

Eric Anderson

# Social Media Marketing

Game Theory and the  
Emergence of Collaboration

 Springer

# Social Media Marketing



Eric Anderson

# Social Media Marketing

Game Theory and the Emergence  
of Collaboration

 Springer

Eric Anderson  
3036 NE Ainsworth St.  
Portland, Oregon 97211  
USA  
eanderson@whitehorse.com

ISBN 978-3-642-13298-8 e-ISBN 978-3-642-13299-5  
DOI 10.1007/978-3-642-13299-5  
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2010931519

© Springer-Verlag Berlin Heidelberg 2010

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, 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.

*Cover design:* WMXDesign GmbH, Heidelberg, Germany

Printed on acid-free paper

Springer is part of Springer Science+Business Media ([www.springer.com](http://www.springer.com))

## Preface

Social media marketing has been heralded as a sea change in the marketer-consumer relationship, but its rapid growth and rabid following among marketers has also produced a sea of confusion. Lacking any durable framework for understanding how, why, and on what terms the consumer relationship has changed under social media, marketers pursue new venues for their newness alone – with decidedly mixed results.

This book finds a theoretical framework for social media marketing in the science of game theory, with its focus on adversarial but mutually dependent relationships. Originally developed to guide nuclear brinkmanship policy during the Cold War, game theory provides the foundation for an evolutionary view of social media marketing. Through fascinating game theory concepts like the Prisoner's Dilemma, the Stag Hunt, Self-Command, and Job Market Signaling, this study uncovers the cooperative trends that brought marketing to its present state and points the way toward marketing's future course.



# Contents

- Chapter 1: Surviving the Customer** ..... 1
- 1.1 The Origins of Game Theory ..... 5
- 1.2 Game Theory, the New Media, and the NEW New Media ..... 7
- 1.3 The Payoff Matrix ..... 8
  
- Chapter 2: Zero-Sum Games in Traditional Marketing** ..... 13
- 2.1 Zero-Sum Games and the Problem of Transparency ..... 14
- 2.2 The Zero-Sum of Pricing Strategies ..... 16
- 2.3 The Wisdom of Randomization ..... 18
- 2.4 Randomization and A/B Testing ..... 20
- 2.5 The Hazards of Entrenchment ..... 21
- 2.6 Making Zero-Sum Work ..... 25
- 2.7 Mastering Randomization ..... 26
- 2.8 Doing Better than Zero-Sum ..... 27
- 2.9 The Cautionary Tale of Banner Click-Through Metrics ..... 28
  
- Chapter 3: The Prisoner’s Dilemma and the Emergence of Cooperation** ..... 35
- 3.1 The Great Consumer Opt-Out ..... 37
- 3.2 The Shifting Ground of Consumer Trust ..... 40
- 3.3 The Marketer’s Loss of the Informational Advantage ..... 41
- 3.4 The Prisoner’s Dilemma ..... 43
- 3.5 The Hidden Allure of Cooperation ..... 45
- 3.6 The Iterated Prisoner’s Dilemma ..... 46
- 3.7 The Persistent Problem of Bad Apples ..... 47
- 3.8 The Enduring Relevance of TIT FOR TAT ..... 49
- 3.9 The Dangers of the Death Spiral ..... 50
- 3.10 The Marketer’s Dilemma ..... 51



<b>Chapter 4: Consumer Revolt and the Rising Cost of Defection</b> .....	53
4.1 The Unreasonable Standard of Immediate Action .....	54
4.2 The Missing Link in the Evolutionary Chain: Paid Search .....	55
4.3 How Paid Search Gave Free Riders the Boot. ....	56
4.4 Enforcing Rules of Cooperation. ....	58
4.5 Social Media and the Peculiar Pleasures of Punishment .....	59
4.6 “United Breaks Guitars” Breaks Through .....	60
4.7 Recalculating the Cost of Defection .....	62
4.8 Changing the Defection Stakes for Banks: Debtor Videos .....	63
4.9 Attention Saturation and the Limitations of Punishment. ....	64
4.10 The Motrin Moms and Social Media Backlash .....	65
4.11 TIT FOR TWO TATS and the Virtues of Forgiveness .....	68
4.12 The Enduring Appeal of Punishment .....	69
<b>Chapter 5: Sustaining Marketer-Consumer Cooperation through Coordination Games</b> .....	71
5.1 Cold War Coordination Games. ....	72
5.2 The Stag Hunt. ....	74
5.3 Dell and the Exposed Flank Coordination Strategy. ....	75
5.4 Greenpeace Exposes Its Flank .....	77
5.5 The Neutral Ground Coordination Strategy. ....	79
5.6 Social Media as Neutral Ground .....	80
5.7 Coca-Cola’s Neutral Ground .....	82
5.8 Skittles and the Limits of the Neutral Ground Strategy .....	83
<b>Chapter 6: Crowdsourcing and Schelling’s Theory of Self-Command.</b> .....	87
6.1 Introducing Self-Command .....	88
6.2 Self-Command in Marketing .....	90
6.3 Blogging as Self-Command .....	90
6.4 Starbucks’ Crowdsourcing as Self-Command .....	93
6.5 What Crowdsourcing Can and Can’t Do .....	95
6.6 Dealing with Free Riders and Bad Actors .....	96
6.7 Schelling’s Focal Point. ....	97
<b>Chapter 7: Content Popularity and Spence’s Theory of Costly Signaling</b> .....	101
7.1 The Theory of Costly Signaling .....	102
7.2 Traditional Advertising and Costly Signaling .....	104
7.3 The Erosion of Costly Signaling in Super Bowl Advertising .....	105

- 7.4 How Web Transparency Disrupts Costly Signaling . . . . . 107
- 7.5 The Evolution of Costly Signaling on the Web . . . . . 108
- 7.6 Google Changes the Costly Signaling Currency . . . . . 109
- 7.7 Paid Search and Popularity-Based Signaling . . . . . 112
- 7.8 Noise in the Signaling System . . . . . 113
- 7.9 Popularity-Based Signaling in Social Media . . . . . 115
- 7.10 The Disruptive Effects of Popularity-Based Signaling . . . . . 116
- 7.11 The Perils of Negativity in Popularity-Based Signaling . . . . . 119
- 7.12 Popularity-based Signaling in Presidential Politics . . . . . 120
- 7.13 Noka and the Disruption of Costly Signaling . . . . . 122
- 7.14 Sock Puppetry as Noise in the System . . . . . 126
- 7.15 Target Learns the Rules of Popularity-Based Signaling . . . . . 127
- 7.16 Target and the Problem of Capacity . . . . . 130
- 7.17 Content Aggregators and the Evolution of Costly Signaling . . . . . 132
- 7.18 The Three Rules of Popularity-Based Costly Signaling . . . . . 133
- 7.19 Paying the Costly Signal to Go Viral . . . . . 136
- 7.20 A Popularity-Based Success Story: Blendtec . . . . . 137

**Chapter 8: Kapferer’s Prism and the Shifting Ground**

- of Brand Identity . . . . . 141**
- 8.1 Branding by Definition . . . . . 142
- 8.2 Branding as Reciprocal Relationship . . . . . 144
- 8.3 The Traditional Limits on Brand Engagement . . . . . 146
- 8.4 Consumers Assert Brand Ownership: Nike . . . . . 148
- 8.5 Consumers Assert Brand Ownership: Apple . . . . . 149
- 8.6 Kapferer’s Prism . . . . . 150
- 8.7 Rapaille and the Brand Imprint . . . . . 152
- 8.8 The Limitations of Traditional Branding . . . . . 153
- 8.9 Introducing Brand Monitoring . . . . . 154
- 8.10 Brand Monitoring and Ethnographic Analysis . . . . . 156
- 8.11 Diffusing Defection through Brand Monitoring . . . . . 157
- 8.12 Crowdsourcing Brand Identity . . . . . 160

**Chapter 9: Maxwell’s Demon and the Dwindling Supply**

- of Consumer Attention . . . . . 165**
- 9.1 The Problem of Overtaxed Attention . . . . . 166
- 9.2 The Volunteer’s Dilemma and the Tragedy of the Commons . . . . . 167
- 9.3 Social Media’s Capacity Problem . . . . . 169
- 9.4 The Risks of Quantification . . . . . 170
- 9.5 Social Traps and Counterreinforcers in Social Media . . . . . 172
- 9.6 Voting with their Feet: Why Quality Matters in Social Media . . . . . 174

9.7 Pancake People and the Problem of Information Overload . . . . .175

9.8 Social Media’s Answer to Google’s Flattening Effects . . . . .177

**Bibliography** . . . . .181

# Chapter 1: Surviving the Customer

*ABSTRACT: The last decade has born witness to rapid and sweeping changes in the marketer-consumer relationship. Marketers have gained entirely new advertising and marketing platforms, but consumers have simultaneously gained an unprecedented degree of empowerment in marketing relationships, beginning in the dot-com era and culminating in the present social media era. The science of game theory, which is used to analyze mutually dependent conflicts, has particular relevance to this changing landscape, because its focus on conditions for cooperation and defection aptly describe the choices available to marketers and consumers in social media marketing. My evolutionary – as opposed to revolutionary – view of social media marketing holds that the marketer-consumer relationship can evolve toward mutual cooperation, and that an examination of digital marketing’s evolution will yield clues as to the conditions necessary for the success of social media marketing.*

When the history of early 21<sup>st</sup> century marketing is finally written, the present era is sure to be remembered as proof of the ongoing relevance of Georg Hegel’s dialectic – the theory that history moves forward in a cycle of *thesis-antithesis-synthesis*. Marketing’s version of this dialectic is something like *hype-backlash-reality*. In the last decade alone, we have witnessed a continuous cycle of feverish embrace followed by strict disavowal, followed by sober acceptance, as marketers rushed headlong into Web marketing, then ran screaming from it, then crept gingerly back to it. The first cycle was ushered in with the new millennium, in the dot-com gold rush. It reached its antithesis in the spring of 2001 when the bubble burst, and one dot-com business after another found its foosball tables and espresso machines in hock mere months after their raucous IPOs.

From the perspective of a decade’s distance, it is easy to forget that in many ways, the era’s excesses were not so much a denial of the inevitable crash but an earnest, if ill-fated, attempt to rise above it. It has become fashionable to wonder at the fantastically bad business models that somehow

garnered heaps of venture capital, but it's worth pointing out that nearly every failed dot-com business model was flawed in precisely the same way: at the end of the day, marketers couldn't say for sure how customers would behave in a brand new marketing and sales environment. Would they buy impulse items like books and CDs online? Yes. Would they buy their groceries online? No. How was one to know?

At the time of this writing, the marketing world finds itself in the throes of a new dialectical cycle: the current explosion of interest in social media marketing has some of that heady feeling of the dot-com era. But it feels different, too – due in no small part to its arrival coinciding with a global economic downturn that has forced brands to weigh *every* marketing investment with careful deliberation. But there has also been evolutionary progress, as Hegel's dialectic implies: marketers have, consciously or not, learned things from the previous digital marketing era that temper their approach to this one.

The question of why social media marketing constitutes an *evolutionary* stage in the marketer/consumer relationship, and how such an evolutionary view can produce more effective marketing, are the main subjects of this book. In order to make the thread of progress visible, I'll start with a basic argument about what the dot-com era was really about.

Dot-com marketing is often remembered as a kind of bacchanalia of over-spending that reached its apex with Super Bowl XXXIV, the so-called "Dot-com Bowl," in which 17 of the 36 advertisers were newly minted dot-coms, paying an average of \$2.2 million for 30 seconds of air time. (Of these 17 brands, only 3 remain intact today) (Elliott). In this rarified environment, the marketing excesses we now decry followed a twisted but not incomprehensible logic: just *act* like a winner long enough to win. In nearly every instance, the goal was to get as many prospects as possible to visit the site, then determine how to monetize them later. Companies that were, for a variety of reasons, able to keep the wolves from their door long enough to shift customer spending habits in this new epoch generally survived.

The problem was, of course, that consumers proved equally adept at maximizing short-term gains in this new environment without necessarily developing any bankable loyalties. Take the notorious case of the company that became the poster child for dot-com short-sightedness: Kozmo.com, the urban delivery service that fulfilled the ultimate consumer fantasy of instant gratification – a candy bar, a video, a pint of ice cream – delivered to your door in under an hour.

At its peak in 2000, Kozmo boasted 400,000 customers in 11 cities. In this same period, its average order size was \$5, against an average delivery cost of \$7.50. As its former director of logistics John Wu noted in *Supply*

*Chain Management Review*, in what could serve as the definitive understatement of the era: “Eventually, the math caught up with us.” Indeed.

Wu’s article describes the company’s failed efforts to shift to a sustainable value equation, in which customers would be forced to accept higher minimum orders and delivery fees. It didn’t work. Why not? Because their customers, acting in their own short-term self-interest, had no real incentive to change. Many customers quit using the service, and fewer new customers joined. Given the opportunity to *cooperate* by modifying their behavior and continuing to enjoy the service, customer chose instead to *defect*, i.e., to fulfill their snacking needs elsewhere, even if it doomed the company.

The tendency of customers and prospects to *cooperate* in some instances and *defect* in others is the basis for the evolutionary theory outlined in this book. The terms I am using to describe these two binary modes of behavior belong to a very different field of study, one that bears little obvious relevance to the field of marketing. Cooperation and defection are the core concepts at the heart of *game theory*, a field of mathematics and logic that has primarily concerned itself with the study of geopolitical maneuvering and macroeconomics.

To illustrate what such a field might have to say about the dot-com implosions and other marketing phenomena, consider again the example of Kozmo.com. In the September 4, 2000 edition of *The New Yorker*, the magazine’s financial columnist James Surowiecki singled out Kozmo.com as the canary in the dot-com coal mine. Titling his essay “How Kozmo is Getting Killed By Its Customers,” Surowiecki heralded an era of customer tyranny, claiming “Never before have companies so gleefully abased themselves by subsidizing their customers’ purchases, catering to every whim, and burning up tens of millions of dollars in pursuit of that elusive thing called ‘loyalty’.” Surowiecki warned that customers had become “little terrors” whose abased behavior, once encouraged, would be impossible to modify, necessitating a “New Economy mantra: Know when to fire your customers” (Surowiecki 2000).

Surowiecki’s analysis of Kozmo’s shortcomings was not unique or even particularly prescient, but the language he used to frame the problem – the language of global conflict – offers a fresh perspective on the era’s marketing overkill. If the customer and the marketer are really at war, don’t they both lose when one or the other is vanquished? Kozmo’s CFO took particular umbrage at the article’s stark terms; in a letter to the editor a few weeks later, he sniffed, “We are not ‘getting killed’ by our customers. Our commitment to them...is the reason we’ve not only survived but grown.” He promised to invite Surowiecki to the company’s tenth-anniversary party. Exactly five months later, Kozmo shut down operations for good.

I draw attention to the terms used to describe this particular marketer-customer interaction because they are more than a journalist's dramatic device: they point to a fundamental shift in the relationship from one-way conquest to two-way contest. The language of battle has long been the *lingua franca* of marketing: marketers talk about "killer creative" or a "dead list," and the term "campaign" itself is military in origin. Marketers must have adopted this discourse for the same reason that comedians talk about an act that "killed" or "died;" it's a way of being reminded that success or failure is in the hands of an audience whose shifting moods, loyalties, and interests can be lethal.

Multiple game theory concepts can be used to illuminate Kozmo's dilemma, but for the sake of merely introducing ideas that will be explored in detail in later chapters, I'll highlight just one: The scenario of the lazy Kozmo customer, who causes the company to lose money on every transaction by satisfying a series of \$5 whims, is a good example of game theory's *volunteer's dilemma*. The dilemma is simply that the customer's short-term self-interest – getting a quick pint of ice cream – is in direct conflict with his long-term self-interest – continuing to get those pints delivered by ensuring that Kozmo stays in business. By modifying his purchase behavior, the customer could, in effect, volunteer to help Kozmo develop a more sustainable business model. Game theorists and sociologists are keenly interested in what makes the subject set aside his short-term interest for a long-term reward; start-up marketers seeking to sustain customer relationships past multiple rounds of funding ought to be interested as well.

On the surface, the notion of the Kozmo customer curbing his impulse purchasing in order to help the company survive is patently absurd. It's unlikely that the customer is even aware that his purchase habits are bad for the company, and it's even more unlikely that he would modify those habits if he did know. Doing so would be tantamount to insisting on paying full price for a sale item because we believe that such discounting is bad for the store's bottom line. We naturally assume that the store knows better, and is acting in its own self-interest, just as we are.

The existence of such dilemmas in the business world is unremarkable; businesses often depend on "loss leading" with customers in the short term in the hopes of gaining their loyalty in the long term. What *is* remarkable is that we have entered an era in which it is no longer unprecedented or absurd for customers to *consciously forego their short-term interests* in cooperation with a brand, in order to serve some longer term mutual interest. I'll introduce examples of this cooperation in later chapters. This is not an outbreak of consumer altruism; consumers cooperate on the basis of a well-defined, if not well-understood, set of rules and incentives, and

a marketer who understands game theory will be able to put these rules to good use.

It is the combative nature of the marketer-customer relationship that make it ripe for game theory analysis, and it is the shift in the playing field that's taken place over the last decade that makes such analysis not simply valid but enormously useful at this moment. In traditional marketing, a campaign might die, but not because the customer killed it. At worst, the customer might choose not to play, which would, in effect, mean not being a customer at all. For a customer to "kill" in the way that Surowiecki means – an active participation in shaping the brand's fate – they would need to be "armed" to a degree that hasn't traditionally been available to them, with knowledge of the playing field and with a willingness to engage their opponent.

## 1.1 The Origins of Game Theory

This book will plumb the depths of individual game theory concepts as it trace the evolution of social media marketing, but it's worth spending a few pages on the field's origins and ground rules here at the outset. As a subset of the field of mathematics, game theory is a relatively recent development, originating with John Von Neumann and Oskar Morgenstern's 1944 study *Theory of Games and Economic Behavior*, and a relatively narrow field, having been defined and refined by just a handful of main theorists over the last half-century.

