Survey by the International Monetary Fund (IMF). Nodes of each economy are located in their BMUs and their size is scaled as to the sum of exposures to other economies. The thickness of the edges represents the size of external exposure to total exposures per economy, where the color of the edge indicates the address of the exposure holder. Indeed, Fig. 8.13 combines the state of financial stability, or probability of a crisis, with the system-wide exposures of each economy. In 2010, the size of financial linkages to high-risk economies (e.g., Brazil and Mexico) enlighten about both past and present: high levels of previous financial stress in the US may have impacted their current state and they still have a high risk of current and future shock propagation from the US. Likewise, the strong connections of economies to the euro area is an indication of transmission channels in the case of distress in the euro area (e.g., Denmark, Poland, Turkey and United Kingdom UK).

While crises are often transmitted through asset-based contagion channels, such as financial linkages, they may also be propagated through similarities in macro-financial conditions, something particularly important when dealing with data of changing nature. When assessing the SOFSM, the concept of neighborhood of a country represents the similarity of the current macro-financial conditions. Hence, independent of location on the map, an economy adjacent to countries in crisis could through shock propagation experience a similar wave of financial distress. This type of representation may help in identifying events surpassing historical experience and the changing nature of crises. Thus, an economy in the upper left part of the network for 2010 in Fig. 8.13, say Mexico, could propagate financial instabilities to countries with similar macro-financial vulnerabilities, e.g., Argentina and Brazil. While this is particularly useful for visual real-time surveillance, we can also test this by letting locations of crises in period t be signals of crises in that location in period $t + s$. More precisely, this creates a leading indicator that signals a crisis in unit m_b in period $t + s$ if a country that experienced a crisis in t was located in m_b , where $s = 6, 12, 18, 24, 48$. As the indicator is a point forecast P_j , it needs no transformation through threshold values. Table 8.2 shows the predictive performance of neighborhoods on the SOFSM with forecast horizons of 6–48 months, where a horizon of 24 months outperforms the rest. A policymaker with $\mu = 0.8$ derives the largest Usefulness. While the table confirms the usefulness of detecting the spread of crisis, the nature of the shock-propagation measures suggest that they are rather complements than substitutes to standard early-warning models.

8.7 Outlier Analysis with the SOFSM

The SOM paradigm provides a simple measure of extremity. Relating to the general task of risk identification, one may assess whether or not, and to what extent, distances between each datum and its mean profile on the SOFSM (i.e., BMU) is an indication of financial imbalances. The computational rationale for this is that the units of the SOM, while being topologically ordered, tend to approximate the probability density

Table 8.2 Predictive performance of spillover and outliers on the SOFSM

Model	Horizon	Threshold	RP	RN	PP	PN	Accuracy	U_r ($\mu = 0.7$)	U_r ($\mu = 0.8$)	U_r ($\mu = 0.9$)	AUC
Contagion	6	BMU	0.23	0.92	0.21	0.93	0.86	0.30	0.27	0.03	-
Contagion	12	BMU	0.22	0.91	0.32	0.86	0.80	0.30	0.26	0.01	-
Contagion	18	BMU	0.19	0.89	0.35	0.78	0.73	0.38	0.35	0.12	-
Contagion	24	BMU	0.17	0.88	0.39	0.71	0.67	0.56	0.57	0.45	-
Contagion	48	BMU	0.13	0.85	0.47	0.48	0.48	0.42	0.36	0.13	-
QE (d_j)	6	0.65	0.49	0.66	0.12	0.93	0.65	0.19	0.02	-0.38	0.56
QE (d_j)	12	0.66	0.48	0.69	0.23	0.87	0.65	0.17	0.00	-0.40	0.55
QE (d_j)	18	0.66	0.54	0.71	0.30	0.86	0.67	0.13	-0.03	-0.44	0.60
QE (d_j)	24	0.63	0.72	0.71	0.35	0.92	0.71	0.11	-0.06	-0.49	0.73
QE (d_j)	48	0.71	0.53	0.79	0.44	0.84	0.73	0.05	-0.12	-0.56	0.69

Notes: The table reports results for contagion and quantization errors (QEs) on the SOFSM. QE is a percentile transformation of d_j and contagion neighborhood is the BMUs of previous pre-crisis periods. The following measures are reported: TP true positives, FP false positives, TN true negatives, FN false negatives, Precision positives (PP) = $TP/(TP + FP)$, Recall positives (RP) = $TP/(TP + FN)$, Precision negatives (PN) = $TN/(TN + FN)$, Recall negatives (RN) = $TN/(TN + FP)$, Accuracy = $(TP + TN)/(TP + TN + FP + FN)$, AUC area under the Receiver Operating Characteristic curve (not computed for contagion as it is binary), and Usefulness $U_a(\mu) = \text{Min}(\mu, 1 - \mu) - (\mu(FN)/(FN + TP)) + (1 - \mu)(FP/(FP + TN))$, where μ stands for cost of FP and FN. For QE, the threshold is chosen as for optimal Usefulness. Best values per measure and method are bolded

function of data, which relates the distance to a BMU to a fit of a single datum to the multivariate data distribution, i.e. its degree of extremity. On the other hand, the economic rationale, when monitoring financial stability, is that one could assume that large distances represent financial imbalances in macro-financial conditions. Our approach goes beyond applying a predefined threshold value on the distance to assess whether or not a datum is an outlier [see, e.g., Vesanto et al. (1998), Saunders and Gero (2001)], by setting the threshold to optimize predictive performance. More formally, the units m_i of the SOM, while being topologically ordered, tend to approximate the probability density function of data $p(x)$ (Kohonen 2001). The standard quantization error (QE) can be seen as the correspondence between m_i and x_j . However, a more meaningful estimate of event rarity is computing the distance of individual data points x_j to their BMU m_b . An outlier, and its degree of extremity, can thus be estimated by the distance to the SOM in a multidimensional setting, i.e., $d_j = \|x_j - m_{b(j)}\|$. To be precise, given that m_i approximate the probability density functions, then the individual QE represents in a temporal setting the fit of a single data point to the historical multivariate data distribution. Finally, the distance d_j is turned into a probability forecast p_j through a percentile transformation, on which a threshold $\lambda \in [0, 1]$ is chosen to optimize Usefulness $U_a(\mu)$.

I assess whether or not, and to what extent, distances between each data vector and its BMUs is an indication of financial imbalances. Table 8.2 shows the predictive performance of QEs on the SOFSM with forecast horizons of 6–48 months. The table illustrates that, while the aim is conceptually different, outliers do not provide equally good means to predict financial instabilities as contagion does. The predictive capability improves with shorter forecast horizons and yields the largest Usefulness for a horizon of 6 months. The measure is most useful for a policymaker with $\mu = 0.7$. The weak performance may reflect the fact that it provides only information of possible impending instabilities rather than information on the exact timing of a crisis, as the imbalance may be located in any state of the financial stability cycle. This is, however, an important property as this does not restrict modeling to the precise nature of crises in the past.

8.8 Combining the SOFSM with Predictive Methods

The aim of this section is to illustrate how the SOFSM can, and should, be complemented by other tools for risk identification and assessment alike. In particular, this section combines the SOFSM with a model for predicting systemic financial crises with the aim of risk identification. The approach herein follows that in Sarlin (2014a) by applying a standard Genetic Algorithm (GA) for finding the optimal configuration of an artificial neural network (ANN)—and coin it the neuro-genetic (NG) model. The rationale for this is to test whether, and to what extent, ANN-based models are better than the SOFSM and logit models and the effect of automated calibration of the NG model.

Even very simple ANNs have been shown to be universal approximators by following any continuous function to any desired accuracy (Hornik et al. 1989). This said, the focus of data-driven ANN applications in real-world settings with noise and uncertainty (e.g., financial markets) should rather be on parsimony and generalization than on fitting models to all non-linearities and complexities in data. Another common concern is the extent of data dredging when conducting data-driven analysis. To this end, an ANN-based early-warning model is built using two objective training, or early stopping, schemes:

- (i) *Scheme 1*: Training is performed until in-sample performance of a conventional benchmark model has been reached.
- (ii) *Scheme 2*: In-sample data are divided into two datasets: train and validation sets. Models are trained on the train set and the one with optimal performance on the validation set is chosen.

As in Sect. 7.4, the in-sample dataset is used for estimating a logit model. The estimates of the model are then used to solicit the probability of a crisis and the threshold chosen as to maximize Usefulness for policy action. In *training scheme 1*, the in-sample Usefulness of the logit model for policy action, $U_r(\mu) = 0.5$, is used as a stopping criterion when training the ANNs. The rationale behind this is twofold: it attempts to prevent overfitting and enables testing whether an ANN that is equally good on the in-sample performs better on out-of-sample data. The performance of ANN configurations is tested over a wide set of possibilities as well as of the automated NG model. In *training scheme 2*, the in-sample dataset is randomly split as follows: 80 % train set and 20 % validation set. This gives us three datasets: train (in-sample, 80 %), validation (in-sample, 20 %) and test (out-of-sample) sets. Then, ANN and NG models are trained by optimizing Usefulness for policymakers on the validation set. In practice, models are trained for 200 epochs, evaluate them at each epoch and choose the one that maximizes Usefulness on the validation set. This allows testing how much better, if at all, the ANN-based models perform when attempting an optimal model. As ANNs are sensitive to initial conditions of the weights, the training of ANN and NG models is repeated ten times with randomized starting weights and biases, and then the one with the fastest convergence (least epochs) is chosen.

As for the SOFSM, a benchmark policymaker is assumed to be substantially more concerned about missing crises than issuing false alarms ($\mu = 0.8$), whereas model performance is also shown for a slight variation in preferences ($\mu = 0.7, 0.9$). For all models using training scheme 1 (in-sample $U_r(\mu = 0.8) = 0.5$), the final ANN elements and GA parameters, as well as their out-of-sample Usefulness for policymakers, are shown in Table 8.3. While models ANN1–9 represent manual configurations, the model A-ANN represents an average of all different manually chosen ANN configurations. When manually parametrizing the ANNs, common practices and rules of thumb in the literature have been followed, in addition to testing variations to the most common choices. The results clearly depict differences in model performance (best models per μ are **bolded**). The ANN-based models, while having similar in-sample performance by definition, show consistently better

Table 8.3 Specifications of the neural network and neuro-genetic models

	<i>Logit</i>	<i>ANN1</i>	<i>ANN2</i>	<i>ANN3</i>	<i>ANN4</i>	<i>ANN5</i>	<i>ANN6</i>	<i>ANN7</i>	<i>ANN8</i>	<i>ANN9</i>	<i>A-ANN</i>	<i>NG</i>
<i>NN configuration</i>												
No. of input indicators	-	14	14	14	14	14	14	14	14	14	14	9
No. of nodes in the output layer:	-	1	1	1	1	1	1	1	1	1	1	1
Number of hidden layers	-	1	1	1	2	2	1	1	1	2	1	1
No. of nodes in layer 1	-	2	2	8	2	8	2	2	2	20	5	4
No. of nodes in layer 2	-	0	0	0	2	8	0	0	0	20	0	0
Activation function	-	Sigmoid	Sigmoid	Sigmoid	Sigmoid	Sigmoid	Gaussian	Hyperbolic	Linear	Sigmoid	Sigmoid	Sigmoid
No. of training epochs	-	15	11	9	22	11	21	13	12	5	19	6
Learning rate α	-	0.9	0.8	0.95	0.9	0.9	0.9	0.9	0.9	0.9	0.83	0.97
Momentum m	-	0.9	0.2	0.8	0.9	0.9	0.9	0.9	0.9	0.9	0.63	0.15
Input noise ε	-	0.01	0	0	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01
<i>GA configuration</i>												
Generation count	-	-	-	-	-	-	-	-	-	-	-	10
Population size Ω	-	-	-	-	-	-	-	-	-	-	-	3
No. of crossovers	-	-	-	-	-	-	-	-	-	-	-	1
Mutation rate	-	-	-	-	-	-	-	-	-	-	-	0.10
Fitness criterion	-	-	-	-	-	-	-	-	-	-	-	Train error
<i>Out-of-sample usefulness U_r</i>												
$\mu = 0.7$	0.08	0.16	0.12	0.14	0.12	0.16	0.12	0.14	0.16	0.16	0.14	0.20
$\mu = 0.8$	0.27	0.36	0.30	0.32	0.32	0.32	0.28	0.28	0.34	0.32	0.32	0.38
$\mu = 0.9$	0.13	0.16	0.08	0.10	0.12	0.08	0.04	0.02	0.14	0.10	0.09	0.16

Notes: A-ANN refers to the mean or mode, as applicable, of all ANN configurations and their performances. The A-ANN configuration is rounded so that each entry is meaningful (e.g., the second layer has no nodes as on average there is only one layer). The bold evaluation entries represent the best-performing model per μ -value

Table 8.4 Predictive performance of the neuro-genetic model

Model	Epochs	In-sample			Validation set			Out-of-sample		
		$\mu = 0.7$	$\mu = 0.8$	$\mu = 0.9$	$\mu = 0.7$	$\mu = 0.8$	$\mu = 0.9$	$\mu = 0.7$	$\mu = 0.8$	$\mu = 0.9$
Logit	–	0.39	0.50	0.30	–	–	–	0.08	0.27	0.13
ANN1	36	0.38	0.58	0.38	0.18	0.38	0.20	0.18	0.38	0.18
NG	17	0.40	0.60	0.42	0.26	0.44	0.24	0.26	0.44	0.22

Notes Except for the epochs, which stand for the number of training iterations, the entries represent $U_r(\mu)$ for different policymaker's preferences μ

out-of-sample performance than the benchmark logit model. Table 8.3 shows that ANNs outperform the logit model not only in specific cases, but also on average (A-ANN). The best ANN model (ANN1) follows parametrization practices common in the early-warning literature, in particular by having 2 hidden nodes and a learning rate $\alpha = 0.9$ [see, e.g., Peltonen (2006), Fioramanti (2008)]. However, best overall performance is shown by the NG model. Most notably, the optimal GA configuration for the ANN uses only 9 indicators, rather than all 14. From an economic point of view, dropping credit growth and current account deficit, as well as global real GDP growth, real credit growth and leverage (see Table 7.1), contradicts the results of a recent study based upon the signaling approach (Alessi and Detken 2011). In contrast to the present analysis, they do not, however, attempt to identify optimal indicators in a multivariate framework, but rather conduct it in a univariate manner.

Further, key parameters describing the NG model are a learning rate $\alpha = 0.97$ and 4 hidden nodes. When examining differences in performance for different policymakers' preferences, one can observe that the logit model fails for those more averse to giving false alarms ($U_r(\mu) = 0.08$), the ANN results are somewhat mixed, with ANN1 performing particularly well for a policymaker more concerned with false alarms ($U_r(\mu) = 0.16$), and the NG model yields Usefulness for all three types of policymakers ($U_r(\mu) = 0.20, 0.38, 0.16$ for $\mu = 0.7, 0.8, 0.9$). For the benchmark preferences $\mu = 0.8$, the NG model performs 11, 2 and 6 % points better than the logit, ANN1 and A-ANN models, respectively. While ANNs are heuristic in nature, the consistency in the slight superiority is likely to be a result of the highly parsimonious training scheme.

Training scheme 2 attempts a better generalization by being less restrictive in terms of parsimony but still attempting to prevent overfitting. Table 8.4 summarizes in-sample (includes both train and validation sets), validation and out-of-sample performance of the logit model, the best-performing ANN model (ANN1) and the NG model, all using training scheme 2. The table shows that, in principle, when allowing for longer training and thus also a better fit to data, model performance improves on all datasets (best performing model per μ and dataset is **bolded**). This is obvious when it comes to in-sample data, but the validation as well as out-of-sample data still need sufficient parsimony for decent performance. Compared to training scheme 1, while there is only a minor increase in out-of-sample Usefulness of ANN1 (2 % points), the NG model experiences a performance increase of 6 % points. Out-of-sample Usefulness of the NG model is 17 % points better than that of




Fig. 8.14 Probabilities of a financial crisis in the US. *Notes* The vertical lines represent the occurrence of crisis (black) and pre-crisis (red) periods. The area highlighted in gray represents periods when the probability of a crisis has exceeded a threshold such that an early-warning signal is given

the logit model and 6 points better than that of the ANN1 model. This also depicts superior performance of the NG model.

Generally, outputs of early-warning models are time-series of country-specific crisis probabilities and can be visualized as line graphs. The line graphs in Figs. 8.14 and 8.15 shows the probability of a crisis in the US and the euro area within 18 months as an output of the NG model. The vertical lines represent occurred events, where red lines are pre-crisis periods and black crisis periods. Figure 8.14 illustrates that the model correctly called at an early stage the dot-com bubble in 2001 and the recent financial crisis in 2007–2008—both with a longer horizon than 18 months as was the definition of the predicted variable. Thus, this model would already in 2006Q1 have signaled the global financial crisis that commenced in the US in 2007. The prediction in Fig. 8.15 shows model performance for the euro area. The early crises stemming from the Russian collapse in 1998 and the dot-com bubble in 2001 are both correctly called. In fact, the recent financial crisis was already signaled in 2004Q2 (even when accounting for publication lags). Yet, it is worth noting that while model performance in these two cases is appropriate, many of the early-warning signals are actually given before the ideal leading indicator or during a crisis period, which both in fact are false alarms. These types of errors, while not having large adverse effects in terms of policy actions, lead to imperfect accuracies.

A partly valid limitation of these conclusions is, however, that the models were built *ex post*, and hence the design of the early-warning model might have benefited from hindsight bias. However, the inputs used in this study were commonly used in the macroprudential literature already before the crisis [see for instance the work




Fig. 8.15 Probabilities of a financial crisis in the euro area. *Notes* The vertical lines represent the occurrence of crisis (black) and pre-crisis (red) periods. The area highlighted in gray represents periods when the probability of a crisis has exceeded a threshold such that an early-warning signal is given

by Borio and Lowe (2002, 2004)] and the definition of a crisis, while still being somewhat subjective, is validated by being highly correlated with financial crises over the entire sample period. Hence, a policymaker could, in principle, have used similar specifications to derive a model for predicting crises at the turn of the century. Interestingly, the real *ex ante* predictions, i.e., the final points in Figs. 8.14 and 8.15, show that vulnerabilities in the US and the euro area have increased in 2011. In fact, it signals for the second consecutive quarter a crisis in Europe within 18 months. As this model was derived in the final quarter of 2011 for Sarlin (2014a), we may based upon today's experiences judge whether or not the prediction was correct.