For the layperson, and certainly for the marketer, it's largely an unknown field, and that obscurity would suit its founders and apostles just fine. They foresaw very limited applications for the theory based on a stringent set of conditions. Nevertheless those conditions have been loosened over time, so that game theory analysis has been applied to subjects as diverse as business decision-making and pop culture phenomena like reality television.

Both game theory purists and general practitioners seem to agree on this basic definition: wherever two self-interested parties have *both* opposing and mutual interests in the outcome of a conflict, game theory potentially has something to say about it. The marketer and the customer, for instance, have both opposing interests – each wants to maximize their return on the deal – and mutual ones, as both hope to make a deal in the first place. In the same way, the Cold War players had opposing geopolitical interests – each wanted to maximize their share of global influence – but also a mutual interest in avoiding global annihilation.

From that basic premise, game theory in its pure form provides a rather strict set of requirements, mostly by insisting on an almost mechanical



degree of rational behavior. The game theory pioneer and Nobel Laureate John Nash contributed the crucial concept of equilibrium – explored in several later chapters – that defines scenarios where each player cannot improve their outcome by acting unilaterally and therefore achieves a stable solution. Equilibrium, by its very nature, requires strict conditions in order to be true. After all, exceptions to *any* rule create instability in the rule, i.e. they challenge its “ruleness.” So Nash insisted on both players’ perfect knowledge of the options available to them and the absolute rationality – in the sense of advancing the player’s self-interest – of every move in the game.

In a mathematical formula, absolute rationality ensures that the same results can be reproduced each time, because psychology is taken out of the equation. When we put psychology back in by trying to apply game theory to the real world, we lose some stability, but we find that the ideas still hold true. Nash’s and von Neumann’s employer, the RAND Corporation, understood this: as a quasi-governmental think tank, their mission was to find practical applications of game theory to the geopolitical scene. Nash’s concept of equilibrium, for instance, is the underlying logic for the dominant policy of the nuclear arms race: Mutually Assured Destruction, or MAD. The MAD theory held that if both the United States and the Soviet Union maintained stockpiles of nuclear weapons sufficient to destroy the world many times over, neither side could advance their own interests more by defecting – launching a pre-emptive nuclear strike – than by cooperating. And so cooperate they did.

I mention this example because it illustrates very well how the introduction of human psychology into the equation can de-stabilize things but still allow the theory to hold. Over the course of four decades of MAD, human irrationality did indeed de-stabilize the system. When Khrushchev banged his shoe on the desk at the UN and shouted, “We will bury you,” or when Reagan did a live sound check for a radio address by joking that he had outlawed Russia forever and would begin bombing in five minutes – these were not international diplomacy’s finest moments. But we survived these tremors of irrationality, the equilibrium held, and today we’re alive to pursue new applications for the products of game theory’s founders.

Ironically, the dean of game theory, John von Neumann, was criticized in his day for his uber-rational approach to the prospect of nuclear annihilation. He was savagely parodied as the Dr. Strangelove character in the eponymous film, in which the logic of game theory produces the absolute rationalism of a doomsday device that destroys the world. But the shortcomings of these two extremes – absolute rationality and runaway emotionalism – highlight the middle path that game theory can offer: it helps us to explain, in rational terms, phenomena that are themselves at least partly irrational.

Indeed, as game theory increasingly finds its applications in the social sciences, this middle path approach has become not just permissible but essential. Avinash K Dixit and Barry J. Nalebuff's *The Art of the Strategy* makes the case simply: "In the social sciences, multiple causes often coexist, each contributing part of the explanation for the same phenomenon." (Dixit & Nalebuff 2008).

I'm taking pains to justify the sociological approach because I want to be clear on what game theory can and can't do for the science of contemporary marketing. I believe it has very broad *explanatory* power, in its ability to make sense of how social media marketing evolved out of earlier forms of marketing. This analysis reveals a deep structure to the marketer-consumer relationship through examination of the mutually dependent conflicts at its basis. And I believe that game theory has more limited but still very useful *predictive* power, in its ability to guide the choices that we make in marketing experiments. That limitation is, of course, the pure cussedness of human nature, which ensures that we'll never completely remove the risk from marketing, that customers will always surprise us.

## 1.2 Game Theory, the New Media, and the NEW New Media

The purpose of this study, then, is to use game theory to make sense of the rapid changes that have taken place over the last decade in the digital marketing landscape, with a particular eye toward the changes of the last three or four years – what is referred to as "social media" as a general phenomenon and "social media marketing" as a label for marketers' participation in this phenomenon for commercial purposes. Analysis of how we got to this place should provide a tool for marketers and marketing scholars to make better decisions about what games to play and how to play them in this shifting new environment.

I also hope to provide something of an antidote to the breathless accounts of the so-called social media marketing revolution, which have dominated the discussion so far. Influential studies like Forrester Research's *Groundswell* make social media marketing sound like a party train leaving the station, with only two choices available to marketers: chase after the train and hurl themselves on board or get stuck at the station forever. *Groundswell* calls social media marketing "an important, irreversible, completely different way for people to relate to companies and to each other." (Li 2008). By contrast, my study views social media marketing as a tectonic shift in the landscape, but one formed inevitably by converging forces that have taken shape over decades.

More importantly, social media marketing will, for the foreseeable future, exist alongside traditional marketing and the established media models it supports – sometimes as a challenge to it, but often as a complement to it, and no complete analysis can ignore this interplay. Nor do I regard social media marketing as train with a ticket for every passenger; while every brand’s marketing will be affected by the growth of social media, the degree of participation by brands will and should vary. My argument for an evolutionary rather than revolutionary theory is meant to provide a sound basis for evaluating *when* and *how* to apply social media marketing tools, based an analysis of the consumer’s and marketer’s self-interest.

### 1.3 The Payoff Matrix

Take the seminal example of General Motors’ Fastlane blog, persistently cited among social media marketers (including *Groundswell*) as an early indicator of major shifts in brand behavior toward consumers. When the blog launched in 2005, it was by no means a pioneer among corporate blogs – that honor belongs to tech companies like Microsoft and Dell, who preceded GM by several years – but it received attention in part because it came from a company that seemed least likely to take such a gamble. Indeed, GM’s recent near-death experience is generally attributed to the company’s ham-fisted inability to evolve, so a deft blog strategy in 2005 was a surprise to nearly everyone.

But when viewed through the lens of game theory, GM’s Fastlane blog was a perfectly logical move, not the astonishing act of bravura that *Groundswell* sees. Game theory allows us to cut through the hype and evaluate such opportunities based on the mutual dependence of what each player realistically stands to gain or lose. This becomes obvious when we examine the Fastlane strategy with a common game theory tool – the payoff matrix.

A payoff matrix is simply a way to map each player’s stake in a given contest. In its most basic form, it involves win-lose binaries, and in more complex versions, it includes numerical ratings for degrees of payoff. Since marketing almost always involves degrees of success rather than absolutes, overly simplified models won’t do. But for the sake of illustration I will try to reduce the GM blog decision to its core elements without doing too much violence to its true complexity.

Let’s assume that GM’s decision came down to the question of whether to have a consumer blog or not, and that they could rate the potential payoffs on that question on a scale of 1 to 4, with 4 being the highest or

greatest payoff. Then let’s assume that some consumers make a decision about whether to provide feedback – either positive or negative – on GM and its vehicles if given the chance, and that their payoff could be illustrated on the same scale.

For the sake of illustration, let’s further assume that non-response is not part of the equation, i.e., that there is a group of consumers that want to give *some* form of feedback to GM. Of course we can’t make that assumption for all brands confronting the blog decision, and we won’t. In later chapters I’ll explore non-response as a form of response, which is a very real risk for many companies facing decisions about collaborative marketing venues. I feel justified in making that assumption with GM, however; at the time, the company sold 9 million cars a year and was one of the largest companies in the world. There was never much chance that consumers wouldn’t have anything to say about their cars or their corporate policies.

In the matrix below (Table 1), GM’s highest payoff occurs in the upper left quadrant – they build a blog, consumers come there and shower them with praise, and everyone is delighted (Note that a payoff matrix is only concerned with *possible* outcomes, not probabilities). Their next-highest payoff, shown in the upper right quadrant, may be tough to accept at face value: what kind of payoff involves investing in a blog only to have customer gripe at you?

**Table 1:** GM Fastlane payoff matrix

	<i>Customer: Give positive feedback</i>	<i>Customer: Give negative feedback</i>
GM: Build a blog	4-3	3-4
GM: Don’t build a blog	2-1	1-2

I’ll lay out this rationale completely in later chapters; for now consider how the blog’s role as a sounding board helps to position GM as a company that listens to its customers, even when the feedback is critical. Further, if GM sells 9 million cars a year and its blog gets a modest but respectable 5,000 hits a day, the impact of critical remarks on sales is negligible, while the positive PR effect of the blog’s existence and GM’s apparent openness is substantial.

By the same token, if GM simply misses the opportunity to capture positive feedback by not having a blog (lower right quadrant), it wouldn’t affect their consumer reputation one way or the other. And in the lower right, the worst-case scenario, customers have complaints but take them elsewhere, e.g., onto auto review site blogs and message boards, where GM has no

voice, no control, and is subject to the slings and arrows of uncontained negative input.

The customer payoff ranking (shown in italics) is more straightforward: for the customer with something to say, having a direct forum is better than not having one, and being able to leave negative feedback is more satisfying and more empowering than leaving positive feedback.

What's remarkable about this particular payoff matrix is that the nature of the feedback, which is what marketing executives mostly worry about when deploying customer interaction tools, is less important than the existence of the tool itself. GM's payoffs actually shift very little (+/-1 degree) based on the type of feedback, but its payoffs shift substantially (+/-2 degrees) depending on the blog's existence. By "cooperating" and providing a blog for its customers where the good, the bad, and the Pontiac Aztek can be thoroughly hashed over, GM reduced the chances of customer defection into forums where they could do more harm and less good.

GM's blog strategy was therefore less about winning the big game than improving on a losing hand in areas of responsiveness and agility. The impact of the blog was never going to be transformative; if it had been, GM might not be bankrupt and under government control as of this writing. The blog is significant as an opening move in a much more complex game of changing the company's brand perception and its relationship with consumers, and it could never do so in isolation – without traditional advertising, without better cars, and without, as it turns out, massive federal intervention.

But as an evolutionary stage, GM's Fastlane blog was indeed a milestone. It illustrates a game strategy that turns on its ear the long-held notion that marketing is primarily a matter of controlling the message. In 1987's *Roger and Me*, GM CEO Roger Smith dodged documentarian Michael Moore for three years; he clearly felt that controlling the message was his best strategy. How did we go from "As GM goes, so goes the nation" to a marketing forum where GM's Vice Chairman willingly gets taken to task by ordinary consumers over the company's slow progress on hybrids? More importantly, when did the latter scenario start to count as a *success*?

It will take several chapters to answer that question fully. For now it's probably enough to say that it has something to do with the unique features of social media forums. Those features – and their consequences for both consumers and marketers – are the subject of the chapters that follow. My goal is for you to emerge with a better understanding of the games that marketers and consumers are actually playing, based on the structure they have in common with games that have already been played. In doing so, I hope to provide a durable theory of social media marketing that remains useful and

relevant even as the content and structure of social media shifts like sand beneath our feet.

In Chapter 2, I'll introduce the most basic game theory concept, the zero-sum game, and demonstrate how it has been used in traditional direct marketing. I'll explain why zero-sum games rely on informational advantages to succeed, and how those advantages have now been disrupted by the transparency of the social media era. I'll examine banner advertising's near-death-experience as an example of how zero-sum tactics can evolve toward a stable, if unsatisfying, equilibrium of interests between marketers and consumers.

In Chapter 3, I'll introduce game theory's most famous and troublesome concept, the Prisoner's Dilemma. I'll show how the Prisoner's Dilemma helps to illuminate marketers' and consumers' long history of mutual defection while offering real hope for the evolution of cooperation. I'll examine how and why cooperation emerges in iterative cycles of the Prisoner's Dilemma and social media's potential to bring about cooperation.

In Chapter 4, I'll describe the challenges posed by consumer revolt in social media, which actually helps marketing evolve toward cooperation by making marketer defection more costly. I'll posit that paid search marketing constitutes an evolutionary step toward cooperative marketing because it involves consumers indirectly in determining the quality of content and punishing defection. In social media, consumer-generated content like the "United Breaks Guitars" video encourages marketer cooperation by punishing defection with greater consequences than ever before.

Chapter 5 takes up the question of whether marketers and consumers can, under limited circumstances, sustain a mutually rewarding relationship without the use of advertising. I'll begin with the premise that advertising is inherently a sub-optimal arrangement for both players, and that the main challenge in achieving mutual cooperation in social media is the need to coordinate the moves of both players. I'll examine the relevance of coordination games, which attempt to make it safer for both players to seek the richer payoff of a cooperative solution. I'll demonstrate how brands that have succeeded in their use of social media have done so by placing themselves at a calculated risk in order to induce consumer cooperation.

Chapter 6 introduces the concept of "self-command," by which a player deliberately constrains their own actions in order to gain influence over the other player's actions; marketers can use self-command very effectively to bring consumers to the table for collaborative engagement in social media. I'll trace the evolution of self-command in social media from blogging through its maturation in "crowdsourcing," by which brands actively engage consumers in shaping business decisions and even brand identity.

Chapter 7 uses Michael Spence's concept of costly signaling to explain the changing transactional terms of social media marketing. In traditional advertising, brands signal their prominence and worthiness by paying the "costly signal" of access to consumers through major media outlets. Social media is undermining this system and replacing it with a new, popularity-based model of costly signaling, in which a brand's ability to attract and sustain interest, often by grassroots means, determines its success. I'll examine both the pitfalls of this new system and its potential for promoting cooperation.

In Chapter 8, I'll examine the unprecedented degree of control that consumers now wield over brand identity itself. While brand theorists have always claimed that consumers are equal players in the formation of brand identity, in practical terms branding has traditionally been a one-way conversation. Social media conversations about brands have the potential to assert far greater control over brand identity than traditional brand vehicles like advertising, because they take place in the highly influential arena of peer-to-peer relationships.

Chapter 9 warns of the potential for over-saturation of content to bring about the collapse of social media marketing even as it is getting underway. Social media itself competes for the increasingly scarce commodity of consumer attention; social media marketing risks exacerbating this problem by gauging its success in quantitative terms. I'll lay out the parameters for a sustainable approach to social media marketing that doesn't overtax consumer attention, and I'll make the case for my belief that many brands will fail – and *must* fail – at social media marketing.

## Chapter 2: Zero-Sum Games in Traditional Marketing

*ABSTRACT: The most basic game theory concept, the zero-sum game, describes conditions in which each gain by one player produces an equal and corresponding loss for the other. Zero-sum games have limited applicability to marketing, because marketing does produce dividends for both players when the right message reaches the right audience at the right time. But marketers have relied on zero-sum in direct marketing, especially when pricing promotions are involved. The Web has disrupted marketing zero-sum strategies because of the degree of transparency it provides and the corresponding insight that consumers gain into marketing tactics. The shifting of the zero-sum equation – the minimax point – in the consumer’s favor can be seen in the rapid decline of click-through rates in banner advertising.*

I’ll begin with the earliest and most basic of game theory concepts, and the only one to have made its way into everyday speech: the “zero-sum game.” Unfortunately the meaning of zero-sum has suffered in its translation into popular culture; it’s often used to refer to “lose-lose” scenarios in which all parties involved end up worse off for having played the game. In fact, zero-sum games have as much potential to be played to a stable point of equilibrium – where both players have achieved their best possible outcome, given the conflict – as any other type of game. An example from popular culture will help to illustrate this misconception and point us in the right direction.

In the 1983 film *War Games*, a kind of Cold War parable for teens, the sentient supercomputer JOSHUA, on the verge of launching Armageddon, is forced instead by the intrepid teenage hero to play multiple rounds of tic-tac-toe against itself, at supercomputer speed. Obviously, this results in endless stalemates, prompting JOSHUA to achieve the human insight that “The only winning move is not to play.”

Stalemates are common enough in zero-sum games, as we’ll see, but they are by no means inevitable. Tic-tac-toe is indeed an example of a zero-sum game, as are chess, checkers, poker, and virtually any game in which in which the potential gains are a fixed quantity, and all gains or losses are directly at the expense of or to the benefit of the other player(s). The term



“zero-sum” refers not made a lack of gain, but to a condition in which all gains and losses between participants, when summed, will equal zero. In tic-tac-toe, a stalemate means both participants have made their best available moves, and neither has lost; the stalemate is a point of equilibrium, albeit a rather unsatisfying one. Tic-tac-toe is therefore not very useful for learning how to play a full-scale nuclear exchange, since both sides would lose in such an exchange. And unlike nuclear *détente*, zero-sum games provide no dividend for cooperation.

Similarly, zero-sum is too stark of a concept to get us very far in analyzing marketing scenarios, even in traditional marketing. Even the types of marketing that consumers most readily defect from – telemarketing or direct marketing, for example – involve some degree of cooperation. Consumers may choose to ignore most telemarketing calls and toss out most direct mail pieces, but when the right message reaches the right consumer at the right time, it produces a net gain for both the consumer and the marketer. At that point, the game is no longer strictly zero-sum.

So why start with zero-sum? Because it’s essential to understanding the evolution of marketing toward other types of games – ones with cooperative solutions that produce dividends for both players. While zero-sum doesn’t fully explain traditional marketing, it *is* a feature of certain forms of direct response marketing. The fact that those forms of marketing are now on the wane is part of the evolutionary process that marketing is now undergoing, for reasons that this chapter will attempt to make clear.

## 2.1 Zero-Sum Games and the Problem of Transparency

The question of transparency or disclosure, i.e., whether each player knows the other’s strategy, is important in zero-sum games. Strictly speaking, zero-sum games can be played with or without that transparency, but in most cases it has a significant impact on the outcome. In poker, players speak of keeping their cards “close to the chest,” emphasizing the importance of non-disclosure to that particular zero-sum game. The biggest gains in poker are made not by holding the best hand, but by bluffing your opponent – convincing them that your chances of winning a hand are better or worse than they actually are. Imagine a game of poker played with full transparency, with every hand visible to every player. The “luck of the draw” would still be in play, but the game would quickly reach a point of relative stalemate. Players would simply fold whenever a better hand appeared, so that very little money would change hands. Most players would go home with something close to their original stake. As with tic-tac-toe stalemates, such an

outcome would be stable – assuming equal skills, each player would average about the same winnings over the course of many poker nights – but not very much fun.