8.9 Concluding summary

This chapter has utilized the SOFSM for tasks of interest in macroprudential oversight, particularly for risk identification and assessment, with the ultimate aim of risk communication. Risk identification relates to early-warning models, whereas risk assessment relates to macro stress-test models and contagion and spillover models. Yet, in relation to previous literature, the approaches herein stress risk communication by focusing on performing the tasks on a visual two-dimensional display of a high-dimensional financial stability space. In particular, the SOFSM has been

exploited by the means of the following eight approaches (where the *parenthesis* represents the addressed systemic risk):

- (i) Assessing distributions of the macro-financial indicators and all class variables with the help of the feature planes of the SOFSM (*imbalances*).
- (ii) Mapping the state of financial stability for individual data and aggregates by the means of labels and trajectories on the SOFSM (*imbalances*).
- (iii) Fuzzification of the SOFSM for visualizing temporal belongingness to financial stability states of individual data and class distance structures on the map (*imbalances*).
- (iv) Probabilistic modeling of state transitions on the SOFSM for visualizing probabilities of transition to financial stability states of individual data and for assessing the cyclical and temporal structure of the financial stability cycle (*imbalances*).
- (v) Scenario analysis for economies on the SOFSM by assessing the effects of positive and negative shocks, both domestically and globally (*aggregate shocks*).
- (vi) Using superimposed portfolio network topologies and neighborhoods on the SOFSM to assess the spread of financial distress and shock propagation (*contagion and spillover*).
- (vii) Computing distances between data and their mean profiles on the SOFSM to find extreme events and imbalances in economies' macro-financial conditions (*imbalances*).
- (viii) Complementing the SOFSM with a solely predictive model that uses genetically optimized neural networks for the identification of risks (*imbalances*).

One task, obviously among many other tasks of importance, that this chapter has overlooked is the identification of the build-up of widespread imbalances in the entire cross section. This is the focus of the SOTM in the following chapter.

References

- Alessi, L., & Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: a role for global liquidity. *European Journal of Political Economy*, 27(3), 520–533.
- Borio, C., & Lowe, P., (2002). *Asset prices, financial and monetary stability: Exploring the nexus*. BIS Working Papers No. 114.
- Borio, C., & Lowe, P., (2004). *Securing sustainable price stability: Should credit come back from the wilderness?* BIS Working Papers No. 157.
- Card, S., Mackinlay, J., & Schneidermann, B. (1999). *Readings in information visualization, using vision to think*. San Diego, CA: . Academic Press Inc.
- Dornbusch, R., Park, Y., & Claessens, S. (2000). Contagion: how it spreads and how it can be stopped. *World Bank Research Observer*, 15, 177–197.
- Fioramanti, M. (2008). Predicting sovereign debt crises using artificial neural networks: a comparative approach. *Journal of Financial Stability*, 4(2), 149–164.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2, 359–366.

- Keim, D., Mansmann, F., Schneidewind, J., & Ziegler, H. (2006). Challenges in visual data analysis. In *Proceedings of the IEEE International Conference on Information Visualization (iV 13)* (pp. 9–16). London, UK: IEEE Computer Society.
- Kohonen, T. (2001). *Self-organizing maps* (3rd ed.). Berlin: Springer.
- Peltonen, T., (2006). *Are emerging market currency crises predictable? A test*, ECB Working Paper No. 571.
- Pericoli, M., & Sbracia, M. (2003). A primer on financial contagion. *Journal of Economic Surveys*, 17, 571–608.
- Sarlin, P. (2013). Exploiting the self-organizing financial stability map. *Engineering Applications of Artificial Intelligence*, 26(5–6), 1532–1539.
- Sarlin, P., (2014a). On biologically inspired predictions of the global financial crisis. *Neural Computing & Applications*, 24(3–4), 663–673.
- Sarlin, P., & Peltonen, T. (2013). Mapping the state of financial stability. *Journal of International Financial Markets, Institutions & Money*, 26, 46–76.
- Saunders, R., & Gero, J. (2001). Designing for interest and novelty: Motivating design agents. In B. de Vries, J. van Leeuwen, & H. Achten (Eds.), *CAAD Futures* (pp. 725–738). Dordrecht: Kluwer.
- Thomas, J., & Cook, K., 2005. *Illuminating the path: Research and development agenda for visual analytics*. New York: IEEE Press.
- Vesanto, J., Himberg, J., Sisonen, M., & Simula, O. (1998). Enhancing som based data visualization. In *Proceedings of the International Conference on Soft Computing and Information/Intelligent Systems (IIZUKA 98)* (pp. 64–67). Iizuka, Japan.

Chapter 9

Decomposing Financial Crises with SOTMs

The provided models for macroprudential oversight have thus far concerned assessing the cross-sectional or temporal dimensions in close to isolation. In this chapter, we turn the focus to exploring cross-sectional dynamics. The Self-Organizing Time Map (SOTM) provides means for visual dynamic clustering and thus also for illustrating dynamics in cross sections of multivariate macro-financial indicators. This is one of the very key tasks in risk identification, when the focus is on build-up phases of imbalances in the entire cross section, such as the global dimension in country-level risks and a system-wide focus on data concerning individual financial intermediaries. With respect to the visual analytics mantra, the SOTM can be positioned similarly as the previously discussed Self-Organizing Financial Stability Map (SOFSM).

The SOTM performs visual dynamic clustering through temporal data and dimension reduction. The approach differs from traditional static exploratory analyses in that the SOTM dynamically adapts to structural changes in cross-sectional data over time, as well as visualizes the temporal cluster structures. In short, the decomposition is enabled by data compression into clusters and twofold topology preservation, where one direction preserves time and the other data topology. The first decomposition applies the standard SOTM to describing the global financial crisis that started in 2007 in a manner that would be applicable for real-time surveillance. The second section uses a SOTM on time-to-event data to generalize patterns before, during and after financial crises. The following two sections draw upon Sarlin (2013a, b), respectively.

This chapter is partly based upon previous research. Please see the following work for further information: Sarlin (2013a, b)

9.1 A Decomposition of the Global Financial Crisis

In this section, the SOTM is applied for decomposing the global financial crisis that started in 2007 in order to identify temporal structural changes and their location in the cross section. The use of the SOTM is illustrated with an abstraction of all the data $x_j \in \mathbb{R}^{18}$ before, during and after the global financial crisis of 2007–2009. We first discuss the parametrization of the SOTM, and then focus on univariate and multivariate properties of the SOTM, including a second-level clustering of the SOTM units.

9.1.1 Parametrizing the SOTM

Similarly as an unsupervised Self-Organizing Map(SOM), the SOTM may use parts of the data in training and only associate parts. The indicator vector $x_{j(in)} \in \mathbb{R}^{14}$ is used to train the SOTM and the spread of the class vectors $x_{j(cl)} \in \mathbb{R}^4$ is only associated to the model. The association is done by computing for each unit an average of class variables for the data attracted by that unit. The model architecture is set to 8×22 units, where 22 units represent the time dimension and 8 units represent the cross-sectional structures. The units on the time dimension are set as to span periods before, during and after the crisis that started in 2007 (i.e., 2005Q2–2010Q3), while the number of units at each point in time is determined based upon its descriptive value. It is worth noting that the SOTM, likewise the SOM, is not restricted to treat each unit as an individual cluster. Due to the property of approximating probability density functions $p(x, t)$, only the dense locations in the data tend to attract units. A further motivation of the number of units on the vertical axis to exceed the number of expected clusters is the second-level clustering of the SOTM.

When choosing the final specification of the SOTM, three quality measures introduced in Sect. 6.3.2 (ε_{qe} , ε_{te} and ε_{sc}) are used. For a SOTM with 8×22 units, a neighborhood radius $\sigma = 2.4$ is chosen, as it has the highest quantization accuracies and no topographic errors (see dashed vertical line in Fig. 9.1). Topographic error is stressed as the interpretation of a SOTM relies heavily on topology preservation, not least the time dimension.¹

9.1.2 A Univariate View of the Crisis

The output of the SOTM, while being a two-dimensional grid, is a set of multidimensional reference vectors. For a better understanding of the above trained SOTM, and its characteristics, we begin by an illustration of the feature planes for individual inputs (Fig.9.2). Feature planes are layers of the two-dimensional SOTM in Fig. 9.3

¹ Beyond the static representations herein, the implementation developed by infolytika provides an interactive, web-based interface to the SOTM (<http://risklab.fi/demo/macropu/fsmt/>). For a description, see Sarlin (2014a).