You can see this transparency-stalemate effect for yourself by playing a game of computer chess and asking your computer-opponent to choose your moves for you. Because the computer chess program simply calculates the best possible move for itself, and then the best possible move for you, it will play itself to a draw. This exercise is illuminating but about as much fun as tic-tac-toe and cards-on-the-table poker.

Risk is a big part of what makes games like poker fun, and the risk comes from our inability to know exactly what our opponents are thinking; correspondingly, the skill necessary to win the game is mostly a matter of deducing what our opponents are thinking. But risk is usually an undesirable element in geo-political games. A key feature of stalemates, after all, is their stability: nobody moves, and nobody gets hurt. The original game theorists understood this, of course, and the history of the Cold War abounds with examples of using transparency to produce a positive stalemate, like a hostage negotiator holding her open hands in plain view when she steps in to negotiate. The entire prevailing system of Mutually Assured Destruction was a kind of zero-sum game played to a stalemate or equilibrium: as long as each side remained convinced that the other possessed an arsenal sufficient to annihilate everyone, a preemptive strike was pointless. When Reagan's proposed "Star Wars" missile defense system threatened to de-stabilize the MAD stalemate by giving the U.S. a defense against a preemptive strike, the Reagan administration offered to allay fears by sharing the technology with the Soviets. (Stabilization was ultimately restored when it became clear that the system wouldn't work to)

Indeed, any change in the 40-year nuclear stand-off that threatened to upset the stalemate had to be accompanied by a strong dose of transparency. When Reagan and Gorbachev agreed to dramatic reductions in both nuclear arsenals, "Trust, but verify" became the mantra of the disarmament process. And the Cold War stand-off most intensely studied by game theorists, the Cuban Missile Crisis, featured a dramatic example of transparency. During negotiations at the height of the stand-off, both Kennedy and Khrushchev took pains to ensure the other that no preemptive strike would occur while dialogue was underway by keeping their bombers in defensive positions. A U2 spy plane out of Alaska inadvertently strayed into Soviet airspace, prompting the Soviets to scramble MiGs to intercept it. A flurry of reassurances changed hands, and the crisis was averted. When Kennedy learned of the incident, he is reported to have lamented, "There's always some SOB who doesn't get the word" (Carlton 1975).

Given the global stakes, game theorists were, not surprisingly, very interested in understanding what conditions led to stability or equilibrium in zero-sum scenarios. One of the bedrock concepts in game theory – really the basis for everything that followed – was Von Neumann’s 1928 publication of *The Minimax Theorem*. Von Neumann showed that a rational player in a zero-sum game *will always pursue the minimax condition, which occurs when they have minimized their maximum possible loss*. Thus in the Cold War, the U.S. mainly pursued the strategy of *minimizing* the chances of the *maximum* loss – nuclear annihilation – rather than pursuing victory through a risky first strike.

The minimax condition is everywhere. We use it in poker when we fold on a bad hand, or in tic-tac-toe when we counter the opening move (X) with a defensive move (O), as in Figure 1. In this familiar sequence, the stalemate outcome is pre-determined after the opening move, assuming each player makes their most rational choice in the moves that follow. Defecting from the minimax strategy not only won’t allow you to win, it’ll cause you to lose.



**Figure 1:** Tic-Tac-Toe

## 2.2 The Zero-Sum of Pricing Strategies

Just as zero-sum is often a feature of board or card games played with fixed quantities, it is often a feature of sales and marketing scenarios involving pricing strategies. The transparency of the Web as a content medium has had a profound effect on pricing strategy, and nowhere is this more evident than in eBay’s global marketplace, where thousands of minimax scenarios get played out every day. As a marketplace, eBay offers unprecedented levels of pricing transparency: as it accumulates transactions, those transactions become a repository of knowledge for future buyers and sellers. Both players know how similar items fared in past auctions and can set their strategies accordingly. The seller will set their minimum price (the “reserve”) based on the minimum amount they’ll accept, stabilized by knowledge of past transactions, and the buyer will behave similarly in setting a maximum bid.

Of course, there are still wildcard scenarios based on scarcity and irrationalism – the same forces we see in other markets. A buyer might be desperate to win an item for sentimental reasons, or because he is a collector, and so he will set his maximum buy much higher than the seller’s minimum sell, and the advantage will go to the seller as the bids escalate. But overall this marketplace that has evolved toward greater pricing stability through transparency: the most prevalent type of transaction is now the “Buy it Now” in which the buyer can skip the bidding process altogether and accept a fixed price. In a “Buy it Now” scenario, a rational seller sets the price not at their bottom line minimum but at their *maximin* – the maximum they believe they can achieve over their minimum price and still attract a buyer.

If the transaction is successful, then the “Buy it Now” price was also the buyer’s *minimax* – the minimum amount they believe they can get away with paying beneath their maximum price. The seller has mitigated the risk of an underheated auction, and the buyer has mitigated the risk of an overheated one. The game is still zero-sum, with the seller’s gain equal to the buyer’s loss, and vice versa.

Let’s say I decide to sell my car on eBay, and I want to get \$10,000 out of it, but I’d settle for \$8,000. I have the option of setting a *minimax* point – a reserve – at that \$8,000 mark, thus *minimizing* my *maximum* loss at \$2,000, while reducing the risk that I won’t sell it at all. I can also set a “Buy it Now” price at my hoped-for \$10,000 mark. That buy-it-now price is my *maximin* point because it helps me to *maximize* my *minimum* gain. I am, in classic zero-sum fashion, hoping for the best while guarding against the worst.

The prospective buyer has a corresponding set of options. By disclosing my *maximin*, I’ve provided the buyer with her own *minimax*: a price that minimizes her maximum outlay while removing the worst-case scenario, i.e., that she’ll miss out on the chance to buy the car. The buyer has the choice of exercising the *minimax* option or bidding up past the reserve, which bears the risk that another buyer will take the “Buy it Now” option in the meantime, or that the final bid will exceed that amount. As the seller, I have an incentive to set a fair *maximin* price, and the buyer has an incentive to accept it as her *minimax*. We’ve reached equilibrium.

My father used to say that a successful negotiation is one in which everyone ends up a little bit disappointed, and that’s not a bad description of the *minimax* theorem. When the buyer’s maximum threshold and the seller’s minimum threshold are the same, a successful transaction occurs, but neither side enjoys the triumphalism of having vanquished their opponent. The trade-off is stability: in a buyer-seller relationship, a vanquished buyer is not a repeat buyer, so equilibrium stabilizes the relationship for future transaction. In marketing, the iterative nature of the exchanges between marketers

and consumers makes finding a point of equilibrium very valuable for both sides, as we'll see repeatedly in this study.

### 2.3 The Wisdom of Randomization

At face value, minimax strategies don't appear to be very useful to the marketer. An effective price-promotion strategy relies on convincing the prospect that the best available deal is the one in front of them; disclosing your bottom line virtually ensures that you're not going to do better than your bottom line. It would also provide the prospect with a lopsided informational advantage – like a poker game with one player's hand displayed and the other hidden – because the marketer cannot fully predict what the prospect will do.

There *are* some simple zero-sum scenarios where your optimal strategy is the same regardless of what you know or don't know about the other player's intended moves; these are somewhat unimaginatively called “no-knowledge” zero sum games, and they include the tic-tac-toe strategy mentioned earlier. But one of Von Neumann's important contributions to our understanding of zero-sum scenarios was to show that a point of equilibrium exists in every such game, regardless of the players' knowledge.

Whether the players can *uncover* that point of equilibrium, based on limited knowledge, is entirely another matter, and that's where things get interesting for marketers. A traditional marketing program like a direct mail campaign, especially one using price promotion tactics, is seeking a minimax point that provides a reliable, measurable rate of response. In this traditional scenario, the prospect is clearly disadvantaged in terms of knowledge, but they do hold one ace in the hole: their willingness to respond. While the marketer cannot know with absolute certainty how much of a discount to provide in order to prompt a response from the prospect, they can reduce their risk by observing the behavior of prospects over time.

To understand how marketers could and should play the zero-sum game when their knowledge of the prospect is limited, let's take a look at a zero-sum example often used by game theorists to illustrate this scenario: the heads-or-tails game. Suppose that Reagan and Gorbachev, instead of relying on a complex negotiation, treaty, and ratification process to reduce their nuclear arsenals, instead decide to settle the matter with a simple game. Each has a stack of quarters representing their respective nuclear arsenals. Each will simultaneously put down one quarter in each round of the game. When both put down heads or both put down tails, Gorbachev wins the round and both quarters. When one puts down heads and the other puts

down tails, Reagan wins the round and both quarters. Their simple payoff matrix, with its zero-sum nature spelled out starkly, looks like this (Table 2):

**Table 2:** Heads or tails payoff matrix

	<i>Reagan: Heads</i>	<i>Reagan: Tails</i>
Gorbachev: Heads	1, -1	-1, 1
Gorbachev: Tails	-1, 1	1, -1

If either Reagan or Gorbachev adopts a “pure” strategy – playing the same way in every round – they’ll quickly lose their quarters as soon as the other recognizes the pattern. So this game features an obvious minimax solution: as long as each randomizes their moves, playing heads and tails at an equal (but non-patterned) rate, the law of probability favors them both equally. In game theory, this is known as a “mixed” strategy. Neither player gains an advantage over the other, and both get to keep their nuclear arsenal. This is also an example of a “no-knowledge” strategy, because both Reagan and Gorbachev could announce their strategy at the start, “I intend to throw down heads or tails completely at random!” without hurting their chances one bit.

The problem with this solution is that human beings aren’t very good at doing things in purely randomized ways; we are naturally inclined toward patterned behavior, based on innate prejudices and preferences. In fact, the game theory scholar Robert Aumann has argued that randomization strategies are not useful applications of game theory, because humans are inherently incapable of acting at random. Aumann’s argument doesn’t undermine the relevance of randomization for the marketer, however, since such patterning is useful to the marketer, as we’ll see.

Returning to the example of Reagan and Gorbachev, we would have to say that on a practical level, both Reagan and Gorbachev will adopt some exploitable pattern of behavior that’s non-random. When this happens, the advantage goes to the better poker player, i.e., the one who first recognizes the patterns – the innate preferences and prejudices – of the other.

Suppose that the canny Gorbachev recognizes that Reagan is slightly favoring heads over tails, perhaps because Reagan unconsciously enjoys gazing at the visage of the father of the nation. Now Gorbachev has the advantage. He no longer has to play at random: he can play heads more than 50% of the time and is guaranteed to increase his winnings, all the way up to 100% (i.e., he plays heads every single time). The optimal strategy for Gorbachev is to exploit Reagan’s weakness for heads *up to, but not beyond*, the point at which Reagan catches on and changes his strategy.

## 2.4 Randomization and A/B Testing

The use of these winning zero-sum tactics by marketers is as old as marketing itself, and much older than game theory. Gorbachev's attempt to uncover Reagan's latent prejudices in order to gain a competitive advantage is the same tactic used in A/B testing – one of the fundamentals of direct response marketing. In an A/B test, the marketer is attempting to gain an informational advantage in a no-knowledge game by randomizing a set of isolated creative variables and observing which ones produce the best response. Every direct mail piece you receive is the product of either direct or cumulative testing of elements like headlines, color, offer, shape, etc.

In a true no-knowledge scenario, the marketer is *always* better off randomizing by testing a broad spectrum of options, because randomization prevents the marketers' own prejudices from excluding possible advantages that may be uncovered through testing. Over time, direct response marketers develop a set of conventions or "best practices" based on accumulated knowledge, but randomization is always a component, or should be.

Why? Because the consumer is playing the same game, trying to minimize their maximum exposure to intrusive marketing messages, and consumers quickly become inured to certain tactics. For instance, direct marketers once held a penchant for the use of Post-It® notes pasted to direct marketing letters, to call attention to certain elements of the letters. Once consumers became accustomed to this tactic, its effectiveness waned. In direct marketing, this waning effect – the point at which response rates reach their apex and begin to decline for a given piece of creative – is often spoken of as "creative exhaustion," but it's not the creative that's exhausted. It's the consumer. The consumer has crossed the saddle point at which they are willing to trade their attention for the marketer's offer, and it's time for a fresh game.

As consumers, we adopt these conditioned responses unconsciously; most of us are probably not aware that we've internalized rules that say, for instance, "Don't be fooled by the handwritten note on the letter." But on the basis of these experiences, we develop broad heuristics – internalized rules – that make us more cautious in the next round.

A/B testing, then, is simply a way to shift the point of equilibrium – the minimax point – toward a more favorable solution for the marketer through the use of randomization and optimization; the game theory equivalent of this practice is called a "mixed strategy." Such a strategy not only prevents the consumer from becoming inured to tactics, it reveals the consumer's own patterns and conditioned responses, allowing the marketer to adjust accordingly. This dual move of blocking the opponent's knowledge of your

moves while gaining knowledge of theirs is quite common in other disciplines.

One of the great innovations in crime prevention in U.S. in the last decade is the growth of community policing, a strategy that involves, among other tactics, the random, visible presence of police officers at various times throughout the city, so that no particular area could be seen as a safe zone for committing crimes. Accompanying this policy is the accumulation of trend data on where crimes are occurring, so that police can target these zones as needed (Beito). The IRS pursued a similar strategy with the random audit, designed to prevent taxpayers from sleuthing out which tax filing practices were likely to trigger an audit; at the same time, the IRS could accumulate data on the true red flags for tax cheats.

In practice, though, few marketers pursue randomization strategies in the zero-sum aspects of their marketing game. Rather, they rely on accumulated and acquired knowledge reified into “best practices.” Accumulated knowledge is a vital component, to be sure, as it allowed Gorbachev to exploit Reagan’s “head” preference, but it would have become a liability for Gorbachev if he didn’t change tactics as soon as Reagan caught on.

Marketers’ weakness in this arena is a natural one, and it can’t be chalked up to a simple lack of awareness. The largest obstacle is the built-in inertia of large organizations. Marketing teams thrive on sure bets, and there is little appetite for trying new tactics when “proven” tactics are readily available. Knowledge of these sure bets is part of the intellectual capital that marketers use to maintain their relevance and hold onto their jobs in competitive organizations. Indeed, this resistance to innovation in zero-sum games is symptomatic of a much broader resistance in the application of new social marketing techniques, as I’ll explore in later chapters.

## **2.5 The Hazards of Entrenchment**

In direct marketing, failure to recognize and break entrenched patterns, i.e., failure to randomize, is deadly. The deadliness, of course, lies in the consumer’s ability to gain exploitable knowledge, such as learning to ignore new direct mail techniques. As consumer knowledge accelerates, so does the deterioration of effective tactics, forcing direct marketers to evolve more rapidly. Let’s look at this effect in action.

Imagine a high-end fitness equipment manufacturer – we’ll call them Manufacturer X – that has always sold its products to consumers through direct channels, never through retailers. Like most direct marketers, this manufacturer has embraced the Web’s direct marketing potential in all of



the obvious ways: interest generated in other channels like DR TV or direct mail is funneled to the call center and/or to the Web site, and online direct-response media like banner advertising and paid search marketing are a key part of the marketing arsenal.

Because this is a high-end, high-consideration piece of equipment, the manufacturer has naturally built some pricing flexibility into their conversion strategy. Most leads don't convert to purchase immediately, so the manufacturer uses incremental offers delivered through direct mail or email to extract maximum value from its leads. The piece of equipment lists for \$2000, but its cost to the manufacturer is \$1000. The manufacturer's follow-up strategy to non-converted leads is to send a \$400 discount offer at a two-month interval, followed by an \$800 discount offer after four months. Over time, the response rate has become predictable. For every 100 leads,

- 10 convert at the \$2000 offer for a \$10,000 profit
- 20 convert at the \$1600 offer for a \$12,000 profit
- 30 convert at the \$1200 offer for a \$6,000 profit


Manufacturer X clearly has a minimax point at the \$1600 level; they can stick to retail pricing and lose incremental leads, or they can discount aggressively and lose margin. Instead they minimize their maximum loss by focusing their energies on making sales at the \$1600 level. In doing so, they've optimized to the customer's maximin point: most customers won't convert at the list price because the product is on the high end of the market, and customers instinctively believe they can do better than the list price. When the customer gets the \$400 discount offer, however, they have ample reason to act: they risk losing the offer altogether on the chance that a better discount may be down the road.

Like many such zero-sum scenarios – poker again is a useful comparison – both players are operating on limited knowledge of the cards the other is holding. The marketer has historical data on customer performance to tell them what the customer is most likely to do, and the customer has performance data, based on past shopping experience, of what the manufacturer is most likely to do. Some cards are showing; others are hidden.

On the surface, this looks like an equilibrium worth maintaining: the marketer is making money on the product, and customers are buying it. The problem is that equilibrium in sales volume generally doesn't sit well with shareholders, and giving them a primer in game theory isn't going to help the situation either. The internal demands of the organization will be to improve on the \$1600 minimax threshold. The burden rests on the marketer to make some magic out of an otherwise static state.

The goal is simple enough: shift a percentage of the \$800 respondents into the \$400 category, and the cost-per-sale drops. But the reliance on incremental discounts, effective though it has been, has painted the manufacturer into a corner. As I've already described, they can and should randomize their tactics, avoiding established patterns of response.

But most marketers take a very narrow view of what it means to randomize (if they even think about it at all). Think of marketing as a pyramid, with everything we can do with messaging – including ad creative, subject lines, etc., occupying a limited span at the top. The media – what channels to advertise in, what mix to use, etc. – occupies a much wider sphere of influence in the center, but the overall strategy – say, price promotion vs. blogging – gives us the widest latitude of all.



**Figure 2:** Hierarchy of marketing factors

What you'll find is that most marketers scurry back and forth across the top of the pyramid, performing nips and tucks on their creative, while leaving the base of the structure unexplored and unimproved. Why? Because that's the nature of equilibrium: safe bets beget more safe bets. True randomization involves risk, and the fitness manufacturer is going to be unwilling to overturn decades of received opinion about price promotions when the fundamental approach remains profitable. After all, they can jiggle the needle on results enough through creative testing – which remains perfectly viable in its own right – to keep the shareholders placated. Throwing open their deeper strategy would introduce risk that, to date, has felt unnecessary.

But while the fitness manufacturer perches unsheltered on the top of the pyramid, storm clouds are forming on the horizon. Remember that limited-knowledge stand-off between marketer and customer? It's gone. This time, when the customer tells the sales representative, "I'll think about it," he doesn't wile away his hours dreaming of that gleaming chrome exercise machine, weakening to the point that the first discount offer through the mail persuades him to make the leap. Instead he hangs up the phone and taps "X-Machine discount" into Google, producing pages and pages of results from deal aggregator sites. Each of these sites cheerfully compiles the deepest available discounts and promotional offers on thousands of products, all

in exchange for a few seconds of the customer's scant attention paid to their advertisers.

It takes the customer less than 5 minutes to find the \$800 offer that would have otherwise reached him 4 months down the road. If the customer is sufficiently interested, he'll steer right past the now-obsolete equilibrium point and buy the product at a deep discount.

This scenario is based on a true story, and it doesn't have a happy ending. Manufacturer X was slow to respond to the changing market conditions and was forced to abandon its price promotion strategy. Sales fell 15% a quarter, and the company made deep job cuts as it struggled to reorganize. Only when faced with imminent catastrophe did the company attempt the kind of randomization of tactics that it ought to have contemplated in sunnier times: it explored retail partnerships, changes to the media mix, different product packages, etc. The company may yet survive, but new marketing programs are not an instant cure in the best of times.

The plight of Manufacturer X, if not their ultimate fate, is a microcosm of the changing marketing landscape – a landscape littered with companies undone by the rapid evolution of consumer behavior. Large, complex organizations don't turn on a dime, so sudden changes to the minimax point in the zero-sum game aren't easily countered. For Manufacturer X and companies like them, there are, very simply, just two available strategies: change the way you play the zero-sum game, or play a different game. Neither is mutually exclusive.

From a game theory perspective, marketers that develop a reputation for discounting inexorably shift the point of equilibrium toward the consumer, because in trying to induce the consumer's immediate action they perversely create a "best is yet to come" mindset that delays the desired action. This mindset writ large produces the macroeconomic cycle known as "deflation," in which consumers delay purchases in anticipation of falling prices.

The fall-off in demand temporarily shifts the equilibrium toward the consumer in the form of deeper discounts, but as manufacturers are forced to make cuts, the economy goes into a downward spiral, and all players suffer. Just as in the case of Kozmo, there is no reason to believe that self-interested players would or should pay more in the short term in order to avoid making a minor negative contribution to the deflationary cycle; the onus is on the manufacturer/marketer to change the game plan.

This reckless patterning of behavior into predictable outcomes occurs not because marketers are lazy, or because they lack information about what works and what doesn't work in a zero-sum scenario, but rather because they have *too much* information, and they are overly reliant on it. Marketers

are victims of their own success in success measurement. Take note of the fact that the advertising media that are the *most* measurable – digital, direct TV, and direct mail – often suffer from a numbing kind of sameness, while the media that are the *least* measurable – broadcast and print – enjoy infinite variety. This occurs because marketers using measurable media develop ideologies about what works, even when it doesn't work for long. Randomization feels risky, even though patterned behavior is *demonstrably* risky.

## 2.6 Making Zero-Sum Work

Zero-sum analysis will show that *any* single direct marketing technique used over a long enough span of time will produce an inexorable shift in the equilibrium point toward the consumer, i.e., the marketer giving up more in terms of cost, impressions, or incentive to drive the same result. But if that's the case, how has the industry even survived? For several reasons, I think:

First, to paraphrase Churchill's famous quote on democracy: direct marketing has been the worst possible way to promote goods and services, except for every other way that has been tried. Prior to the advent of social media marketing, the continuous exhaustion of direct-response techniques was simply the cost of doing business, and it drove continuous innovation – so much so, in fact that a fully mature and complex e-marketing model could evolve in less than a decade.

There is, after all, a base level of effectiveness in every marketing technique – a final minimax point, if you will – simply because *some people want the product*. Banner ad click-through rates may decline asymptotic to zero, but they'll never hit absolute zero, because someone wants the product. What happens instead is that the industry self-corrects – dropping costs, improving targeting, etc. – to make it possible for marketers to stay in the game.

Secondly, the accelerated decline of direct-marketing techniques is a recent phenomenon, brought on primarily by the information glut of the present era, which prompts consumers to tune out marketing noise. Consider that the Web medium alone bombards consumers with over 3 billion advertising impressions *per day*, compared to zero a decade ago.

Thirdly and most importantly, this zero-sum game doesn't *have* to be a race to the bottom. Lots of marketers do well for their clients and their companies without feeling the clammy hand of Career Death on their shoulder, because they innovate in precisely the way that zero-sum analysis teaches us to do: they randomize their tactics continuously in order gain incremental

improvements in their minimax point. The industry may still be sliding, but it can gain some footholds along the way.

## 2.7 Mastering Randomization

Marketers who continue to focus primarily on zero-sum tactics do have some moves, such as randomization, available to improve their position even in the face of a massive shift in consumer knowledge. A case study for randomization arrives in my inbox each morning, in the form of a promotional email from Sierra Trading Post, a cataloguer and e-commerce site for discount outdoor apparel and gear.

STP is a master of randomization. Each and every day is a new discount – on a different product type, in a different formulation, a different shipping incentive, etc. There is no discernible pattern to the discounts, and each is treated as momentous, screaming at me in boldface type. Each promotion lasts only 24 hours, which not only forces me to act on urgency but prevents the lead aggregator sites from over-exposing the discounts.

In general, this randomization is highly effective. The promotions are dramatic and varied enough to compel me to cooperate with some frequency, and when I find a product that I want at a good discount, I'm far less likely to adopt a "best is yet to come" mentality – for all I know, the discount will disappear entirely the next day. STP and I are in a state of equilibrium: I will ignore most of their messages, but I will respond with sufficient frequency to make it worthwhile for both of us. They will not provide me with incremental discounts, as in the case of Manufacturer X, but they'll discount enough items with sufficient variety to hold my attention.

By definition, an equilibrium point occurs when neither player can unilaterally improve their position by defecting. If one truly exists between STP and me, we should be able to map these conditions to a payoff table. Let's posit that my options are to respond to the discount emails or not respond, and STP's options are to continue sending daily discounts, which prompt action but cut into margins, or not send them and rely instead on the occasional sale, as most retailers do. For the sake of clarity, I'll use a simple 1-4 scale to rate our respective options.

**Table 3:** Randomization payoff matrix

	<i>Me: Shop STP</i>	<i>Me: Don't shop STP</i>
STP: Random discount emails	3-3	1-2
STP: No random discount emails	4-1	2-4

This payoff matrix is a good illustration of finding equilibrium at a sub-optimal point. The optimal situation for me (the “4” rating) could only occur if STP offered no sales and I wasted no time shopping there; that scenario costs me no time and no capital. STP’s optimal situation is for me to shop there without the cost to them in time and capital to provide me with discounts. But the optimal point is not stable for either one of us; STP loses if I shop elsewhere, and I lose if STP offers discounts’ and I fail to take advantage and instead pay more somewhere else.

At the point of equilibrium (upper left quadrant), neither of us can improve our lot by defecting unilaterally. I’m better off getting these occasionally annoying emails because of the potential for discounts, and STP is better off sending them because of the potential business. I can’t rationally defect because I can’t be sure what discounts are available outside of this email correspondence; randomization keeps me in check at the point of equilibrium.

The weakness for STP is that they are now in a corner. I have expectations that can’t be reset easily: I expect that I will receive daily discounts, that no one discount is more important than any other (so that I can afford to ignore most of them), and that my relationship with the brand will always be framed by these discounts. STP can prevent the equilibrium point from shifting *further* toward me, but they can’t shift it back while the game is zero-sum. They’ve made discounting the basis of the relationship, and that’s a penurious way to build a brand, as we’ll explore in later chapters. But their randomization strategy could still keep them above water while other outdoor retail brands sink under economic pressures.

## 2.8 Doing Better than Zero-Sum

If tweaking the zero sum game is a precarious strategy at best, what other options are available? To illustrate, let’s return to the example of Manufacturer X and assume they’ve seen the writing on the proverbial wall. Their exclusive reliance on the zero-sum game of direct response left them vulnerable to all the ways their opponent-customer could *defect* – in this

case, by gaining new knowledge and using it to leapfrog the price promotion strategy. What Manufacturer X really needs in the marketing mix is some *cooperation*.

Suppose that Manufacturer X begins parsing their customer list – legions of fans that bought the exercise machine and love it – and they identify 100 people who are passionate about the product, have great success stories on how they lost 30 pounds or brought their blood pressure down, and don't mind telling the world about it. These people are, in the growing parlance of cooperative marketing, “buzz agents,” and they may be worth an army of discount peddlers in this new marketing environment.

Suppose Manufacturer X tries a range of tactics with its buzz agents: giving them incentives to recommend the products to their friends, asking them to blog about their experiences, encouraging them to put the word out in fitness forums, recording testimonials. Some tactics work better than others, but that's how the game is played, and Manufacturer X is now out of the business of relying on received opinions and in the business of finding out what works. They're figuring out how to turn the hyper-informed, hyper-connected customer to their advantage. Their strategy of pursuing the sub-optimal in order to mitigate *long-term* risk, even if it creates a loss of efficiency in the short term, is the essence of the equilibrium concept I'll explore in more detail in the next chapter.

## 2.9 The Cautionary Tale of Banner Click-Through Metrics

There is one final zero-sum scenario worth exploring because it specifically illuminates the general (and generational) shift away from zero-sum marketing as its tactics begin to erode in the face of new consumer behavior. And as the story also illustrates, the catalyst in this shift is the Web, which provides consumers with the transparency and agility to see past traditional direct-response tactics and make different choices.

It all began innocently enough, with a fuzzy rectangular graphic perched atop a Hotwired.com page on October 25, 1994. The world's first banner ad read, “Have you ever clicked your mouse right here? YOU WILL.” With stunning prescience, AT&T had extended to the Web its popular “You Will” campaign, which predicted future consumer technology, into a prediction that users would blindly click on a banner ad that offered nothing specific in return (D'Angelo).

Remarkably, users *did* click, and that first click set Web marketing down a zero-sum path from which it is only now recovering. For nearly a decade,

the click was all that mattered. It was a measurable action that brought the user in direct contact with the offer. In other words, it most closely resembled the zero-sum game of direct mail, with even better measurability. And because banner ads could be switched out easily, the ability to improve the minimax point through randomization was vastly simplified, if often overlooked.

The obvious problem is that banner ads are only *partly* like direct mail. For the most part, direct mail's practical purpose is simply to get consumers to respond. If the consumer throws the envelope unopened in the trash, it accomplishes nothing. But banners could do more. As with print and broadcast advertising, the banner appears alongside free or subsidized consumer content and helps to offset its cost. As in these other media, consumers can absorb a "brand impression" while they focus on other content.

And marketers generally agree, though they may lack the game theory framework to describe it, that a brand impression sits outside of the zero-sum game. Branding is not directly transactional; it demands no immediate action by the consumer, allowing instead for the cumulative impact of repeat exposure. In its purest form, branding is a form of cooperation, inviting the consumer to participate emotionally in defining the product's meaning. The brand marketer seeks a long-term relationship that depends on consumer goodwill in a way that direct response marketing does not.

There'll be more on where branding fits in to game theory later. The point here is that banner advertising stood at those divergent paths from the start, and it took the path more travelled, consigning itself, perhaps forever, to the realm of direct response. The allure was irresistible: here was a medium that offered immediate, highly measurable feedback on its effectiveness, allowing the marketer to track the *actual value* of a given ad and media placement.

If marketers had known how that value would fluctuate, they might have chosen a different path for the medium from the start. Recall the previous axiom that *any* single direct market technique over a long enough span of time will produce an inexorable shift in the equilibrium point toward the consumer. It's also axiomatic that marketers will chase their losses with more aggressive direct response tactics, producing short term gains but ultimately making a bad situation worse.

And that is, in essence, what happened to banner advertising. Fearful of missing out on the next big thing, advertisers threw money at the Web. Publishers, trying to gain dominance quickly in the race to monetize content online, obligingly raised rates. In 1998, advertisers could expect to pay an



average of \$37 for every 1,000 impressions (Morgan Stanley Dean Witter), which was made digestible only by the 1-2% response rates that the ads still commanded.

But from 1998 onward, that response rate slid. To sate advertisers' appetite for impressions, publishers began saturating their content with ads. When Microsoft's car-shopping portal, Carpoint, debuted in 1997, there were no ads on its home page. By 2001, there were at least eight, not including sponsored links and pop-ups. As a matter of simple mathematics – even the most willing user can only click on one ad at a time – click-through rates declined accordingly.

But there were other factors that hastened the decline. The most obvious is the axiomatic one: consumers in a zero-sum game become inured to marketer's tactics over time. Tactics that produced incremental gains quickly become overused dogma, whereupon they become ineffective. Because advertisers now had to compete for eyeballs in much bigger arenas, their methods became increasingly intrusive and deceptive: strobing ads, fake interfaces, and ads camouflaged as real content.

The most notorious example, still spoken of ruefully among Web marketers, is Treeloot.com's "PUNCH THE MONKEY AND WIN 20 BUCKS" ad, which invited the user to brandish a virtual boxing glove to punch a virtual monkey. Millions of users were duped into clicking, only to discover that they'd won 20 "banana bucks" that could be parlayed into real money only by playing even more games. The ad was so often decried by the industry's doomsayers that some still hold it accountable for the near-death of the medium.

The truly tragic aspect of the direction that Web advertising went is that marketers saw the writing on the wall very quickly. From its debut in 1999, the Web marketing forum *Clickz* began fretting about the industry's over-emphasis on direct response, believing it would lead to a crash. Topics covered the first year included "Escaping the Cult of the Click-Throughs" (Graham 1999), "Tracking Non-Click Conversions," and "Between a Rock and a Hard Place," which contained the quaint observation that click-through rates were "at an all-time low" (Hespos 1999). (The average response has since declined another 500%.)

It's easy to be smug about the inevitable consequences of the new medium's direct-response myopia, but in truth individual marketers were simply powerless to invert the widely accepted perception that banner advertising's primary function was as a direct response medium. The industry produced study after study showing how exposure to banner ads increased brand awareness by some measurable delta. The Internet Advertising Bureau was formed mainly to advance that agenda, by standardizing ad sizes around

more brand-friendly specifications and running studies on the impact of rich media.

Certainly the evidence was persuasive, but it didn't matter, because of another axiom: given the choice between hard and soft data, marketers will always choose hard. So unless the entire industry simultaneously stopped measuring click-throughs, it remained the only metric universally accepted as an indicator of campaign performance.

Then the crash came. Advertisers were more or less content to throw bad money after good in banner advertising as long as the Internet economy was strong. But when dot-coms started to bomb with greater intensity in late 2000, dragging the rest of the economy with them, online ad money dried up overnight. Start-up online media companies canceled IPOs, and public ones like rivals Avenue A and Doubleclick watched their value vanish. The mainstream media wasted no time in declaring the era of online advertising well over, and the Web's ad volume shrank for the first time since its inception. It remained in decline for nearly two years.

In retrospect, it seems unfair that Web marketing was sent into the desert like a scapegoat, carrying marketers' sins on its back. To this day Web marketers still complain, and quite justifiably, that the level of accountability between online and offline advertising is badly misaligned. We still argue about brand impact and still tout statistics to persuade advertisers to accept other metrics.

But none of that really matters when we look at this story through the coolly objective eyes of the game theorist. Web advertising went the zero-sum route, and zero-sum is what it got. Its zero-sum mathematics went the only direction such mathematics can: the minimax point shifted toward the consumer. But it's also true in game theory that that which does not kill us helps us find equilibrium, and that's what happened here.

Interestingly, at least one business journalist observed the relationship between game theory and banner advertising's race to the bottom early on. In a piece for *Business World* entitled, "The Unbearable Lightness of Ad Revenue," Frank Yu declared, "Ad budgets are a zero-sum game and so are users' attention spans." He predicted that as "jaded, cynical consumers" learned to tune ads out, only the top content providers could afford to stay in the game, and severe "clustering" of content and media revenue would occur. He further predicted that new platforms like PDAs would challenge the Web and force new content monetization models (Yu).

Yu was at least partly prescient, if too cynical. Web traffic did indeed cluster around top content providers, but smaller players were able to stay in the game as a result of the Web's transparency. Media planning tools like Nielsen Online (formerly Nielsen NetRatings) were able to ascertain the

dimensions of the audience on more niche sites and allow advertisers to trade volume for relevance.

The predicted changes brought on by new platforms are only now beginning to occur, with marketers taking notice of the growth of mobile applications as a small but rising threat to the now-traditional online advertising model. But the fundamental problem Wu raises – that of consumers tuning out – remains the industry’s greatest challenge.

What truly saved Web advertising was the equilibrium that occurred between response rates and media costs. While the minimax point shifted inexorably toward the once-bitten-twice-shy consumer during this period, the industry survived because the cost model shifted too. The cost has stabilized around a proportional rate of return that direct-response marketers can live with; in other words, the cost of impressions dropped alongside the rate of response. This has, in turn, eradicated most of the least tolerable tactics. Pop-under ads are largely a thing of the past, and fake interactions are mostly passé.

The limitations of this outcome are the same as they are for Sierra Trading Post: a more stable zero-sum game is still a zero-sum game. It leaves marketers with the basic problem of trying to eke out performance gains from a medium that is shifting inexorably away from direct consumer engagement. The stark reality of this marketer-consumer relationship was made plain by a 2007 study that sent shock waves through the digital marketing community. A joint study by media research company Comscore and media agency Starcom showed that a stunning 50% of all clicks on banner ads came from one small slice of the Web population: Web users aged 25-44 with a household income of less than \$40,000 per year. Dubbed “Natural Born Clickers,” these users spend four times more time online than average users but purchase products at significantly lower frequency.