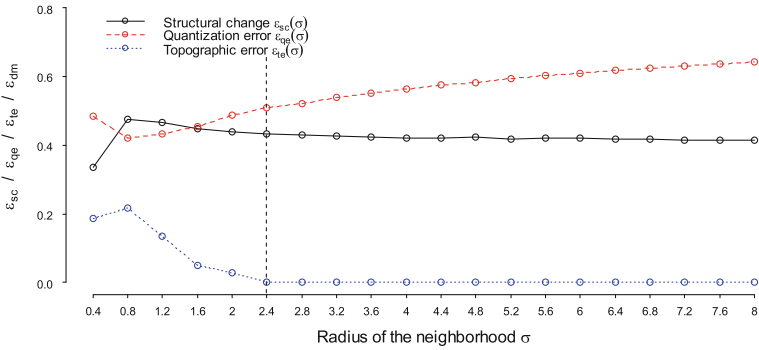


Fig. 9.1 Quality measures of the SOTM. *Notes* For models with a 8×22 array of units, the errors ϵ_{qe} and ϵ_{te} are computed as aggregates of all time units $t = 1, 2, \dots, T$ and structural changes ϵ_{sc} of time units $t = 2, 3, \dots, T$ over neighborhood radii σ

and show the spread of individual inputs using a constant blue hue and variations in luminance. Again, each feature plane has its individual scale on the left and a timeline below. With a focus on univariate structures, feature planes are particularly useful for monitoring the evolution of individual inputs on the SOTM, especially for discovering the spread of values in the cross-section and their variation over time. The last four feature planes represent the class variables $x_{j(cl)} \in \mathbb{R}^4$, while the rest represent the macro-financial indicators $x_{j(in)} \in \mathbb{R}^{14}$, including both domestic and global measures. The distribution of the class variables, in particular pre-crisis periods, illustrates that during the early pre-crisis periods vulnerable economies were mainly located in the lower part of the SOTM, whereas the crises occur throughout the cross-section in 2008–2009. The spread of the input variables also indicates larger vulnerabilities in the lower part of the SOTM. For instance, real credit growth, leverage and current account deficit generally take higher values in the lower part and government deficit lower values, which all can be seen as build-ups of risks, vulnerabilities and imbalances. Moreover, the feature planes of the input variables also illustrate a number of temporal changes. For instance, one can observe a loss in equity growth in 2008Q3 across the entire cross-section, an increase in 2010Q1–2 and somewhat decrease in 2010Q2–3. Losses in Losses in gross domestic product (GDP), while also occurring throughout the entire cross-section, react only in 2009Q2. Credit growth may be seen as an imbalance that decreased during the crisis period, and has not experienced any significant increases after the crisis. Government deficits are shown to have widely increased in the latter part of the analyzed period, as could be expected. This highlights the importance of the use of the SOTM for assessing events of changing nature, as government deficits were clearly not a signal of the first wave of distress. Yet, the deficits may obviously be related to increases in government debt, and thus also to the current sovereign debt problems. Interestingly, leverage is shown to increase during the sample period, but does not show significant decreases during or after the financial crisis, rather the opposite. Globally, we can

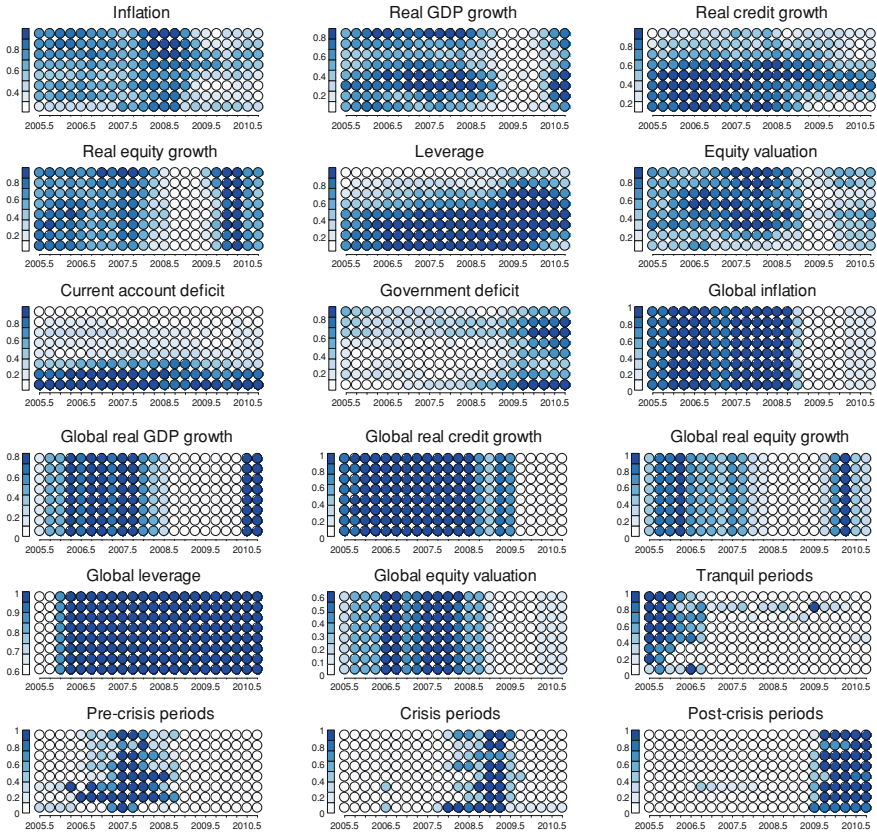


Fig. 9.2 Feature planes for the SOTM. *Notes* The figure shows feature planes for the 14 indicators and the benchmark class variables. The feature planes are layers of the SOTM in Fig. 9.3. While the indicators are defined in Table 7.1, the four main class variables are Pre-crisis (C18), Crisis (C0), Post-crisis (P18) and Tranquil periods (T0). As each data vector consists of 14 indicators and 4 main class variables, these feature planes show the distribution of each data column on the SOTM grid. In the case of binary class variables that take values 1 and 0, high values represent a high proportion of data in that unit (pre-crisis, crisis, post-crisis or tranquil periods)

see similar patterns to the domestic ones, but with a surprisingly strong slowdown in credit and equity growth, as well as an even higher level of leverage.

9.1.3 A Multivariate View of the Crisis

Now, as we have an understanding of the components of the SOTM, we can assess its multivariate structures. This is done with three approaches. First, Fig. 9.3 illustrates the two-dimensional SOTM, where the timeline below the figure represents the time




Fig. 9.3 A SOTM of the global financial crisis. *Notes* The figure represents a SOTM of the global financial crisis, where the cluster coloring shows changes in multivariate cluster structures. Labels above the figure define the classes in data, i.e., the stages of the financial stability cycle, and the trajectory on the SOTM represents the evolution of macro-financial conditions in the United States (US)

dimension in data (as for the feature planes) and the labels above represent occurrences of the events in the cross section. The labels simply refer to averages of the classes at each point in time. The coloring of the SOTM in Fig. 9.3 illustrates the proximity of units as approximated by the Sammon’s mapping in Fig. 9.4. Thus, in Figs. 9.3 and 9.4, differences between units along the vertical direction show differences in cross-sections and differences along the horizontal direction show differences over time, where the former figure illustrates differences with color coding on the SOTM grid and the latter with a standard plot of the Sammon’s topology. The trajectory on the SOTM represents the evolution of macro-financial conditions in the US, which clearly illustrates that the US were characterized by relatively large risks and vulnerabilities throughout the period. Figures 9.3 and 9.4 illustrate one key phenomenon. Shifts in the color scale towards yellow indicate a start of structural changes during the early phases of the crisis in 2008, whereas the structural changes reach their peak in 2009 and the structures move back in mid-2010. The interpretation of the backward shift in 2010 is, however, somewhat ambiguous. The shift, while being an indication of decreased financial stress, may also be an indication of future risks, as the structures clearly resemble those during the pre-crisis peak.

To further assess the evolution of the multivariate structures, a second approach applies a second-level clustering to the units of the trained SOTM [as proposed in Sarlin and Yao (2013)]. With no predefined number of groups and with the aim of examining the structures in these data, we can explore clustering solutions with different K . This is a common exercise with hierarchical clustering methods as the agglomeration process provides insights about the structures. Thus, the cluster validation is not used for choosing one optimal clustering solution, but rather to




Fig. 9.4 A Sammon's mapping of the SOTM. *Notes* The figure illustrates the proximity of units as approximated by the Sammon's mapping. This is the input to the coloring of the SOTM in Fig. 9.3




Fig. 9.5 Cluster validation of the clustering of the SOTM. *Notes* The figure shows the Silhouette coefficient for $K = 3, 4, \dots, 10$, where large values indicate a good cluster compactness

identify which solutions have the largest explanatory power. Figure 9.5 shows the Silhouette coefficient for $K = 3, 4, \dots, 10$. While it indicates that $K = 3$ is optimal, the Silhouette coefficient shows only minor differences for larger K . The optimality of the 3-cluster solution, as it splits data as per only the time dimension, is a clear indication of significant inherent temporal differences.