Such users tend to favor gambling, employment, and auction sites – a much narrower pattern of surfing behavior than the Web population as a whole. A 2009 update to the study showed that the minimax point was continuing to slide. The percentage of monthly clickers fell from 32 percent in July 2007 to 16 percent in March 2009, with only 8% of Web users accounting for 85% of clicks (Comscore 2009).

From a game theory perspective, the implication of the “Natural Born Clickers” phenomenon is that it undermines the precarious equilibrium in click-based banner advertising. That equilibrium is based on the idea that the cost of finding and prompting action from the right targets compensates for banner advertising’s low response rate. If, however, that low rate of response also falls short of finding the right targets, the advertiser is no

longer in equilibrium. Advertisers are then paying too much for the wrong kind of results.

Obviously the industry is in need of a game-changer – a shift in the use of the medium that moves it outside of the stark give-and-take of zero-sum. Fortunately for the banner ad medium, that game-changer has come in the form of more advanced metrics that account for the effects of advertising beyond direct response. Any of us can recall an instance of having seen an ad or a series of ads and having some later decision, e.g., which cars to research, informed by those previous impressions.

This is, in fact, the way that advertising has always been understood to work: as one of many factors that add up to a purchase decision. Banner advertising, by contrast, had been operating under the fallacy that only a direct and immediate action, irrespective of whatever else the user might be doing, is the only way to account for the ad’s impact. Such an outrageous supposition easily leads to the Natural Born Clickers phenomenon, as clicking on an ad bears the lowest cost for a user who is at their leisure and has no intention of purchasing.

But the advent of advanced metrics disposes of this fallacy. Advertisers can now account for “view-throughs” of an ad, i.e., the perfectly natural phenomenon of a user seeing an ad and responding later. In rich media advertising, one can now account for interaction with the ad – certainly important in making a brand impression – as well as the brand impact of the ad. And banner advertising can be evaluated for its contribution to sales rather than to the fallacious clicks metric.

The digital marketer might rightfully protest that no other advertising medium is required to justify its existence in this way; it is the equivalent of demanding that billboard advertising account for consumers that spotted the sign and then later went to the store and purchased the advertised item. But again, game theory provides a ready explanation: once the payoffs in a game have been established, no single player can unilaterally change the rules. No bottom-line focused marketer wishes to give up hard metrics in favor of more logically persuasive but softer arguments concerning brand impact.

This is precisely why the advent of social media marketing is so important to the health of digital marketing as a whole: it provides the game-changer that demands different metrics, none of them easily obtainable, for how online conversations with consumers impact brand relationships. When viewed *in the context of* (as opposed to in conflict with) now-traditional tactics like banner advertising, social media marketing becomes a way of continuing a conversation that may be initiated in traditional ways.

How precisely social media marketing works in symbiosis with other forms of advertising is a topic for a later chapter. The main point of

recounting banner advertising's tumultuous journey is that its evolution away from direct response and toward a more nuanced role has led the way for more radical evolutionary stages represented by social media. And that evolution is reflected in the numbers: while marketers' investment in banner advertising dipped, then stabilized, at a fraction of its former value, their *total* investment in the Web has grown year over year.

This has occurred because interactive media has begun, albeit slowly and with no shortage of false starts, to offer a way out of the zero-sum game of direct-response marketing. The chapters that follow will demonstrate how zero-sum has evolved into more complex gaming scenarios that involve varying degrees of cooperation. These games offer an alternative to the uneasy truce of mutually assured destruction and pave the way toward a very different future for both players.

## Chapter 3: The Prisoner's Dilemma and the Emergence of Cooperation

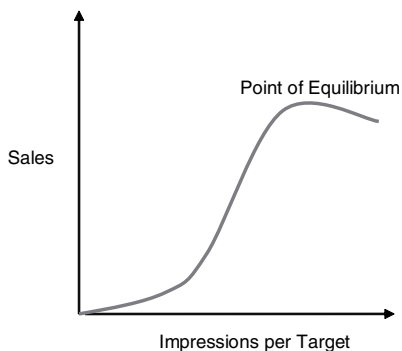
*ABSTRACT: John Nash's concept of equilibrium demonstrates how games may have sub-optimal solutions that are nevertheless stable, because neither player can improve their condition unilaterally. Advertising is inherently a sub-optimal condition for both players: marketers would prefer to win consumers without spending money on advertising, and consumers would prefer to enjoy content without being advertised to. The concept of the Prisoner's Dilemma illuminates this condition, because it suggests that consumers and marketers could reach a more satisfying relationship if they could coordinate cooperation. Studies conducted on iterative rounds of the Prisoner's Dilemma demonstrate this natural evolution toward cooperation and reveal a set of conditions that make cooperation possible, but they also demonstrate the fragility of cooperation and the potential for downward spirals of mutual defection.*

Just as banner advertising could not evolve into a more nuanced medium until its near-death experience forced it to do so, marketing in general does not willingly forego its short-term gains in favor of long-term stability. Put simply, restraint does not occur in the absence of consequence, and so the advertiser will push their advantage in reaching a set of consumers until that negative consequence is achieved in the form of diminished returns.

This tendency may be dramatized as greed or blindness in the face of consumer resentment, but in fact it is perfectly logical and explainable within the terms offered by game theory. As noted before, consumers would find it optimal to find good products and enjoy free media content without being marketed to at all, and marketers would find it optimal for consumers to choose the marketer's products over others without the need to spend a single dollar on advertising. But both sides compromise in the interest of achieving their goals and find a point of equilibrium. In this important respect, *advertising is always the pursuit of the sub-optimal*, a means to an

end for both players. For the consumer it is a Faustian bargain, which puts the marketer in the unfortunate role of Mephistopheles.

The natural tendency of each side to push their respective advantage is illustrated in the traditional “S” curve, by which the effectiveness of a given media spend is often evaluated and optimized. The S curve indicates the impact on sales or some other success metric (the “Y” axis) of a given number of media impressions (the “X” axis). The upper arc of the curve indicates the point of equilibrium – the optimal number of impressions necessary to achieve the best possible sales outcome. The effectiveness of further impressions is diminished beyond that point.



**Figure 3:** Marketing efficiency “S” curve

Thus the marketer has a built-in incentive to pursue the maximum number of impressions possible, provided they produce an incremental return.

In greatly simplified terms, the marketing industry as a whole operates as one enormous “S” curve, pursuing an advantage to its furthest logical point. And since that industry doesn’t operate as a single entity but rather as a vast array of independent players, there is no collective incentive to change course or to sacrifice short-term gains for long-term health.

There are, in fact, specific conditions under which a group of independent players in a game will make such sacrifices; these will be explored in Chapter Five’s examination of the coordination game. For now I will stipulate that these conditions do not presently exist in marketing to any widespread degree. If they did, marketers would not find themselves trapped in a dialectic of hype and backlash. This dialectic occurs because of the instability in the marketer-consumer equilibrium, in which marketers continually press their advantages to compensate for consumers’ increasing disdain for their messages.

This equilibrium has always been delicate at best; it means that advertising impressions and response are in a symbiotic balance, with consumers

tolerating enough advertising to grant them access to free or subsidized content, and marketers are gaining sufficient return from this advertising to limit the number of impressions they impose on consumers. Upsetting the equilibrium would cost the marketer more money and expose the consumer to more advertising – a less optimal outcome for both. Neither has an incentive to defect unilaterally.

### 3.1 The Great Consumer Opt-Out

But like all equilibria in game theory, this one is upset by shifts in each player's knowledge and opportunity. Fully analyzing these shifts would require deeper historicizing of the chicken-and-egg relationship between marketing saturation and consumer disdain than this book will attempt, but suffice it to say that marketers have, in the last two decades, gained vastly greater access to advertising opportunities than in previous eras, and this has upset the balance. We often hear that consumers in the U.S. are exposed to more than 3,000 ad messages per day (Taylor), but the real number, accounting for logo and label exposure, product placement, etc., is probably several times higher. Many of these opportunities are newly minted: not only Web advertising, but commercial email, naming rights, product placement, mobile advertising, and on and on.

The result of this act of defection is a corresponding defection on the part of consumers, with increasing ad tune-out and opt-out. The explosion in non-traditional means of advertising like product placement is, in part, a direct result of marketers seeking alternatives to television advertising, which has been severely compromised by the growth of Tivo and other ad-skipping technologies. One study showed that 90% of consumers that can skip television ads do so, an act of defection made possible by a shift in opportunity – a technological one, in this case – in the consumer's favor (Pasik). And so we find ourselves in a downward spiral of defection, with ad exposure and ad tune-out accelerating at a corresponding rate.

The degree of consumer inurnment to ad messages alone is sufficient to demand a shift in strategy for marketers away from the zero-sum game. To cite a few examples: A study by the Stanford Poynter Institute in 2000, a mere six years after banner advertising's debut, uncovered the phenomenon of "banner blindness," whereby Web users develop the ability to tune out advertising on Web pages they were viewing. Participants in the study saw banner ads only 45% of the time, with an average attention of only 1 second. Since banner ads typically take several seconds to deliver a message, investing in this level of attention is a bad deal for marketers, to say the least.



This tendency has only worsened over time: Web usability expert Jakob Nielsen has conducted multiple banner blindness studies since 2000, using heat-tracking technology to record users' eye movements, and has reached the grim conclusion that "Users almost never look at anything that looks like an advertisement, whether or not it's actually an ad" (Nielsen).

Email is similarly besieged by indifference. Arguably this has much to do with the increase in illegal spam email, which cannot be attributed to bad behavior by legitimate marketers, but the net impact is the same. The anti-spam company Postini reported in 2008 that 94% of all email was spam, with the rate of spam increasing by 1.2% *per day* (Keizer). This has contributed to the discrediting of legitimate emailers, with only 20% of consumers saying they trust email that they've opted in to receive.

In this downward spiral, both players are attempting to regain their position by unilateral defection: the marketer by increasing the number of impressions or ad exposures, and the consumer by decreasing attention. This is a poor strategy for both. The marketer's credibility further erodes, and the consumer is merely advertised to in more and increasingly pernicious ways in order to compensate for the loss of attention. It's also worth pointing out that the consumer's defection also harms their ability to enjoy sponsored content, in a variety of meaningful ways. Some examples:

- Newspapers facing declining revenues reduce their national and global coverage and cut back investigative reporting, thus producing less useful content
- Television programming increasingly relies on cheaper formats like reality programs, reducing the variety and quality of content
- Programming across channels that fails to attract an immediate audience is shelved more quickly, further reducing content variety

When mapped to a payoff matrix, this mutually unsatisfying arrangement becomes obvious (Table 4). The downward spiral is represented by the upper right quadrant, with the marketer increasing their spending, the consumer ignoring the messaging, and the content suffering for it. For the consumer, responding to the ads is marginally better, because doing so increases the health of sponsored content and may prompt cooperation – in the form of decreased exposure – from the marketer. Correspondingly, the optimal payoff for the marketer involves spending less on marketing but getting greater response.

**Table 4:** Advertising response payoff matrix

	<i>Consumer: Respond to advertising</i>	<i>Consumer: Ignore advertising</i>
Marketer: Increase exposure	3-2	1-1
Marketer: Decrease exposure	4-4	2-3

It scarcely requires pointing out that such behavior by both parties – consumers voluntarily responding to more ads and marketers voluntarily decreasing spending in order to get on better terms with each other – is patently absurd. But that is precisely why this game has two points of equilibrium – the worst case scenario in the upper right, and the best case in the lower left. Recall that a point of equilibrium is not defined as the optimal solution; it is merely a point at which neither player can improve their position by acting unilaterally. In the downward spiral, both players are defecting, but unilateral cooperation accomplishes nothing. It is absurd for the marketer to think that lower ad exposure will unilaterally produce a higher response, and it is equally absurd for the consumer to think that unilaterally responding to more ads will ease the ad bombardment or improve content quality. And so the downward spiral continues.

But the game has another point of equilibrium. In the lower left quadrant, the marketer and the consumer find symbiosis. The marketer puts fewer, more relevant messages in front of the consumer, and the consumer responds more frequently. Content quality improves without the need for heavier sponsorship. Neither side has an incentive to defect unilaterally, because all such moves produce sub-optimal solutions.

Getting to this point requires a degree of cooperation that can't occur in the absence of outside factors. And this is precisely where social media enters the picture as the outside factor with the power to change the game. While it has become fashionable to speak of social media as a fast-rising groundswell that began to take shape around 2007, it is more accurate to regard it as the culmination of forces that have been endemic to the Web from the beginning. For consumers, its early manifestations were in tools like product ratings & reviews, message boards, chats, and newsgroups, all of which had been in place for more than a decade. These foundational elements were in fact essential to social media's more recent explosion as a full-blown cultural phenomenon, because they conditioned consumer behavior and expectations to recognize that greater participation and transparency were available to them than ever before.

### 3.2 The Shifting Ground of Consumer Trust

The key element in social media is engagement, and to the same extent that consumers were disengaging in advertising in advertising, they were learning to engage each other. The global PR firm Edelman has been conducting an annual "Trust Barometer" survey to gauge consumer trust in various information sources for the last 9 years, and their 2006 report offered a startling finding: Trust in "a person like me" had risen from only 20% in 2003 to 68%, surpassing all other sources, including doctors and academics. The company's CEO, Richard Edelman, issued the prophetic admonition that "companies need to move away from sole reliance on top-down messages delivered to elites toward fostering peer-to-peer dialogue among consumers and employees, activating a company's most credible advocates" (Edelman).

How could something as fundamental as trust shift so dramatically in only 3 years? Loss of trust in traditional media sources was certainly a contributor; the same Barometer report noted that trust in television as a first source of trustworthy information had fallen from 39% to 29% in two years. But this loss does not automatically privilege trust in "a person like me." Television, after all, is ubiquitous, and like-minded peers with expertise on specific subjects are hard to come by. Or at least they use to be.

Significantly, the same report noted that 34% of consumers take action against a distrusted company by sharing "negative company opinions/experiences online." Therein lies the sea change in behavior that begat a sea change in trust. Consumers gradually, and now naturally, gained easy access to the means of providing feedback online, and their fellow consumers grew to trust these new media as information sources.

Suppose, for instance, that I am planning a trip to Honolulu back in 1996, and I want to find a good beachfront resort. I can rely on travel guides, which will offer me only one perspective on a given resort. I can look at travel magazines, which may contain content about the resorts – sponsored by the same resorts. And I can obtain brochures from the resorts themselves. It's unlikely that I'd be able to gather enough knowledge about a given resort – or even about Honolulu – from friends and family to be able to count on this peer group as an information source at all.

But a decade later, all of these sources are vastly overshadowed – if they are even consulted at all – by a single Web site: TripAdvisor.com, which can offer me dozens of meticulous, detailed opinions and ratings on each resort in Honolulu. I also have access to Honolulu message boards containing peer advisors with local expertise. These opinions are unsolicited, unpaid for, and frank, and most significantly, *there are 20 million of them*. In a keystone,

my trust and reliance on less available, less thorough, and less objective sources plummets.

My point is not that TripAdvisor offers something unique in its subversion of the traditional means of information-gathering about a potential purchase; indeed, it is only the most prominent of many examples of travel review sites. It is the subversion itself that is the game-changer. We can all agree without much need for further analysis that unpaid, unsolicited first-hand accounts of a destination or product are going to be more reliable and therefore more valuable to the consumer than advertiser-supplied information. A single opinion might be too anecdotal to trust, but a dozen opinions have serious weight. So it follows that the ubiquitous availability of such information changes the advertising game in two important ways:

1. In zero-sum games, it provides an informational advantage to the consumer that shifts the equilibrium point in their favor, e.g., on product pricing and discounts, because it provides them with knowledge of the marketer's moves.
2. In non-zero-sum games, which will be the subject of this chapter, consumers' traditional reliance on marketers for product information has been vastly reduced or eliminated, so advertising itself is further diminished in value.

The importance of this latter point cannot be overstated. The essence of the game is that the consumer and marketer are *mutually dependent* adversaries, but the removal of one key area of dependence – if the consumer wanted to learn about the product, they had to hear from advertisers – has radically shifted the game in favor of the consumer.

### **3.3 The Marketer's Loss of the Informational Advantage**

To understand the implications of this, let's return to the example of TripAdvisor and contemplate the dilemma from the point of view of a Honolulu hotel marketer facing a spate of mixed or poor reviews.

To begin with, as the marker I cannot counter this negative information simply by choosing a different playing field, i.e., by reaching the consumer by other means. It is nearly inevitable that the interested consumer will at some point take to the Web to learn more about my hotel. If I'm very concerned about my reviews, I can counter-weight this tendency by focusing on consumers that are less likely to read online reviews – senior citizens, for instance. (Music labels have followed a similar strategy in promoting artists

whose demographic appeal makes their music less likely to be illegally downloaded.) But trading a wide audience for a narrow one still constitutes a shift in equilibrium to the marketer's disadvantage.

By the same token, I cannot simply outspend the problem, because I'm not operating on a playing field that makes that option viable. In the so-called "Cola Wars" of the 1980's, Pepsi-Cola and Coca-Cola played the game in the advertising arena and maintained an equilibrium, in effect, as two marketers posting entirely biased but entertaining reviews of each other's products; the consumer's role was secondary and passive. In this new paradigm, the consumer's role is primary. If my competitors achieve better consumer reviews, then overspending on marketing will not solve the problem, and I am potentially made *more* vulnerable by expending resources in a futile effort.

I could also try to counteract the results of negative reviews by discounting. If my hotel is overpriced at \$200 a night, perhaps it is a bargain at \$100 a night. But this obviously constitutes a shift in equilibrium by the zero-sum terms defined in the last chapter, and it does nothing to counteract the negativity directly. Plus I lose money.