To further assess structure information within these three temporal clusters, Fig. 9.6 illustrates the agglomeration process for $K = 3, 4, \dots, 8$. Interpretation of the agglomeration process when increasing K is facilitated, by defining the color coding from the ColorBrewer scheme to be constant for all clusters except for the split one, for which a new color is introduced. The agglomeration process is summarized from low to high K for illustrational purposes, although agglomeration proceeds in a top-down manner. While the 3-cluster solution mainly shows temporal differences in data (cluster 1, red; cluster 2, blue; cluster 3, green), the 4-cluster solution introduces a cluster (4, purple) which broadly speaking coincides with the largest structural changes identified in Figs. 9.3 and 9.4 and the vulnerabilities prior to crises in Fig. 9.2. Second, whereas the 5-cluster solution only adds a small cluster (5, orange) representing temporal differences in the beginning of the analyzed period, the 6-cluster solution derives from cluster 4 a cluster (6, yellow) that covers in broad terms the entire cross section during the end of pre-crisis times. In addition to the 7-cluster solution introducing a separate cluster of the two last quarters (7, brown), the 8-cluster solution derives from cluster 1, the less vulnerable pre-crisis cluster, one cluster in between the most and least vulnerable economies (8, pink). The second-level clus-

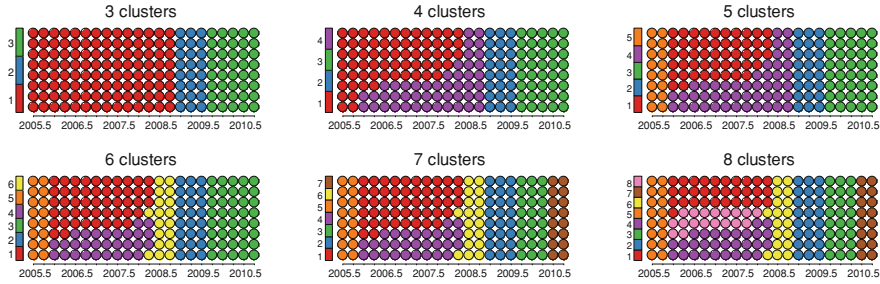



Fig. 9.6 Cluster memberships on the SOTM. *Notes* It illustrates the agglomeration process for $K = 3, 4, \dots, 10$. Interpretation of the agglomeration process when increasing K is facilitated by defining the color coding from the ColorBrewer scheme to be constant for all clusters except for the split one, for which a new color is introduced

Fig. 9.7 Property measures of the SOTM. *Notes* The errors $\varepsilon_{qe}(t)$ and $\varepsilon_{te}(t)$ are computed for time units $t = 1, 2, \dots, T$ and $\varepsilon_{sc}(t)$ for time units $t = 2, 3, \dots, T$ for the final model with an 8×22 array of units and $\sigma = 2.4$



tering of the SOTM illustrates two key messages: (i) temporal trends are strong in these data, and (ii) the increases in univariate vulnerabilities and risks prior to the crisis observed in the previous subsection are illustrated with a cluster (4, purple) that increases in size from 2006Q1 to 2008Q4 in Fig. 9.6.

Property measures for each $A(t)$ provide a third view of the multivariate structures. By exploiting properties of the SOTM, we can observe quantitative characteristics of it. Figure 9.7 shows the temporal variation of the property measures ($\varepsilon_{qe}(t)$, $\varepsilon_{te}(t)$ and $\varepsilon_{sc}(t)$). The quantization errors $\varepsilon_{qe}(t)$ and topographic errors $\varepsilon_{te}(t)$ relate more to qualities across the SOTM, while the structural change $\varepsilon_{sc}(t)$ shows properties in terms of distances between $A(t - 1)$ and $A(t)$. The qualities indicate, as expected, no topographic errors and stable quantization errors over time. Interestingly, the largest structural changes are found in the late-crisis and post-crisis periods, and even more interestingly the largest single change occurs in 2010Q2, when the structures move from the crisis structures towards those during pre-crisis periods. The location of structural changes in Fig. 9.7, while to some extent being illustrated by the SOTM in Figs. 9.3 and 9.4 through bilateral vertical and horizontal differences in colors and Sammon’s topology, are not obvious without an objective quantification of the column-wise distances.

9.2 A Decomposition of Modern Financial Crises

In the previous section, the SOTM was used for visual dynamic clustering over time. In this section, the SOTM is applied to time-to-event data. The SOTM for time-to-event data has a different interpretation for the time dimension t . Rather than representing the time span in data, it represents the time to a specific event. Hence, it takes, for instance, the following form: $t = -T, -T + 1, \dots, T - 1, T$, where T sets the range of time units before and after the event. Whereas the experiments in this chapter use symmetric pre- and post-event spans, the SOTM obviously sets no such restriction. Given an interchanged time dimension, the rest of the functioning of the SOTM follows the standard specifications.

In this section, the time-to-event SOTM is applied for decomposing patterns before, during and after global financial crises from 1990–2011. The standard SOTM in Sect. 9.1 provided an abstraction of patterns in macro-financial indicators before, during and after the global financial crisis of 2007–2008. The time-to-event SOTM herein differs from, or goes beyond, the one above by generalizing the patterns prior, during and after modern systemic financial crises.

Over the period 1990–2011 for the 28 countries in the sample, the approach based upon the Financial Distress Index (FDI) identifies a set of 94 systemic financial crises, of which some crises may last for multiple quarters. These function as the events in our dataset. The time dimension is transformed into time-to-event format by locating all observations from $t - 8$ to $t + 8$, where $t - 0$ are the crisis dates defined using the FDI. In contrast to firm-level failures, countries are bound to experience recurring events. To decrease noise and increase reliability, only observations with one time-to-event stamp are kept in the dataset. For instance, an observation is disregarded if it is a post-crisis period to one event and at the same time a pre-crisis period to a following event, and *vice versa*. It is still worth noting that crises may last for several quarters.

9.2.1 Parametrizing the Time-to-Event SOTM

The architecture of the time-to-event SOTM is set to 6×17 units, where 6 units represent the cross-sectional structures and 17 units represent quarters ranging from $t - 8$ to $t + 8$. The units representing the time dimension is set as to span periods before (8 quarters), during (1 quarter) and after (8 quarters) crises. The number of units representing cross-sectional structures at one period is determined based upon its descriptive value. Again, it is worth noting that the SOTM is not restricted to treat each unit as an individual cluster. Yet, the number of units on the vertical dimension is kept low, as no second-level clustering is applied here. The training phase uses the macro-financial indicators as inputs, while the class variable is only needed for creating the time-to-event data. The final specification of the SOTM is chosen based

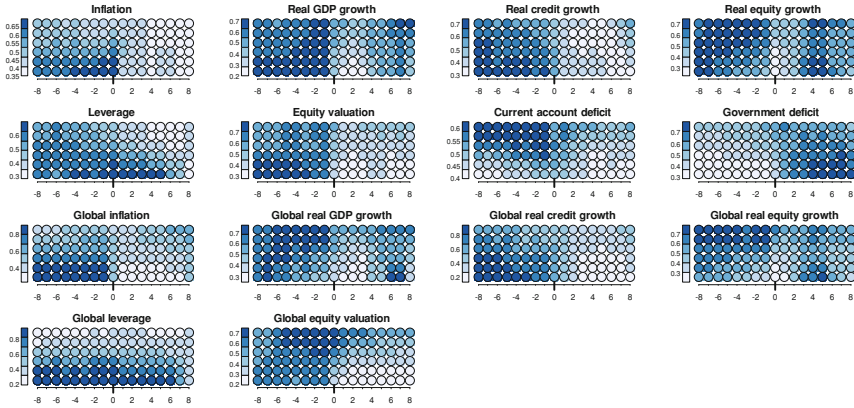


Fig. 9.8 Feature planes for the time-to-event SOTM. *Notes* The figure shows feature planes for the 14 indicators, which are defined in Table 7.1. *Notes* The feature planes are indicator layers of the SOTM in Fig. 9.9. As each data vector consists of 14 indicators, these feature planes show the distribution of each data column on the SOTM grid




Fig. 9.9 A time-to-event SOTM of global financial crises. *Notes* where the cluster coloring shows changes in multivariate clusters structures

upon the quality measures (ε_{qe} , ε_{te} and ε_{sc}). The neighborhood radius σ takes the value 2.8, as it has the highest quantization accuracy, given no topographic errors (Fig. 9.8).

9.2.2 A Univariate View of Crises

The output of the SOTM is a grid of 6×17 multidimensional reference vectors, where the timeline below the figure represents the time-to-event dimension in data (see Fig. 9.9). For an understanding of the univariate patterns behind the multivariate structures, we start by assessing feature planes for individual inputs in Fig. 9.8. The analysis of patterns is divided into three parts: *pre-crisis*, *crisis* and *post-crisis* patterns.

During the *pre-crisis* periods, the lower and the upper part of the grid show different paths to a crisis. The lower part shows high values in inflation, real credit growth, leverage, equity valuation, global inflation, global real credit growth and global leverage. The upper part, on the other hand, shows high values for real equity growth, current account deficit, government deficit, global real GDP growth, global real equity growth and global equity valuation. The *crisis* periods show a decrease in all variables, except for leverage, government deficit and global leverage, as well as a contraction of most variables towards similar values (i.e., a small range). The patterns of *post-crisis* periods can in most cases be divided into early and late patterns, whereas some measures take low values for the entire period. Indicators that take low values from $t + 1$ to $t + 8$ are domestic measures of inflation, real credit growth, equity valuation and current account deficit, as well as global measures of inflation, real credit growth and equity valuation. The pattern of a decrease and a subsequent increase in values illustrates the behavior of many indicators, such as real GDP growth and equity growth, both domestically and globally. Leverage, on the other hand, shows a slow gradual decrease during early periods and a more significant decrease only in the latter part. In addition, government deficits increase significantly in early and late periods, which points to sovereigns being fiscally strained after systemic financial crises.

9.2.3 A Multivariate View of Crises

The two-dimensional SOTM representing the multivariate structures is shown in Fig. 9.9. To assess the multivariate structures of the SOTM, we can again make use of the coloring based upon a Sammon's mapping for illustrating proximity of units. Figure 9.9 illustrates a number of phenomena. Whereas changes from $t - 8$ to $t - 1$, i.e., pre-crisis periods, are gradual, the crisis, i.e., $t - 0$, illustrates a structural break in terms of a shift and contraction in structures. The post-crisis periods from $t + 1$ to $t + 3$, likewise, illustrate the occurrence of large structural changes. Yet, $t + 4$ and $t + 8$ show a contraction of the data and only minor cross-sectional variation and changes over time-to-events. It is worth to note that the orientation of the SOTM is interchanged after the crisis. This is due to the strong contraction that suppresses the values into one, dense cluster. When the SOTM moves to $t + 1$, the initialization based upon t does not guide it enough.