Even if the consumer does a direct search for my hotel, outside opinions are inescapable. (Try this exercise with any hotel. A search for "Ritz Carlton Cancun" on Google returns the hotel's Web site as the top result. The #2 result? The Ritz Carlton Cancun review page on TripAdvisor). I might try advertising on the TripAdvisor site, and if my reviews are good, this would be a highly effective strategy; consumers could go directly from interest to action by reading the reviews then clicking the ad to visit my site. But if the reviews are uniformly good, I never had much to worry about in the first place, and the fundamental alteration remains the same: I have no course of action available to me in traditional marketing to counter the consumer's information-gathering move.

The conclusion is inescapable: as a hotel marketer I am going to have to move outside of my traditional marketing comfort zone and take on the issue of negative reviews head-on. This is going to become as integral to my hotel's marketing strategy as my logo, and it isn't going to be easy. I will work with customer service to respond to negative reviews and elicit positive testimonials. I will coordinate customer care initiatives to address the issues that led to bad reviews in the first place. I will send emails and mailers to past guests soliciting positive responses and surveying their experience. I will tirelessly monitor, catalogue, and respond to customer reactions to my hotel, and I will track the uptick in positivity.

It is altogether obvious that I am describing a very different kind of marketing activity than the ones marketers are accustomed to, and certainly

different from zero-sum. The specific tactics used in this new marketing game will be explored in far greater detail in later chapters; my primary purpose in offering the hotelier example here is to underscore the previously-described limitations of zero-sum and set the stage for a different game theory concept that will help us analyze the challenges faced by the hotelier, and indeed, all marketers in this new era.

As noted, the basic limitation of zero-sum is that presupposes one player's direct gain is the other player's direct loss, and even traditional marketing doesn't usually work that way. Marketers and consumers are mutually dependent because consumers want products and marketers want to sell them, and advertising the right product to the right consumer at the right time constitutes a mutual gain for both players. But as has also been previously described, this entire arrangement has an equilibrium solution that is *always sub-optimal*, always compromised. The optimal solution for marketers would be to sell their products over other products with no investment in advertising, and the optimal solution for the consumer would be to select the right product over other products (and to be able to enjoy free media content) without being exposed to advertising. Neither of these optimal outcomes has long-term viability, and so consumers and marketers do their endless dance.

### 3.4 The Prisoner's Dilemma

The most famous of all game theory concepts, the Prisoner's Dilemma, is precisely what this situation demands, because it is most often used to analyze the conditions of cooperation and defection in situations involving sub-optimal solutions. As a "dilemma," it is much more complex and nuanced than zero-sum, and so it is worth exploring at some length.

The Prisoner's Dilemma has its origins in a paper produced by a pair of game theory's original practitioners, Merrill Flood and Marvin Dresher, in 1949. The paper offered a set of real-life scenarios that were meant to explore the limits of the Nash equilibrium (Dixit & Nalebuff 2008). You'll recall that the Nash equilibrium posits that for every two-person game, there is at least one stable equilibrium point in which *neither player can improve their outcome unilaterally*, given the moves available to the other player. As we saw in the example of Manufacturer X, the point of equilibrium can shift if one player gains information about the other's available moves, but those *information gains always favor one player over the other*. If the consumer learns that Manufacturer X is willing to sell the product for \$200 less than they were prepared to pay, that insight helps the consumer and hurts the

manufacturer. For the traditional hotel advertiser facing poor reviews on TripAdvisor, the availability of those reviews hurts the advertiser and helps the consumer.

The unique feature of the Prisoner's Dilemma is that *both* players accept a sub-optimal solution based on the information they have, but *both* could have arrived at an optimal solution had they been able to coordinate their moves. This potentially turns Nash's equilibrium on its ear, because it uncovers a basic instability: it suggests that in some scenarios, coordinating information-sharing could produce better outcomes for both players. Doing so can be tricky, but the rewards may be worth it.

To understand this, let's look at the dilemma itself. Examples abound, but the most famous is the dilemma's namesake, articulated by game theorist Albert Tucker. Imagine two criminal conspirators are arrested for a robbery. The two are separated by the police, and in a set-up familiar to anyone who has seen a prime-time police drama, each prisoner is invited to implicate the other in order to receive a lighter sentence (Tucker 1983).

The game presupposes that the police need one of the players to implicate the other in order to get a conviction on all charges; if both players stay silent, both players will receive only a one-year sentence for lesser charges. But the players are separated; they cannot coordinate their actions. If one player implicates the other while the other stays silent, the silent one will get a harsh sentence – 6 years – while the other player goes free. If both players implicate the other, no ringleader is established, and both receive lighter sentences of 3 years each.

These stark options are easy to articulate in a payoff matrix if we reverse the polarity of the numbers so they refer to the years in a prison sentence; a zero means no prison sentence, and so on.

**Table 5:** Prisoner's Dilemma payoff matrix

	<i>Prisoner 2: Stays silent</i>	<i>Prisoner 2: Betrays</i>
Prisoner 1: Stay silent	1-1	6-0
Prisoner 1: Betrays	0-6	3-3

According to the payoff matrix, both players are better off defecting (betraying), no matter what the other player does, because they have no way of coordinating what the other player will do. Betraying will result in either no prison time or three years, whereas staying silent carries the risk of 6 years behind bars. In this case, defecting is a dominant strategy for both players

and produces a Nash equilibrium in the lower right quadrant (3-3), because neither player can unilaterally improve on this position.

But the dilemma is a true dilemma for several reasons. First, information-sharing (coordination) would produce a better solution for both players (the 1-1 outcome in the upper left quadrant), which runs contrary to Nash's theorem. If only the prisoner's could signal each other, their lot would be *vastly* improved!

Secondly, and more importantly for our purposes, when this dilemma is taken out of the laboratory and viewed through the lens of human emotion rather than pure logic, the sub-optimal solution is extremely unsatisfying. Game theory teaches us to "Hope for the best, but prepare for the worst" (Dixit & Nalebuff 2008), but in real-life situations, hope often seems to trump preparation. If there is any honor at all among thieves, their sense of morality – typically excluded from game theory analysis – would rebel against betraying their partner in crime. And then there is the starkness of it all – a very light one-year sentence is such a vast improvement over a 3-year sentence that one can hardly bear to imagine that one's fellow thief would fail to reach the same conclusion and unilaterally cooperate.

It should be no great surprise, then, that reckless pursuit of the optimal instead of the more stable sub-optimal is often what we see when the Prisoner's Dilemma is played outside of the lab. In the real-world environment, we find two different PD scenarios at work, each demanding different strategies: one-off games and the more common and important "Iterated Prisoner's Dilemma."

### 3.5 The Hidden Allure of Cooperation

Pure logic and conventional wisdom dictate that in one-off games of Prisoner's Dilemma – as in the example above – *it is always better to defect*. That is the dominant strategy. It follows, then that we would see this result borne out in the countless versions of PD staged by academics over the decades.

But in fact we see the opposite. Dixit & Nalebuff report in that in aggregate across one-off PD experiments, cooperation occurs almost half the time, "even when each pair of players meets only once" (Dixit & Nalebuff 2008). They offer the fascinating example of the TV game show, *Friend or Foe*, in which players competing for a pot of money were simply required to secretly write down "Friend" or "Foe" to indicate their move. If one player cooperated and the other defected, the defector got the whole pot. If both cooperated, they split the pot. If both defected, they got nothing.



Applying the strict logic of “preparing for the worst,” it is quite obvious that defecting – choosing “foe” – is the better strategy in a single round, because you’ll either end up with the full pot or nothing. Choosing “friend” nets you half the pot or nothing, and of course you have no way of knowing if your opponent wishes to cooperate. Yet Dixit & Nalebuff report that almost half of the contestants chose “friend,” preparing for the best instead. When I posed this dilemma to my 8-year-old son, he immediately chose “friend,” and I was torn between admiration for his altruistic tendencies and chagrin at his hasty reasoning. (Attempting to iterate only seemed to annoy him). Or perhaps he simply wasn’t paying attention to the game, which may be the most common scenario for marketers too.

The result is clear: we tend toward cooperation even when the odds are stacked against us. But why? The answer necessarily lies outside of an ultra-rational application of game theory, since it seems to involve either a degree of selflessness or some different transaction – hidden terms in the game, as it were. Such an answer would also help us explain why travelers visiting TripAdvisor would devote so much time – 20 million reviews and counting! – to sharing their experiences with other travelers, with no tangible reward at stake.

This psychological basis for cooperation will be explored thoroughly in the next chapter, when we dig into examples of cooperation in the social media arena. For now, let’s also consider how and why cooperation emerges in iterated games of Prisoner’s Dilemma. The difference is important, because while cooperation may occur in one-off games as the result of altruism, blind hope, or some other emotional cause, it is often a rational response in iterated PD.

### 3.6 The Iterated Prisoner’s Dilemma

Iterated PD occurs when two players face each other in consecutive matches, with the outcome decided by the cumulative score. Most real-life dilemmas are in fact iterative; the U.S. and the Soviets negotiated many arms agreements, and marketers and consumers square off thousands upon thousands of times. It leads us to the question of what strategy is best *in the long run* rather than in short-term self-interest.

The interesting feature of iterated PD is that it brings in punishment or retaliation as a feature; one might reconsider the logic of defection if it makes your opponent more likely to retaliate in the next round. The theory’s originators, Flood and Dresher, uncovered this insight when they staged an iterative PD game among their game theory colleagues, featuring the same

two players squaring off 100 times in a single session. The session revealed a “difficult struggle to secure mutual cooperation,” (Poundstone 1992) which was logical, since both players stood to gain more if they could coordinate their actions and bring about iterative cooperation. When one player defected, the other would punish him with a defection in the next round, and both would return to cooperating thereafter.

In all, the supposed Nash equilibrium – mutual defection – occurred only 14 times in 100 rounds of play. When Flood and Dresher showed these results to John Nash, he complained that the entire set-up was more like one multi-move game, in which his theory would not apply, because ““There is too much interaction, which is obvious in the results of the experiment”” (Poundstone 1992). Indeed, but it is precisely such interaction that interests us when we apply the Prisoner’s Dilemma to the marketing world, because the interaction involves an exchange of information about the other’s intentions.

In the Flood and Dresher experiment, each player’s direct intentions in each round was still kept hidden, but each player could glean some insight into the other’s future intentions based on their past actions, in exactly the same way that poker players observe when other players have a tendency to bluff, even if their actual hand remains hidden. Thus information still changes hands, even if it is not complete information.

We see this play out in consumer responses to various forms of marketing. Consumers’ wholesale defection from banner advertising in the early aughts was a direct response to the perceived defection on the part of marketers. Consumers had been subjected to ads with fake interfaces, which led to unintended clicking, as well as pop-unders, pop-overs, and all manner of dirty tricks. While the majority of advertisers did not engage in such tactics, the impact of the defection was wholesale; consumers mistrusted the intentions of online advertisers in general. They stopped responding, and it required many rounds of cooperation over many years for the relationship to regain its equilibrium.

### **3.7 The Persistent Problem of Bad Apples**

The consequence of early defection, after which many rounds of cooperation are necessary to rebuild trust, is one of the key insights of the iterated Prisoner’s Dilemma. Simply put, defection is a short-term gain but a long-term loss. It leaves a poor first impression, or at least an impression that one is a defector by nature, which is a bad thing to be in an iterated Prisoner’s Dilemma. Thus Poundstone describes the primary ingredient of

the Prisoner's Dilemma as "a temptation to better one's interests in a way that would be ruinous if *everyone* did it" (Poundstone 1992).

And therein lies the rub. In marketing, successful cooperation requires not only that the marketer signal cooperation to the consumer, i.e., "If I you click on this banner ad, you can trust me that it will not trick you," but the marketer must also somehow reinforce cooperation among his fellow marketers, i.e., "We will not fool consumers, and we will punish those who do."

This is the classic problem of the bad apples. In the statistics I cited earlier regarding distrust for commercial email, it is clear that consumer disgust with the bad apples – the illegal spammers – has indeed spread to the whole bunch, and not without reason: marketers are continuously probing the gray areas left open by CAN-SPAM restrictions, once again building a reputation for defection. Similarly, consumer disgust with telemarketing phone calls – many of which were misleading and overly persistent – led to the creation of the National Do Not Call Registry in 2003. Telemarketers had been engaged in iterative acts of defection – reaching consumers in a way that they actively disliked – and consumers finally defected en masse in retaliation. An astonishing 72% of Americans had placed themselves on the list by 2007 (Federal Trade Commission 2007), easily the single greatest act of consumer revolt in the history of marketing. Legitimate marketers vehemently opposed the legislation, arguing that it amounted to punishing all for the acts of a few. And indeed it did, but such is the retributive nature of the iterated Prisoner's Dilemma.

So the tendency of marketers to shoot themselves in the foot has a strong historical basis, even if the benefits of cooperation are equally well established. It is far too easy for individual marketers to act in their own short-term interests; it is much harder, in a highly competitive landscape, to act in the long-term interests of the common good. How, then, could cooperation possibly emerge?

The simple answer is that it emerges in the context of a set of rules that resolve the Prisoner's Dilemma by ensuring that players can properly signal their intentions and cooperation can flourish. But significantly, these have to be rules within the game itself, i.e., between players, and not ones imposed from the outside. Marketers fought legislation like CAN-SPAM and the Do-Not-Call Registry because they naturally feared that such rules would overly restrict their ability to play the game. Rules that emerge in the context of consumer response work the best, because marketers have to pay attention to the moves of their opponent in order to play the game.

What marketers really need is for consumers to be able to signal their cooperative response, i.e., they need the information coordination described

earlier. If the emergent rules aid the transfer of information, so that market-ers feel assured of consumer cooperation, then they themselves are more likely to cooperate. Is it possible to play the game with this level of clarity and conviction? Generally, yes. It has been accomplished in the ground-breaking iterated Prisoner's Dilemma strategy known as TIT FOR TAT.

### 3.8 The Enduring Relevance of TIT FOR TAT

TIT FOR TAT is the product of a competition conducted by the political science professor Robert Axelrod in 1980 and described in his 1984 work *The Evolution of Cooperation*. Axelrod invited academics from the fields of psychology, economics, political science mathematics, and sociology – all of them familiar with the Prisoner's Dilemma – to submit computer programs that would play iterated games of the Prisoner's Dilemma against each other. Each program would be pitted against another, round robin style, and 200 rounds would be played in each match-up. Because the games would be played by computer, emotional effects – like hoping for cooperation – could not taint the results.

Fourteen programs were submitted in Axelrod's first experiment. Many were highly sophisticated, comprising dozens of lines of code. TIT FOR TAT was astonishingly simple, consisting of only 4 lines of code and a strategy so basic that a kindergartner could play it: *Cooperate in the first round. After that, do whatever your opponent did in the last round.*

TIT FOR TAT not only handily won Axelrod's tournament, it won a second tournament among 62 contenders that tried to improve on its initial success, and it has never lost its first-place status after three decades. Its logic is unassailable. It offers no pattern that an opponent can exploit (recalling the importance of randomization in iterative games), because it simply responds to its opponent's actions. The opening cooperative move sets a positive agenda that encourages further cooperation, but if the opponent defects, he is continually punished until cooperation is regained.

Axelrod offers useful conclusions as to what makes TIT FOR TAT such a successful strategy. TIT FOR TAT works, Axelrod claims, because it is "nice, retaliatory, forgiving, and clear." One could hardly hope for a more useful and succinct explanation of how to succeed in mutually dependent conflicts. Axelrod goes on to show how these same features could be found in other conflicts that engendered cooperation, including the *détentes* that emerged in WWI's trench warfare, when enemy soldiers refused to fire on each other, and in the evolution of biological systems. The beauty of TIT FOR TAT, in Axelrod's view, is that it succeeded in engendering better

behavior in its opponents, i.e., it elicited “behavior from the other player that enabled both to do well” (Axelrod 2006).

As we'll see in the next chapter, this overall elevation in the level of discourse and behavior on both sides is critical to the success of social media marketing. We would do well to remember that these strategies are not the least bit altruistic; they serve the interests of both sides, and there is punishment awaiting the defector (social media offers ample means of punishment, as we'll see). So TIT FOR TAT is no patsy strategy, but it also hopes for the best, taking the risk of a cooperative opening move. This turns out to be key: Axelrod reports that “the single best predictor of how well a rule performed was whether or not it was nice.” Niceness was a feature of all of the top eight performers, and none of the bottom seven. From that we can derive another rule that will be critical to marketers approaching the social media space: *In an iterative game, never be the first to defect.*

### 3.9 The Dangers of the Death Spiral

The above rule actually underscores the one potential weakness in the TIT FOR TAT strategy. It is simply this: mistakes can be fatal. In a TIT FOR TAT software program, the chances of a glitch causing an erroneous response are rather low; it is, after all, a very simple piece of logic. But if the program did manage to defect by accident in response to cooperation on the previous move, it could prompt the opponent to defect in response, resulting in another defection, and so on. This outcome has become known as the “death spiral,” and it's a popular device in Hollywood movies like *Reservoir Dogs*, when the characters find themselves in an armed stand-off and some accidental stimulus causes everyone to shoot each other (Tarantino 1992). It's also a distinct possibility when marketers are playing real-life TIT FOR TAT; the history of banner advertising is something like this death spiral scenario.

I mention the death spiral because it will be important in examining social media blunders in the next chapter. It's relatively easy for well-intentioned marketers dabbling in the social media space to overstep their bounds and provoke rather vitriolic consumer backlash. In general, these consumers are not being intolerant; they are making an iterative move in a longstanding game in which marketers have a long history of overstepping. In such cases, marketers would do well to employ a variation on TIT FOR TAT that has proven effective against the death spiral. It's called TIT FOR TWO TATS, and as the name implies, it allows the opponent to defect twice in a row before a retaliatory move, leaving more room for cooperation. Marketers

that feel stung by consumer backlash in the social media space should consider this strategy before backing away from the space.

### 3.10 The Marketer's Dilemma

At this point, if I have done my job, I have established that iterative games hold a great deal of hope for evolving toward cooperation, and that this may have some explanatory power for what's going on in social media marketing. However, there's an important caveat: if you paid attention to the difference in the payoff table for marketers and the payoff table for the Prisoner's Dilemma, then you've noticed that the marketer's dilemma is a much tougher one than the prisoner's. Marketing has an equilibrium at the worst-case scenario (1-1, a death spiral), and a best-case scenario (4-4, a seemingly unattainable goal), whereas the Prisoner's Dilemma finds a point of equilibrium in the middle ground (neither the best or worst case).