Finally, Fig. 9.10 illustrates changes in property measures ($\varepsilon_{qe}(t)$, $\varepsilon_{te}(t)$ and $\varepsilon_{sc}(t)$) of the SOTM over time. The key message of the figure is that the largest structural change occurs between the $t - 1$ and $t - 0$, i.e., the transition from a pre-crisis to the crisis period. This indicates that the crisis periods are very different from the rest of the data. Another interesting pattern of the structural changes occur in the final period $t + 8$, to which there is no direct explanation. In general, structural changes are larger after crises than before them. The zero level of topographic errors confirms the quality of the topology preservation, as was also previously emphasized.




Fig. 9.10 Property measures of the time-to-event SOTM. *Notes* The errors $\varepsilon_{qe}(t)$ and $\varepsilon_{te}(t)$ are computed for time units $t = -T, -T + 1, \dots, T$ and $\varepsilon_{sc}(t)$ for time units $t = -T + 1, -T + 2, \dots, T$ for the final model with an 6×17 array of units and $\sigma = 2.8$

9.3 Concluding Summary

This section has performed two types of visual dynamic clustering to assess cross-sectional dynamics in multivariate macro-financial indicators. The first decomposition applied the standard SOTM to describing the global financial crisis that started in 2007, whereas the second section applied the SOTM to time-to-event data in order to generalize patterns before, during and after financial crises.

From the viewpoint of macroprudential oversight, this aids in the very key task of identifying build-up phases of risks, vulnerabilities and imbalances in the entire cross section. Hence, out of the three types of systemic risks, this is an approach that truly holds promise for addressing the identification of *the endogenous build-up of widespread imbalances*. Whereas the application herein focused on the global dimension in country-level macro-financial data, the SOTM also provides means for assessing a wide scope of applications for similar purposes, not the least a system-wide focus in data on individual financial intermediaries.

References

Sarlin, P. (2013a). Decomposing the global financial crisis: A self-organizing time map. *Pattern Recognition Letters*, 34, 1701–1709.

Sarlin, P. (2013b). A self-organizing time map for time-to-event data. *Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining (CIDM'13)* (pp. 230–237). Singapore: IEEE Press.

Sarlin, P., & Yao, Z. (2013). Clustering of the self-organizing time map. *Neurocomputing*, 121, 317–327.

Sarlin, P. (2014a). On biologically inspired predictions of the global financial crisis. *Neural Computing & Applications*, 24(3–4), 663–673.

Chapter 10

Conclusions, Limitations and the Future

*All models are wrong.
Some models are useful.*
–George E.P. Box

Early identification of financial instabilities is of interest for a wide spectrum of decision-makers for a wide range of reasons: *policymakers* want to avoid economic fluctuations, *financial market participants* want to earn returns, *businesses* want to set production to optimize profits, and *politicians* want to be re-elected. However, this boils down to the challenge that the Swedish Minister of Finance Kjell-Olof Feldt hinted already in the 1980s: “*Challenging decisions are politically too early until they are financially too late*”. As noted by Korkman (2012), the problem of political populism motivates promoting knowledge and understanding among our fellow citizens, who alas most often are laymen in the field.

It is needless to say that the recent occurrences of instability have stimulated efforts in understanding and predicting financial stress. The work in this book has provided a wide range of tools for macroprudential oversight, whose common denominator is a visual representation. The tools focus on risk identification and assessment, with an ultimate aim to aid in risk communication. It is worth noting that the relevance of visual representations of tools for safeguarding financial stability lies not only in external risk communication, but also in generating insights in internal use to support risk identification and assessment.

This chapter summarizes the key findings of the work in this book, discusses the limitations of the findings, and presents ideas for future research.

10.1 Conclusions, Findings and Implications

The work in this book has put forward visual means for risk identification and assessment. Throughout, the overall task has been to represent high-dimensional data concerning financial entities, be they countries, markets or institutions, on



Fig. 10.1 The process of RQs. *Notes* The coloring of the blocks divides the RQs into RTs. The *red blocks* relate to an understanding of the macroprudential domain and data, the *blue blocks* relate to deriving optimal methods and their extensions, and the *green block* relates to applications of the methods to the task at hand

low-dimensional displays to facilitate the identification, assessment and communication of vulnerabilities and risks. In the introduction to this book, the research objectives (ROs) were said to be two:

- (i) RO1: to choose and extend data and dimension reduction methods such that they meet the needs set by macroprudential oversight and data, and
- (ii) RO2: to apply data and dimension reduction methods in macroprudential oversight to be used by and introduced to the policymaking community.

In order to achieve these two ROs, three research themes (RTs) and five research questions (RQs) were introduced. Thus, a discussion of how the RQs have been answered precedes a discussion of how well the ROs have been met. To refresh memory, Fig. 10.1 presents the RQs in a process format, where the RTs are shown by red, blue and green blocks. Below, we first discuss answers to all five RQs, whereafter we turn to the ROs in the two following subsections.

RQ1: What are the needs for macroprudential oversight? The key aim of Chap. 2 was to give a broad overview of financial systems, financial instability and systemic risks, as well as the reasons for financial systems being fragile. While we discussed the complexity of factors affecting financial systems, how fragilities may build up and what form systemic risks may take, as well as empirical and theoretical underpinnings, an obvious focus of this chapter was on tools and models for macroprudential oversight. Given the mandate of multiple macroprudential supervisory bodies, the starting point ought to be timely and accurate measurement of systemic risks. This resulted in three key systemic risks (*and tools for addressing them*):

- (i) endogenous build-up of widespread imbalances (*early-warning models*);
- (ii) exogenous aggregate shocks (*macro stress-testing models*); and
- (iii) contagion and spillover (*contagion and spillover models*).

These set an inherent need for a broad basis of tools for the identification and assessment of the potential risks, vulnerabilities and imbalances. The chapter concludes by relating the fragilities, risks and tools to an overall macroprudential oversight process. The process clearly illustrates the lack of integration of a third component, risk communication, with risk identification and assessment tools, particularly in the case of macro stress-tests and early-warning models. For contagion models, visualizations based upon network models and graph theory have been and are still gaining further interest within the policymaking community. Yet, the task of representing high-dimensional early-warning indicators on a low-dimensional display has not been addressed in a sufficient manner.

RQ2: What form do macroprudential data take? In all above discussed tasks, the quality of a model is highly dependent on the quality of the underlying data. Chapter 3 discussed data needs and demands for macroprudential oversight, with a particular focus on early-warning models. The chapter identified that country-level data commonly used in early-warning models to represent indicators of risks, vulnerabilities and imbalances are of three types: (i) macroeconomic, (ii) banking system, and (iii) market-based. The broad notion of macroprudential data was untangled into a four-dimensional cube representation, where the dimensions represent time, countries, variables and linkages. The data, while being easy to identify, are not unproblematic. In this vein, the chapter also discussed stylized challenges related to macroprudential data. One obvious conclusion is that today large amounts of data representing risks and vulnerabilities are widely available. A task of central importance is, however, to acknowledge and account for their challenging characteristics, such as missing values, skewed distributions, revisions and publication lags and the general issues of provision and integration of various sources. More importantly, rather than aggregating data into composite indices, the chapter further motivates visualizing these complex data in easily understandable formats to support disciplined and structured judgmental analysis based upon policymakers' experience.

RQ3: Which data and dimension reduction methods hold most promise for the task? The answers to the previous questions, in addition to the general task of this book, set forth a definition of the task at hand: to provide low-dimensional representations of high-dimensional indicators of risks, vulnerabilities and imbalances. A starting point for judging the most suitable methods for the task was put forward in Chap. 4 by providing an overview of data and dimension reduction methods. First, the methods were related to knowledge discovery, data mining, information visualization and visual analytics. Then, the chapter reviewed the basics of classical data and dimension reduction methods and related a comprehensive set of methods in a taxonomy.

Chapter 5 discussed the particular needs for and properties of macroprudential oversight and data in relation to the characteristics of data and dimension reductions, and their combinations. The suitability of three classical, or so-called first-generation, dimension reduction methods for the task at hand was illustrated with qualitative comparisons and illustrative experiments. A key implication of the chapter is that the family of topology-preserving methods with a regular grid shape in general and the Self-Organizing Map (SOM) in particular hold most promise for the task at hand.

RQ4: How should the methods be extended and enhanced for the task? The discussion in Chaps. 4 and 5 concluded that the method of preference for the purposes in this book is the SOM. However, the standard SOM, while holding promise for the task at hand, may be extended in multiple directions. Chapters 2 and 3 spell out the needs and demands for the task at hand. A particular focus of the extensions is related to two tasks that not only answer the demands of macroprudential oversight and data, but have also been stated to be in need of future research in the fields of information visualization and dimension reduction. First, Chen (2005), Wong et al. (2012) highlight a paradigm shift from only visualizing structures to visualizing dynamics. An even further step is to assess dynamics of structures. Second, to be

aware of the quality and distortions of dimension reductions, Wismüller et al. (2010), Wong et al. (2012) stress that they are not an end, but provide only a means to display useful information on top of them, such as evidence, uncertainty and individual data.