So unlike the stable Prisoner's Dilemma, marketers and consumers vacillate back and forth over contested territory, each claiming conditional victories. Consumers gain temporary advantages with things like Tivo, which reduces their exposure to unwanted advertising; marketers regain the advantage with things like product placement in films, which replaces a portion of the lost exposure. And the dance continues. The game is perpetually sub-optimal for both players, with the added stress of instability. So how could this game possibly evolve to the optimal equilibrium?

As the next chapter will show, the short answer is: Not every easily, and not all at once. The long answer may lie, at least in part, with the ethical dimension to Axelrod's analysis that goes beyond the material (or at least points-based) rewards that success in the Prisoner's Dilemma promises. As each side gains an understanding of the other's self-interest, something like empathy emerges, so that the act of cooperation is ennobling – thus changing the stakes of the game. Anyone who has developed brand loyalty based on a brand's apparent trustworthiness and care for its consumers has experienced this greater reward. Our loyalty in such cases goes beyond the purely rational; it touches an emotional core. Axelrod found that in these instances, "the very experience of sustained mutual cooperation altered the payoffs for both players, making mutual cooperation even more valued than it was before".

Flood and Dreshers' and Axelrod's experiments showed that sustained, stable cooperation is possible and even preferable in games for which the dominant strategy is to defect – certainly the marketing game is included in this category. Cooperation requires iteration, mutually agreed upon rules,

and most importantly, transparency that breeds trust, allowing all players to signal their willingness to cooperate. For marketers, this evolution will be a long journey, and failures tend to make headlines. But in the examples we'll examine in the next chapter, the worth and the long-term stability of cooperative strategies will prove themselves out.

## Chapter 4: Consumer Revolt and the Rising Cost of Defection

*ABSTRACT: Consumer defection from direct-marketing tactics like traditional banner advertising has served as a catalyst for more collaborative marketing formats to evolve. Paid search advertising represents an important evolutionary step, because it involves consumers directly in assessing the quality of content and compelling advertisers to cooperate by providing relevant content. The growth in opportunities for direct consumer feedback has produced grim lessons for marketers, as consumers are able to take punitive measures against brands that defect. Instances of consumer backlash through social media, often in the form of viral videos that counter brand messaging, have produced tangible results in compelling marketers to cooperate with consumer demands. Consumers' eagerness to punish defecting brands has both a historical basis – in marketing's long history of defection – and a neurological basis, as individuals take pleasure in enforcing rules of engagement.*

As most of the last chapter was spent looking at cooperative strategies through the bird's eye of theory, it's time to come down to earth and put theory into practice in real-life examples. When I speak of an evolutionary theory of marketing, I confess that I am mainly interested in creating a simple counterpoint between traditional direct marketing and the transformations taking place in social media, because my aim is to make game theory useful to social media marketing, rather than to provide a full history of how we got here. But making the leap from direct marketing to social media marketing is a bit like flipping an evolutionary switch from Cro Magnons to modern humans, with no acknowledgment of the steps in between.

So it's worth spending a few pages on how other marketing tools fit into the evolutionary chain of events. I've already noted how banner advertising itself has evolved, primarily by two means: 1) by forsaking (at least partially) the click-through metric that provided a false equilibrium and encouraged defection, and 2) by making itself more relevant through better



targeting. Those changes are prescriptive of changes in digital marketing as a whole, and so they're important to how social media marketing emerged. I'll examine them in order.

## 4.1 The Unreasonable Standard of Immediate Action

Fundamentally, the click-through metric fails because it saddles advertising with a task that it was never meant to fulfill at all, i.e., to be so persuasive as to divert the consumer from whatever they were doing and get them to do something else, i.e., *to willingly absorb even more marketing messages* on a Web site, and often to make a purchase. Since this behavior runs completely contrary to consumers' attempts to *reduce* exposure to advertising messages, it seems entirely unintuitive.

It also ignores decades of market research on the role of advertising in consumer persuasion. The classic debate is whether advertising is a "strong force," capable of changing consumer brand preference, or a "weak force," capable only of reinforcing or defending brand preference (White 1999). Notably, neither of these two schools of thought envisions a role in which advertising's persuasive force is sufficient to induce immediate action, and that the advertising itself is the conduit for that action; such a role might be termed "overwhelming force." The chasm between what advertising can do and what it was being asked to do in the case of click-through metrics means that the death-spiral described in the previous chapter was inevitable. As long as the cost for cooperation to the consumer – stop what you're doing, respond to this ad – was higher than the potential reward, consumers would continually defect, and marketers would desperately chase their losses through reciprocal defection, i.e., with more pernicious advertising.

By rejecting, at least to a large extent, the "overwhelming force" role for banner advertising, marketers have once again made it a viable medium. Consumers now accept that a certain amount of banner advertising is necessary to enjoy free content, and marketers accept that banner advertising may have other beneficial effects besides inducing immediate action.

I raise this issue because the acknowledgement that immediate action is an unreasonable standard for online marketing was a crucial evolutionary step that allowed social media marketing to evolve. More reasonable stakes for both marketer and consumer open the door for cooperation, because both sides can realize a long-term benefit. (Unfortunately, some practitioners of social media marketing are now calling for immediate-action standards, and if successful, the effects will be ruinous. Those who cannot remember the past... you know the rest. I'll address this dilemma in the final chapter.)

The second evolutionary change in online marketing – enhanced relevance through better targeting – plays an equally crucial role in setting the stage for social media marketing to take hold. As with the withering away of the click-through metric, greater relevance has the effect of lowering the *cost* of cooperation for both sides. Starting from the obvious premise that it is sub-optimal for advertisers to put advertising in front of poor targets, and it is sub-optimal for consumers to be exposed to ads that are poorly targeted to them, both sides gain points in the payoff if the right ads reach the right consumers.

In banner advertising, a plethora of innovations – good aggregated consumer data, behavioral targeting – have made it possible for marketers to achieve an astonishing degree of relevance, so that an ad can be served to the individual consumer based on such granular factors as their past purchase behavior, their content preferences, and their location.

This degree of relevance begins to edge us toward the seemingly impossible 4-4 standard, in which marketer and consumer act in concert in fulfillment of the mutual goal of a satisfying engagement. But advertising itself can never reach this standard, because it always exacts a cost from both parties in creating these moments of potential persuasion. Consumers would always rather have their content commercial-free, and marketers would always rather get consumers for no cost rather than some cost.

But advertising can get *closer* to this standard, and it has done so in an evolutionary step beyond banner advertising and before social media marketing – the missing link, as it were. I refer to the evolution and maturation of paid search as an advertising model.

## 4.2 The Missing Link in the Evolutionary Chain: Paid Search

Few marketers may recall that paid search was explicitly developed as an alternative to the click dilemma posed by banner advertising. As Daniel C. Fain and Jan O. Pedersen recount in “Sponsored Search: A Brief History,” search engine listings were originally monetized with banner advertising in the same way as other free Web content, but this model was especially burdened by the problem of how to get users to click away from relevant content (the search results they were looking for) to less relevant content (the advertiser’s pitch).

The solution was to tie the advertising results explicitly to the subject matter for which the user was searching, so that the user might consider the paid content as worthy of a click as the “natural” search results. This further required displaying the paid advertising result alongside the paid results in

text links that made them appear more natural. But advertisers were less willing to pay for ad impressions that were merely text links, not graphical ads, and so the “pay per click” (PPC) model evolved, whereby the search engines would only charge the advertisers for actual clicks on the ads (Fain et al 2005).

This model held great advantages for advertisers, since they could extend their relevancy all the way to consumers actively searching for a given subject, and only pay for those consumers whose interests extended to action, i.e., who were ready to click on the ad because its content was relevant. In this way, marketers could cooperate by serving relevant content to interested consumers, and thanks to the PPC model, the reward outweighed the cost of cooperation.

But for the model to work, the reward had to exceed the cost for consumers as well. If some marketers engaged in short-term defection and served up PPC ads that were irrelevant or misleading, e.g., advertising an electric furnace when the consumer searched for a gas one, then consumers would defect in retaliation, learn to ignore PPC advertising, and the entire model would collapse.

But it has not collapsed. In fact, paid search marketing has experienced double-digit percentage growth each year since the adoption of the pay-per-click model in 2003, and is expected to reach nearly \$11 billion in 2009 (eMarketer 2009). Some of this growth has come at the expense of banner advertising, with its slower-evolving pay-for-impressions model. Such growth is entirely predictable from a game theory perspective: it offers greater rewards for cooperation than defection, and so it moves both the marketer and consumer closer to the 4-4 quadrant. But the thornier and more relevant question for the evolution of social media marketing is *how* it managed to reinforce rules of cooperation and prevent short-term defections from ruining the system. Understanding how such systems enforce order may have a great deal of relevance to the new systems evolving under social media marketers.

### **4.3 How Paid Search Gave Free Riders the Boot**

Everyone intuitively understands the problems created by spoilers or free riders, because we encounter these problems every day. The person who cuts ahead in the bakery line threatens the whole system, because the people who wait patiently in line may decide that they are suckers for sticking to the rules and defect, causing a run on the pastry counter. The problem is exacerbated in large systems like corporations, where collective rather

than individual responsibility makes rule-breaking easier; a competitor who breaks the marketing rules with pernicious tactics and gets short-term results creates a temptation for others to do the same.

In the case of paid search marketing, the search engines themselves act as a governing body, enforcing rules that discourage defection. They have an obvious incentive to do so, since they are fighting to preserve a business model that requires consumer cooperation. To begin with, search engines require relevance and audit their advertisers for compliance, ensuring that the content of the Web site that follows a click on the ad is pertinent to the ad's content. In this way, the rules of participation are made clear to marketers.

Secondly, they offer a system of rewards and penalties for enhanced or degraded relevance. The advertiser's rank – the order in which the ad appears on the page – is determined not merely by the amount the advertiser has bid, but also by something called a “quality score.” The search engines assign a quality score based on the popularity of the content with consumers, i.e., by how often a given ad receives a response. Advertisers with higher quality scores achieve higher rankings for less cost, while advertisers with poor quality scores rank lower. Poorly ranking advertisers have a built-in incentive to improve their scores with more relevant content – which includes bidding on only the most relevant keywords – because they'll save money and get more traffic. Consumers are rewarded with more relevant content, and the search engine is rewarded with revenue.

This small feature – search engines' enlistment of consumers in deciding the worthiness of an advertiser – is a groundbreaking yet overlooked development in marketer-consumer cooperation. Historically, advertising systems have been premised entirely on the advertiser's willingness to pay, not on the consumer's interest. Impressions simply went to the highest bidder. Marketers still had a built-in incentive for relevance, since they wouldn't wish to pay to reach the wrong consumers. But such systems still leave wide latitude for defection, as we've seen.

What the search engines uncovered was that a system that engaged marketers and consumers in the mutual pursuit of relevance could benefit all parties. As in all cooperative games, it required a long-term perspective to develop such a system, since a more traditional system based solely on the highest bidder would produce greater short-term revenue and be easier to administer. The new system also worked because it tapped into the broader sociological change brought about by the Web, in which consumers simply demand more control. As we saw in the case of TripAdvisor, the availability of other information sources requires that advertisers compensate for the loss of their informational role with greater relevance and better content.

And for the consumer, the cost of defection has been greatly reduced; if the marketer defects with a low-relevance paid search ad, the consumer can defect in retaliation by choosing a competitive listing, all in a matter of seconds. The search engines merely recognized that consumer engagement was a necessary – if subtle – ingredient in developing an ad system that worked in this new medium, and they made the evolutionary leap forward.

#### 4.4 Enforcing Rules of Cooperation

The search engine model has some important features besides consumer engagement that make it relevant to social media marketing. It is a system with a clear set of rules, rewards for cooperation and punishments for defection, and the involvement of its participants in rule enforcement, so that marketers are less likely to worry about free riders ruining the system while the search engines aren't looking. But it still requires the outside enforcer – the search engine – to mete out rewards and punishments. In the example of the bakery, by contrast, we're apt to see rule-enforcement occur without the intervention of the bakery's staff. A person cutting in line usually *doesn't* cause a run on the pastry counter, because the rule-breaker is called out by others in line and firmly ordered to go to the back, with lots of accompanying dirty looks. Why?

There is, in fact, an entire field of study devoted to answering this question. The study of “common pool resources,” pioneered by Elinor Ostrom, draws on game theory but also political and social science to analyze cooperative systems and uncover the common set of best practices for enforcing cooperation. In general, these features include a clear set of rules, rewards and penalties, and a reasonable assurance that defectors can be detected and penalized.

We see all of these features in paid search marketing, but only the third feature – detection of defectors – is in the hands of participants, and only in an indirect way, through the search engines' ability to measure consumer preferences and distribute rewards or penalties in the bidding system. For this reason, paid search cannot be said to be a fully evolved system of marketer-consumer cooperation. The rules of engagement are set by a third party that stands to gain the most by their enforcement. The marketer-consumer interaction is still only a means to an end – a way for consumers to get access to search engines for free, and for marketers to gain an ad platform – rather than an end in itself. In order to evolve the relationship further, marketers and consumers must cooperate in common pool scenarios, where the community of participants enforces its own rules.

## 4.5 Social Media and the Peculiar Pleasures of Punishment

To show how marketers and consumers are making this last evolutionary step, I need to begin with the concept of punishment. At the risk of being overly reductive, I will contend that marketers' participation in social media's rewards has generally begun with punishment. In consumers' minds, marketers have historically been the line-cutters in the bakery, seeking to maximize their advantage at a cost to the collective good. In the long-term iterative game that is the marketer-consumer relationship, the marketer has gained a reputation for defection, and consumers have uncovered a new form of retaliation, thus changing the stakes of the game. Several well-publicized incidents of punishment have made marketers even warier of social media participation, but such incidents are merely a natural part of the maturation of a mutually rewarding cooperative system, as we'll see.

We've seen already in the case of TripAdvisor that social media venues can provide a powerful reality check on the truth-claims made by marketers, forcing marketers to contend with the power of consumer reviews on their own terms. But social media offers even more powerful means of consumer retaliation in its capacity as an outlet for consumer-generated content.

To understand how this works, first consider how advertising functions at a level of remove from the brand experience itself. Allow me to state the obvious: advertising is not the brand relationship itself, it is a *signal* of what the marketer wishes the relationship to be. Advertising gives marketers the freedom to idealize that relationship in an act of wish-fulfillment; if it is successful (and if one accepts the notion of advertising as a strong force), it prompts the consumer to aspire to such a relationship too. A Nike ad makes the consumer want to be committed and resourceful in pursuit of their fitness goals, and to see Nike products as the means of achieving those goals.

In traditional marketing, marketers have had the brand megaphone mostly to themselves. A brand has the financial resources and the media access to own most of the arguments about what the relationship should be; consumers are relegated to the passive role of accepting or rejecting that argument.

That is, they were until now. A crucial step in the evolution of cooperative marketing is the power consumers now have to dole out punishment to enforce cooperation. Consumers with access to social media can engage in a kind of counter-advertising or counter-signaling, disputing the messages being conveyed about the brand. Doing so raises the stakes of cooperation, forcing the marketer to consider not merely whether and how to advertise but how to improve the customer experience.

## 4.6 “United Breaks Guitars” Breaks Through

This is best explained by way of example. Let’s begin with one of the most prominent incidents of consumer revolt in recent memory, the “United Breaks Guitars” social media phenomenon. In spring 2008, musician Dave Carroll and his bandmates were traveling from Nova Scotia to Nebraska. While the plane was on the tarmac, Carroll witnessed baggage handlers carelessly tossing his band’s guitars. Upon landing, he discovered that his prized Taylor acoustic had indeed been broken in transit. In his telling, Carroll began a Kafkaesque series of attempts to gain recompense from the airline. He was repeatedly turned away (United claimed he had not registered his complaint at the correct time for it to be honored), and he vowed to the last United employee to deny his claim that he would exact a musical revenge – three music videos that would expose United’s transgressions (Reynolds 2009).

Let’s return to the theoretical level for a moment to make sense of what’s going on here. In game theory, promises and threats are viable means of reinforcing cooperation: promises indicate a future reward, and threats indicate a future penalty. Threats are particularly useful because the player bears no additional cost if the threat is successful, whereas successful promises must be fulfilled. But for threats to be effective, they must be *credible*, i.e., the recipient must have some reasonable assurance that the threat could be carried out (Dixit & Nalebuff 2008).

Needless to say, individual consumers have not traditionally been able to make credible threats to do damage to a brand’s reputation. The consumer would bear great cost in trying to generate sufficient publicity to do harm, with limited chance of success. The economics of customer service account for this low probability; they do not dictate, of course, that the brand does nothing to help unhappy customers, but they also do not add a large multiplier to each unhappy customer, counting on their ability to spread the word of their unhappiness far and wide. According to the *Los Angeles Times*, citing statistics from the Department of Transportation, United Airlines had 13,517 “baggage reports” in April 2009 alone (Reynolds 2009). Like all airlines, United places conditions on its damage compensation policies because the cost of customer irritation is calculated to be less than the cost of making good on every single claim. From a game theory perspective, this is a very rational response, even if it strikes us as cold-hearted or offends our notions of customer service.

Thus in handling customer baggage complaints, United had created, whether they realize it or not, a minimax point, whereby they minimize their maximum loss in paid damage claims, providing compensation in

enough cases to keep them out of serious hot water, but not enough to make every customer happy. They were playing a zero-sum game because it was seemingly the best strategy available in this unpleasant corner of the customer experience, which always starts out with unhappy customers. This small zero-sum game had no effect on the larger game being carried out in United’s marketing. Or at least it didn’t until Dave Carroll came along and changed the stakes.

Carroll proved that his threat was more credible than the airline could have possibly envisioned. Possessed of some natural advantages that many online content producers lack – a gift for songwriting, access to good production resources, and a sly sense of humor – Carroll produced a savagely witty song and video with his complaint laid out starkly in the title: United Breaks Guitars. The song quickly became that most elusive of things, a YouTube phenomenon. It logged more than 100,000 views within a few days, and as is typical of such phenomena, it took hold in mainstream media, showing up in countless news accounts, blogs, TV segments, talk shows, etc. As of this writing, the video has garnered 5,656,458 views on YouTube and 22,456 comments – nearly all of them vitriolic in their condemnation of United.