To this end, with a key focus on temporality, Chap. 6 first discussed the literature on time in SOMs. The discussion, and subsequent extensions to the SOM paradigm, relate to processing data from the cube representation, i.e., along multivariate, temporal and cross-sectional dimensions, where a key focus is on better processing and visualizing time. The enhancements not only aid in analyzing and visualizing individual cross-sectional and/or time-series data on the SOM, but also contribute to the assessment of overall properties and qualities of the SOM. Extensions to be used with a standard SOM comprise approaches for fuzzification, transition probabilities and assessing shock propagation. The chapter also presented the stand-alone Self-Organizing Time Map (SOTM) for assessing how cluster structures evolve over time (i.e., visual dynamic clustering). The motivation and functioning of the extensions is demonstrated with a number of illustrative examples.

RQ5: How should the methods and their extensions be applied to the task?

The core of this book lies in rather technical applications. Still, an essential part is an adequate understanding of the domain and underlying data, including highly practical issues. Even more important is to make use of methods suitable for the aims of the task at hand. A large share of this book focuses on practical applications of data and dimension reduction methods to macroprudential oversight. First, a framework for building the Self-Organizing Financial Stability Map (SOFSM) is put forward. The second topic concerns a number of extensions to the standard SOM-based model. Third, this book has shown applications of the SOTM and the time-to-event SOTM to risk identification.

Chapter 7 described the construction of the SOFSM. The framework consists of five building blocks: (i) data and dimension reduction based upon the SOM, (ii) identification of systemic financial crises, (iii) choice of macro-financial indicators of vulnerabilities and risks, (iv) a model evaluation framework for assessing performance, and (v) a model training framework for creating parsimonious, objective and interpretable models. Then, the chapter illustrated how the training and evaluation frameworks are applied for constructing the SOFSM, including a range of performance and robustness tests. Finally, the SOFSM provides means to monitor macro-financial vulnerabilities by locating a country in the financial stability cycle on a two-dimensional display.

Turning from model construction to more practical applications, Chap. 8 used the SOFSM for an approach that combines risk identification, assessment and communication. Thus, extensions to the standard SOM are applied in macroprudential oversight, including a fuzzification, transition probabilities and shock-propagation analysis. The SOFSM was also used for illustrating results of stress tests and detecting outliers. In addition, the SOFSM is paired with a stand-alone predictive model to illustrate the complementary role of such approaches. Hence, the SOFSM not only provides means for visual early-warning exercises, but also enable superimposed visualizations of stress test results and potential for contagion.

In Chap. 9, the SOTM was applied in macroprudential oversight in general and risk identification in particular. The SOTM performs visual dynamic clustering for decomposing global financial crises from two viewpoints. The first decomposition applied a standard SOTM to describe how risks and imbalances evolved before, during and after the global financial crisis of 2007–2008. The second decomposition used a time-to-event SOTM to generalize patterns before, during and after modern financial crises from 1990–2011.

The answers to the RQs take us to a discussion of how well the ROs have been met.

10.1.1 Implications for Dimension Reduction

The first objective relates to the choice and extensions of data and dimension reduction methods with respect to the needs for the task in this book. While this has implications for both data and dimension reduction, the key conclusions relate to visualizing data.

This book has shown that the family of topology-preserving methods with a regular grid shape in general and the SOM in particular holds most promise for the task at hand. This relates to four key properties: (i) trustworthy neighbors, (ii) low computational cost, (iii) flexibility for problematic data, and (iv) a regularly shaped grid. Obviously, independent of the topic and field, this conclusion is also applicable to tasks with similar needs and properties. Yet, while the stand-alone SOM holds promise for the task, it has also been extended along multiple directions.

The extensions have been approached from the viewpoint of the task and data at hand. In addition to their high dimensionality, macroprudential data consist of two central components: the cross-sectional and the temporal dimension. Whereas a key focus herein has been on better processing and visualizing time in SOMs, the cross-sectional dimension has throughout also been of importance. In addition, the enhancements not only aid in analyzing and visualizing cross-sectional and/or time-series data on the SOM, but also contribute to the assessment of overall properties and qualities of the SOM. The most central extensions to the standard SOM are three:

- (i) **Fuzzifications** aid in visualizing temporal belongingness to clusters of individual data and cluster distance structures on the SOM.
- (ii) **Transition probabilities** aid in visualizing probabilities of transition of individual data and for assessing the cyclical and temporal structure on the SOM.
- (iii) **Network topologies** aid in understanding links between data by illustrating a network topology on the data topology of a SOM.

In addition, the SOM has been illustrated to enable **contagion analysis** through neighborhood relations on the grid structure, **scenario analysis** by visualizing transitions in the case of changes in data, and **outlier analysis** in the form of distances of data to the SOM grid.

A concept of its own is the Self-Organizing Time Map (SOTM) that goes beyond the standard SOM representation for visual dynamic clustering. It enables thus a visualization of how multivariate cross-sectional cluster structures evolve over time. The temporal changes can be assessed univariately and multivariately. The time-to-event SOTM and the second-level clustering of the SOTM enhance it by enabling assessment of patterns before, during and after specified events and by providing an objective means for assessing temporal changes in cluster structures.

10.1.2 Implications for Policy Use

The second objective concerns applications of data and dimension reduction methods for policy use. This book has created the SOFSM, which uses data and dimension reduction methods for mapping the state of financial stability, as well as the above discussed extensions to the SOFSM. The SOFSM is a two-dimensional representation of a multidimensional financial stability space that allows disentangling the individual sources of vulnerabilities impacting on systemic risks and can be used to monitor macro-financial vulnerabilities by locating a country in the financial stability cycle, being it either in the pre-crisis, crisis, post-crisis or tranquil state. The technical qualities and robustness of the SOFSM have been tested by varying the SOM parameters, thresholds of the models, the policymakers' preferences and the forecast horizon. In addition, the model would not only have correctly called the financial crisis of 2007–2008 in the United States (US) and the euro area in mid-2005 (even when accounting for publication lags in data), but also communicated the results in an easily interpretable format. Hence, the SOFSM provides a framework for future works to follow in order to create financial stability maps. Moreover, the SOTM provided a means to observe how risks build up over time in the cross section. From the viewpoint of systemic risk, this is a highly relevant concept as it enables a perspective beyond individual data. Likewise, the time-to-event SOTM provides an overview of alternative roads to and from a crisis.

The use and acceptance of the SOFSM in policy use has been indicated by practical implementations, as well as communication to academics and practitioners. In addition to the academic communication in terms of published papers, the work in this book has been widely communicated to practitioners. The SOFSM has been published as a working and discussion paper both at the European Central Bank (ECB) and the Bank of Finland, has been included as a special feature in the Financial Stability Review of the ECB see ECB (2011), and has been highlighted by the ECB's vice president as a promising approach.¹ In addition, the SOFSM is a project in the Macro-prudential Research Network (MaRs) (ECB 2012). It is also currently being implemented at multiple central banks and financial institutions for a map to be used for external and internal communication, as well as being implemented in an interactive browser-based application by infolytika (<http://risklab.fi/demo/macropru/fsm/> and <http://risklab.fi/demo/macropru/fsmt/>).

¹ The speech can be found here: <http://www.bis.org/review/r120619a.pdf>.

10.2 Limitations

It is needless to say that the work in this book has its limitations. On the one hand, these may be related to simplifications due to various challenges in modeling, such as data provision and availability, or other simplifications related to the context of today's world, such as available computing power. On the other hand, these provide opportunities for further research, such as extensions to methods and evaluations of models. Again, the discussion is separated according to the two ROs: methods and applications.

10.2.1 Limitations of Methods

The method extensions presented in this book are obviously restricted to apply to only specific application domains and data. In addition, the connection to information visualization of the work in this book also needs to be limited, such as whether and to what extent it provides means for interactive visualizations.

One major concern of the SOTM, and the SOM in general, is that instances are most often treated as being of uniform importance. The significance of such an approach can be easily shown with an example from the topic of this book. Think about the SOFSM, which uses data from 28 economies. It is indeed important to have a cross-sectional perspective, as the number of crises in individual countries is rather small and capturing a wide variety of crises is often strived for. Yet, the importance of, for instance, Sweden and the US for such a model is most likely not equal. Likewise, one could assume that the importance of an economy varies over time due to numerous reasons, such as size of the banking sector, structures of the banks and other measures of interdependence. In the case of the standard SOM, this motivates a weighted approach that learns from data based upon instance-specific importance values. Likewise, this also applies to the SOTM, but from another perspective. A key use of the SOTM is for the understanding of vulnerabilities and risks that are building up in the cross section. Here, however, it is of central importance whether or not the imbalances are growing in systemically important economies, which indeed varies over time and across countries.

Another line of limitations relate to information visualization. An essential part of information visualization is the use of a manipulable medium, which allows users to vary parameter values to interactively explore properties of data. As the work in this book provides models, or constructs, the inclusion of interaction techniques is discussable in that they are planned to be included at the level of instantiations. While all products of this book are not ready-to-use tools with user interfaces and interaction mediums, most of the applied and derived methods could still be easily combined with a user interface and a range of parameters for interactively exploring properties of data. In fact, this is an essential part of the above discussed implementations, not the least in the case of Financial Network Analytics.

10.2.2 Limitations of Applications

The applications presented in this book come with a number of limitations. First and foremost, the usefulness of all models for policy use need to be related to the so-called Lucas critique. The rest of the limitations are partly related to the underlying data and the evaluations of the results.

The Lucas critique discusses how model accuracy is dependent upon potential feedback effects of changes in expectations and human behavior. The reasoning behind the Lucas critique has to be acknowledged (Lucas 1976): *“any change in policy will systematically alter the structure of econometric models”*. Yet, as also discussed by Bisias et al. (2012), the aim of more accurate early-warning signals relates to few undesirable effects in terms of changes in behavior and expectations. First, independent of the fact that changes in behavior might discount the impact of policies, the more accurate the risk measures are the more accurate are the inputs to policy. Second, the key intent of early-warning signals is encouraging individuals and institutions to take actions on their own, rather than only relying on actions of governments. Yet, it is also of importance to be cautious in the use and communication of early-warning signals, such as the risk of self-fulfilling prophecies. In order to have indications of gradual changes in the state of financial stability, and allow for individuals to take own corrective actions, this further motivates the use of timely and frequent data.

Measuring financial instability is indeed a challenging task, especially when the aim is a global approach. Accordingly, the data are limited by multiple challenges. First, the dataset is entirely missing indicators measuring aggregate risks in country-level banking sectors. This is a result of having a global approach, as all emerging market economies do not report these types of data. Second, the euro area is included as an aggregate. While the focus in this book was to have a global approach, it is obvious that it would be of interest to analyze individual euro-area countries as well. Likewise, any other excluded country could be of interest. Third, the class variable is discretized from the Financial Distress Index (FDI). While discretization leads to some loss of information, the fundamental idea of predicting vulnerabilities prior to financial crisis, i.e. pre-crisis periods, does not allow modeling a continuous index measuring contemporaneous stress. Further, one could also claim that the results are dependent on the time and country frame and thus not generalizable. Indeed, the results are only restricted to the used data, yet the dataset covers a global set of economies from 1990 onwards. Thus, the results may be said to apply to modern financial crises.

The second set of limitations relate to the qualities and validity of the models. As user satisfaction and perceived usefulness of the models has not been tested, this book concerns no claims related to measuring how good the proposed models are in terms of applicability for visualizing data. This relates to the question of internal and external evaluations. While a wide range of internal measures have been used to assess and calibrate qualities of the models, this book has not shown thorough

evaluations of the models in terms of external measures. Yet, the communication to, interaction with and acceptance by domain experts provides an indication of external quality.

10.3 Future Research

Future research is to a large extent, yet obviously not entirely, directed by the above discussed limitations. While this book already includes a large share of material on the SOTM, there exists large potential in both extending it to many directions and applying it to a wide range of tasks, particularly in macroprudential oversight.

10.3.1 Future Methods

One idea for future research is the above discussed weighting scheme for the SOM. A starting point to this can be found in Sarlin (2013), which may be extended in a number of directions. In particular, the weighting could be applied, as again discussed above, to the SOTM. Additionally, while the x axis of the SOTM has represented time and time-to-event dimensions, it should not be restricted to any specific variable. Depending upon the task at hand, the SOTM can be performed over a wide variety of dimensions, such as age of customers in customer segmentation and steps of a process in industrial process monitoring. Moreover, whereas the standard SOTM reduces both data and dimensionality by projecting data onto a two-dimensional grid of units, the reduction of the dataset could still be enhanced by also reducing the time dimension. This enables a focus on only temporally relevant parts of the time dimension. Obviously, the SOTM could also be extended to a three-dimensional case, but this would involve strong interaction techniques to be able to exploit the details provided in a three-dimensional cube. Indeed, the simple, two-dimensional representation of the SOTM is a merit.

Another line of research could be to combine the approaches presented in this book. For instance, the SOTM representation shows currently only changes in cluster structures, but neglects the transition patterns. Transition probabilities would be a straightforward approach for understanding patterns of who is changing, in addition to only knowing the occurrence of changes.

The two-dimensional SOFSM is mainly illustrated with four crisp clusters (i.e., the financial stability states), yet structures in real-world data are seldom crisp. One promising approach is the projection-based coloring scheme for revealing cluster structures provided by Kaski et al. (2001). The approach has, however, a number of limitations: (i) the objective function takes a complex form and involves a number of parameters to be specified, (ii) the coloring method is not flexible for different types of projection methods, and (iii) most variation is restricted to occur in two dimensions of hue, which implies a slight distortion to the mapping. In this vein, it is

of interest to explore possibilities of finding a general, yet simple, solution to cluster coloring. Along these lines, an interesting approach would also be to better integrate the elements of information visualization with dimension reduction methods. In practice, this not only relates to the tasks of accounting for perception and cognition in data graphics, as has briefly been discussed herein, but also, and in particular, the task of interaction techniques. As noted by Wismüller et al. (2010), and illustrated in this book, dimension reductions should only be treated as a starting point for the general visualization process and visual analytics in particular.

10.3.2 Future Applications

While the scope of future applications is qualitatively unlimited, and mainly dependent upon the aims of the modeling task, there exists a range of interesting directions in applications that are judged to be worth pursuing. To start with, when data provision and availability become better in the future, as they are expected to do, the above discussed limitations related to challenges in macroprudential data can be improved. In particular, banking sector data should be collected to also include imbalances in country-level banking sectors. Another improvement is to include individual euro-area countries. One direction related to early-warning models and indicators is related to a comprehensive comparison of approaches. The literature is missing an objective evaluation of the relative performance of different methods, given the same indicators, split of in- and out-of-sample data, evaluation measures and model-building schemes.

In general, the SOM can be seen as a promising approach to communicating any types of multivariate panel data, where entities may be firms, countries, assets, individuals, etc. In particular, it provides means for building a low-dimensional display on top of which individual data may be visualized. Hence, the SOM, as well as the SOTM, could equally well be used by, for instance, the European Banking Authority at the level of financial intermediaries and the European Securities and Markets Authority at the level of financial markets and securities.

In this vein, an interesting application of the SOTM would be to assess how risks and vulnerabilities have built up in the cross section over recent years, as is done in this book, but from more granular perspectives. In contrast to macro-level applications, which provide a global view in the cross section, micro-level data would enable more detailed information related to accumulated risks and their changes over time. If one wants to focus on even more granular data, one could move from firms to individual securities or assets, whose structures might be of interest. One task could be to assess how asset correlation structures have evolved over time, with a focus on illustrating dynamics during recent shocks, such as market reactions to the flash crash and the failure of Lehman Brothers, where significant contractions in correlation structures are to be expected. Likewise, the SOTM could be applied to a wide range of other tasks. Related to the visual models in this book, a common task missing is a user

evaluation to measure the perceived usefulness of the models. This would obviously be of interest in the context of policymakers involved in macroprudential oversight.

Still, given methods for identifying risks and vulnerabilities, the key question for future research to answer is: *how should policymakers be persuaded to take correct(ive) actions?* This boils down to the fact that behind policy lies a mishmash of politics and economics, not the least in the recent European decisions. Some of these challenges are exemplified by the below quotes.

“Often, public officials have two unfortunate incentives: to give undue attention to worst-case scenarios and to pay no attention to them at all. Sometimes their electoral prospects, or their overall popularity, depend on one or the other. Before the attacks of 9/11, almost all American officials neglected the need for better security at airports, not least because the public would have strongly resisted significant additional burdens on air travel.”
–Cass R. Sunstein, Worst-Case Scenarios

“We all know what to do, we just don’t know how to get re-elected after we’ve done it.”
–Jean-Claude Juncker, President of the Euro Group and Prime Minister of Luxembourg, Le Soir, 2 July 2007

This being said, one could argue that means for better risk communication, such as some of the visuals put forward in this book, promote the knowledge and understanding among our fellow citizens. This implies that a soar in research on external risk communication could provide an improved basis for persuading policymakers to early enough corrective actions.

References

- Bisias, D., Flood, M., Lo, A., & Valavanis, S. (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics*, 4, 255–296.
- Chen, C. (2005). Top 10 unsolved information visualization problems. *IEEE Computer Graphics and Applications*, 25(4), 12–16.
- ECB. (2011). Mapping the state of financial stability. In: *Financial Stability Review*. Germany: Frankfurt: European Central Bank.
- ECB. (2012). *Report on the first two years of the macro-prudential research network*. Frankfurt, Germany: European Central Bank.
- Kaski, S., Venna, J., & Kohonen, T. (2001). Coloring that reveals cluster structures in multivariate data. *Australian Journal of Intelligent Information Processing Systems*, 60, 2–88.
- Korkman, S. (2012). *Talous ja Utopia*. Jyväskylä, Finland: Docendo.
- Lucas, R. (1976). Econometric policy evaluation: a critique. *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.
- Sarlin, P. (2014). A Weighted SOM for classifying data with instance-varying importance. *International Journal of Machine Learning and Cybernetics*, 5(1), 101–110.
- Wismüller, A., Verleysen, M., Aupetit, M., & Lee, J. (2010). Recent advances in nonlinear dimensionality reduction, manifold and topological learning. *Proceedings of the European Symposium on Artificial Neural Networks (ESANN 10)* (pp. 71–80). Belgium: Bruges.
- Wong, P., Shen, H., Johnson, C., Chen, C., & Ross, R. (2012). The top 10 challenges in extreme-scale visual analytics. *IEEE Computer Graphics and Applications*, 32(4), 63–67.