Carroll’s video is only the most prominent example of a broader set of behavior in which ordinary consumers engage in counter-signaling through social media, and they are effectively able to disrupt the signal being sent by the brand through traditional marketing. By escalating his complaint to social media, Carroll transformed his private complaint into an argument about the brand itself and its relationship to its customers. It would be difficult to understate the strength of this signal to social media users; as of this writing, United’s own most recent advertising, by contrast, has garnered 48,315 views on YouTube, or less than 1/100<sup>th</sup> of Carroll’s audience. Who wins?

The answer to that question is more complicated than it first appears. To anyone who has ever felt deeply wronged by an airline – which is to say, nearly anyone who has ever flown on an airline – Carroll’s musical revenge is deeply satisfying. The fact that Carroll eventually did receive full compensation from the airline is further proof of the rough justice that social media is capable of exacting. But brand-consumer relationships are iterative games, and this is only one round. The “United Breaks Guitars” incident generated a great deal of overheated commentary about consumer revolt, including a claim by Chris Ayres of the *Times Online* that the PR fallout from the incident caused United’s stock to fall 10 percent, “costing shareholders \$180 million” (Ayres 2009).



## 4.7 Recalculating the Cost of Defection

If such incidents did indeed produce \$180 million swings in stock values, brands would rightly be in full-blown panic over social media's negative impact, and no equilibrium would even be possible. Consumers could effectively hijack brands with their counter-signals, bringing them to their knees. But that hasn't happened. Analyzing the more subtle effects of the incident is not only more accurate but more useful. Incidents like Carroll's, as well as others we'll examine, first and foremost have the effect of forcing brands to reevaluate how they play the game. United could not use its traditional signaling method – paid media – as a way of correcting the perception created by Carroll's video, i.e., they could not defect in retaliation without making their long-term prospects worse.

Rather, Carroll's original threat was intended to compel cooperation, and it achieved its purpose, so that the long-term prospects of further cooperation between United and its customers are actually improved by the incident. United declared that it intended to use the video in training employees on how to better handle future incidents, and this is perhaps the most important effect of all: it suggests that United now believes such threats are credible and is more likely to make cooperative moves in future iterations, thus bringing its signaling and its actual behavior in closer alignment.

To review: counter-signaling through social media arms the consumer with a new weapon in the iterative marketing game. The consumer's usual means of defection – signal-blocking and brand rejection – are augmented by new punitive tools that raise the stakes of the game and are more likely to compel cooperation in subsequent rounds of play. The counter-signal exacts a cost that is a multiplier on the defection of a single consumer, because it causes other consumers to defect and/or send up their own counter-signals. Marketers' recalculation of the cost of defection becomes an important evolutionary step, one that brings marketers directly into the social media playing field.

But before we consider marketers' forays into social media, there's more to examine on the phenomenon of consumer use of social media to compel cooperation. On a much smaller scale than Carroll's social media juggernaut, consumers have quietly taken up the practice of what amounts to a kind of public shaming of brands that defect from their brand promises. Some of the better-known videos include a cable repairman asleep on a customer's couch, captured on the home's nanny-cam, and a phone recording of an AOL customer service representative that refuses to allow a customer to unsubscribe. By gaining viewership, commentary, and media attention,

these counter-signals have functioned as a kind of public shaming of errant players – one of history’s oldest forms of punishment.

#### **4.8 Changing the Defection Stakes for Banks: Debtor Videos**

It may be stating the obvious to point out the importance of the public or social component to the success of these counter-signals; if Carroll had made a great video that no one saw, United would have been far less likely to act. I raise this point in order to highlight the commodity being contested in these exchanges: social attention is a commodity in finite supply, and consumer revolt works only as long as an adequate supply of this attention can be obtained. But the supply will, in fact, run out, and counter-signaling as a tactic will be diminished. It is a short-term tactic in a long-term game.

How do I know this? I am thinking of the very recent phenomenon of “debtor revolt,” in which consumers post videos and write blog entries about their struggles with banks and lenders in the midst of the worst credit crunch since the Great Depression. In the standard narrative, the consumer complains that the lender or credit card issuer won’t renegotiate terms that would allow the consumer to continue paying their debt, instead forcing them into default. The more popular videos are the more outrageous cases: a student whose education loan deferment is mistaken for a delinquency, or a landlord seeking a point-and-a-half- reduction in his mortgage rate in order to keep his rentals afloat.

The posting of the video is seen as the logical next step beyond the consumer’s failed efforts at 1-to-1 negotiation, but from a game theory perspective, it is much more. As with United Airlines’ baggage claim department’s handling of Dave Carroll’s complaints, these banks and lenders are applying a minimax theorem to the problem of delinquency, setting a threshold at which they believe it will cost them more to extend additional consideration or leniency. On an individual basis this minimax theorem produces outrageous cases, because the lender also loses if an otherwise viable customer defaults over their failure to renegotiate terms. But on a macro level, the minimax theorem prevents the lender from giving up so much ground to consumers that they fall back into the “toxic asset” problem that created the credit crunch to begin with.

For the consumer, then, the debt video is an attempt to force the other player out of the zero-sum game and into one with a cooperative solution by changing the stakes of the game. In this new game, the consumer implicitly argues, the bank’s cost of defection is not merely the financial loss – which the bank has already calculated in its algorithms for dealing with

delinquency – but the compounded loss of positive brand equity. Without knowing for certain whether the consumer can pay back the loan, the bank believes it has a sub-optimal but stable solution in not renegotiating, and this calculation is correct only insofar as the bank has taken all of its possible losses into account. But the consumer destabilizes that solution by increasing the penalty for defection and compounding the potential loss.

This strategy appears to be working. The landlord received a response from his lender within 4 hours of posting his video, with a promise to investigate his case. Other celebrated cases also received high-level and direct communications from bank personnel empowered to renegotiate. In the case of the landlord, a bank spokesperson claimed that “if he had sent that letter without the video, he would have gotten the same response,” but this is unlikely. The bank’s calculation changed only when its cost of defection changed, as we would expect (Delaney 2009).

The consumers in these debtor revolt games are winning, at least temporarily, not because their complaints are reasonable (they may be, but that is largely irrelevant), but because they are successfully seizing control of a commodity that the bank brands prize very highly: attention. Positive attention in the form of favorable brand impressions is the commodity the brands purchase in order to win the marketing game, and if consumers can acquire negative attention for that same brand for free, the game changes drastically. But as I noted, attention is a finite commodity in a saturated media environment. A few debtor revolt videos can hold the attention of banks because they can also hold the attention of a sufficient number of Web users and media outlets. But there were 358,471 foreclosure filings in August 2009 alone (RealtyTrac 2009). How scalable is this tactic?

## **4.9 Attention Saturation and the Limitations of Punishment**

Recall that a standard feature of the non-cooperative game is that defection would be disastrous if everyone did it. In the competition for fickle consumer attention, consumer revolt videos on a large scale would quickly lose their luster, which would mean a loss of attention as a commodity. The banks’ cost of defection would suddenly be reduced to its previous proportions. Indeed, one can easily imagine consumer backlash against the complainants, because their defection would eventually be perceived as the “free rider” syndrome – players cutting ahead in the bakery line while everyone else has to wait.

I raise the issue of saturation in part because it is, I will contend, the single greatest delimiter of the effectiveness of social media marketing for

both marketer and consumer players. Saturation would have a deflationary effect on the value of the video, especially in a medium that places a premium on novelty. I've devoted the final chapter to the issue of saturation and information overload, so I'll explore this issue in greater detail there. The main impact of saturation in this hypothetical case is the sudden drop in the punishment stakes, as brands learn to triage customer complaints that appear in this format.

The optimal solution for consumers as a whole is for this initial defection move – the debtor revolt video – to compel a reevaluation of policies among the banks and a more flexible and accessible approach to individual cases. Since it would be sub-optimal for consumers if the banks either became insolvent or decided they could safely ignore consumer revolt videos, continuous defection would be a very poor strategy. As with Carroll's video, the goal in this game of TIT FOR TAT is to first compel cooperation from the other player, then respond with cooperation of your own.

Unfortunately, consumers don't behave as a monolith any more than marketers do, and coordinating total cooperation would be impossible even if it were a conscious strategy, which it is not. But that is why iteration is so important; as Axelrod showed, the potential for cooperation is directly related to how many times the game is played. As these social media exchanges become more commonplace, brands and consumers both evolve a set of unwritten rules – slowly and painfully, but inexorably – that dictates acceptable behavior, as we'll see in the next example. As debtor videos became more commonplace, *both* players would eventually turn a jaundiced eye on the phenomenon, in the interests of longer-term cooperation.

## 4.10 The Motrin Moms and Social Media Backlash

My third example of the use of punishment in compelling marketer cooperation is meant to bring us a step closer to a discussion of mutual cooperation, because it involves the increasingly familiar scenario of a marketer earnestly trying to make good use of social media marketing, overstepping the bounds, and incurring consumer backlash. As I noted previously, successful cooperative scenarios invariably require a widely understood, well articulated set of rules – a non-distorted transmission of information between players. Suffice it to say that such rules are neither well understood nor well articulated in the burgeoning social media space, and many brands that take the extra risk of going first take the brunt of the painful lessons for the rest of the industry.

One of the unfortunate side effects of the astonishing viewership of viral videos like Dave Carroll's is that marketers have tried desperately to emulate its success, with widely varying results. The question of what makes a viral video successful is indeed a fit subject for game theory, but it's also complex enough to belong in a later chapter. Here I'll simply offer a brief analysis of motives: marketers that develop their own viral videos can be justly credited with trying to talk to consumers in the social space where the content of their videos can be judged on their own merit, in collaboration with the consumer, as opposed to a video being thrust upon consumers by virtue of the marketer having purchased airtime. In the long continuum of cooperative marketing, it constitutes progress.

But the marketer is also interested in getting something for nothing – namely, ad impressions without media cost. Doing so would obviously improve the marketer's position in the marketing game, because they'd be able to increase the ratio of consumer response vs. the cost of exposure. In paid advertising, the marketer needs a certain number of consumers to like the video (or at least not actively dislike it) in order for the investment to be worthwhile. But in viral advertising, the marketer needs much more: the consumer must genuinely *like* the video in order for it to get any exposure at all, because the marketer relies on the consumer not just as the video's audience but as its distribution channel. This is a crucial lesson: the marketer's stakes are not actually *reduced* by using viral video; the marketer simply *replaces* an investment of actual capital – the cost of running an ad – with social capital, i.e., the effort of engaging consumers through a compelling viral video.

Unfortunately, many marketers that make the leap into viral video have not performed this analysis. They are still playing a mutual defection game focused on maximizing impressions at the lowest possible cost, and they miss the need for collaboration. This is in substance what occurred in one of the most notorious cases of social media backlash, popularly known as the “Motrin Moms” incident.

In November 2008, the pain reliever brand Motrin, a product of McNeil Consumer Healthcare, posted a cheeky viral video to their Web site. The ad poked fun at the trend of moms carrying babies close to their bodies in slings, wraps, or “schwings,” as the ad put it. The ad teased that this was an attempt to look like an “official mom” and offered Motrin as a product that could help mothers with the aches and pains of slinging their children.

The tone of the ad was not malicious, and when stepping outside of the controversy, one can even imagine that the ad was intended to “demonstrate genuine sympathy,” as McNeil later claimed. But it debuted in a medium that leaves scant room for error in tone, especially where satire is concerned

(Belkin 2008). The ad may well have been unpopular had it run on television, but the consumer's primary weapon of passive non-response – ignoring the ad or changing the channel – would have generated little heat. Online, the ad created not just heat but a wildfire.

Motrin posted the video on their Web site and YouTube on a Friday, and by Saturday, mothers using the microblogging site Twitter began spreading the word among their peer groups about the ad's controversial content. It is the nature of controversies spread through Twitter to grow exponentially, because each recipient of the controversial information can spread the content to a unique circle of friends, who each have their own circle, and so on. Since every consumer equipped with a personal computer is also equipped with the means to make their own video, protest videos began appearing on YouTube by Saturday afternoon. The most popular of them merely showed screenshots of Tweets posted by angry consumers, set to background music. Each of these videos in turn garnered dozens of comments and thousands of views, all in the same weekend. News accounts and blog entries quickly followed, and thousands more consumers with no direct stake in the controversy witnessed the conflagration (Belkin 2009).

The incident is especially chilling for marketers because the controversy reached a fever pitch before Motrin even became aware of it – over the course of a single weekend. There is no precedent for this in traditional advertising; far more controversial ads would take weeks to generate comparable reactions, and the controversy could be expected to ebb as soon as the offending ad stopped airing. In the Motrin case, the ad itself, all protest videos and associated comments, the angry Tweets, and the dozens of blog entries and news articles on the topic are all *still on public display a year later*, accessible by a simple Google search. As a form of counter-signaling, the protest achieved far larger brand awareness and perception effects than the original video could ever hope for. This level of defection is unquestionably a game-changer.

It's easy to see the Motrin incident as a prime example of consumer defection, but the lessons it holds for the marketers depend a great deal on whether one sees Motrin's original viral video as an instance of cooperation or defection. It fits much more neatly into my analysis if we label it "defection," because the logic of TIT FOR TAT then applies: Motrin defected, and Motrin Moms defected in response, and forced Motrin back to cooperation, starting with the lavish apology that appeared on the Motrin home page in the aftermath of the backlash.

Further supporting the "defection" label is the fact that the video itself is so tone-deaf to its audience, despite being reliant on that audience for distribution. Like many marketers, Motrin seems to have misread the signals

coming from viral video success, which suggest (distortedly) that “edginess” is a key ingredient to getting a video passed along. In this view, Motrin merely wanted to enjoy the lower hard costs that viral marketing could offer, without being willing to pay the social costs of understanding its audience and crafting something they’d like. And there’s a key lesson available in this view: *participation by marketers in social media does not itself constitute an act of cooperation*. It will require deeper engagement, as I’ll explore in the next chapter.

But I’m reluctant to consign Motrin to the dustbin of social media defectors so quickly, because the reality is a little more complicated. Social media often behaves like an echo chamber, in which an initial negative reaction pings around endlessly, and subsequent content consumers never see the original content outside of that negative context. In retrospect, Motrin’s effort to produce a hip viral video appears hamfisted and slightly embarrassing at worst – like a dad showing off in front of his teenage daughter’s friends.

Therefore I propose we assign to the Motrin gaffe the more precise label of “accidental defection,” which will in turn describe many of the more notorious social media miscues that have come from brands in recent years. You may recall that accidental defection is the Achilles Heel of the otherwise cooperation-friendly TIT FOR TAT strategy; it can create the endless cycle of retaliation described in the last chapter as the death spiral. Motrin’s response to its disastrous first foray into social media could be very long retreat away from the space and back toward its traditional marketing, and this would be a defection, even if it doesn’t raise hackles the way the viral ad did. This would be a long-term loss for both players, since Motrin clearly *wanted* to play well in social media, and the “teaching moment” produced by the backlash could help them to do so.

#### **4.11 TIT FOR TWO TATS and the Virtues of Forgiveness**

The solution would be for Motrin to play the modified TIT FOR TAT strategy known as TIT FOR TWO TATS, described earlier as a proven antidote to the death spiral. Under this strategy, an extra round of forgiveness prevents a single defection from derailing the entire game. Instead of retreating in response to consumer backlash, Motrin would try again, correcting the mistakes of the previous round.

It could be argued that the burden of TIT FOR TWO TATS belongs with the consumer, since Motrin made the original accidental defection. The consumer should forgive the transgression and give Motrin another chance.

After all, in the case of the embarrassing show-off dad, we do not expect the daughter to disown her father in retaliation. Yet many of the “Motrin Moms” did exactly that, swearing off the product for life. And while their anger certainly seems outsized in retrospect, it’s not atypical for acts of consumer backlash in social media. So how does game theory account for such reactions?

## 4.12 The Enduring Appeal of Punishment

The first explanation is the premise that began this chapter: consumers generally feel that they are the object of a very long-term pattern of defection on the part of brands and marketers, as evidenced by the very low trust in advertising shown in studies like the Edelman Trust Index. Marketers simply cannot treat their entrée into the social media space as though it is a clean slate – a brand new game apart from the one they’ve played for decades. Consumers will take advantage of the medium’s natural suitability for grassroots backlash on the basis of very little provocation. Motrin found this out the hard way.

The second explanation may, in fact, explain a great deal about the behavior of individuals in social networks in general. It springs from the study of player behavior in iterative coordination games, i.e., games in which players must coordinate their moves in order to achieve a common good. Dixit & Nalebuff describe the willingness of players to punish defectors even if the punishment came at their own personal expense. Players given the opportunity to lower their own payoffs in order to impose an even lower payoff on defectors did so *eagerly*. In a purely mathematical view of game theory, such behavior is highly irrational; it is, as Dixit & Nalebuff point out, “a dominated strategy,” since the punisher ends up paying more than the group as a whole.

We need to understand this behavior in order to make sense out of much of what goes on in social media. The Motrin Moms who spent the weekend making protest videos ostensibly gained nothing, not even personal fame, from their sacrifice of time and effort. Web community members that troll forums and message boards looking for rule-breakers seem similarly out of balance with their best interests. A player acting purely in their self-interest wouldn’t bother with punishment; they would simply tune out or ignore the rule-breaker.

But as it turns out, there is a personal benefit to punishing rule-breakers that balances out the cost, allowing us to reassign this behavior to the realm of the rational. In “The Neural Basis of Altruistic Punishment,” published in



*Science*, a group of researchers from the University of Zurich documented the results of brain scans conducted during cooperative games. They discovered that doling out punishment activated the pleasure centers of the brain in the dorsal striatum, and that even the *anticipation* of being able to punish defection was pleasurable (de Quervain 1994). Some subjects derived much more pleasure than others, and one can deduce that such subjects are also the people most eager to punish defectors in social networks. (The anticipation factor also explains why some players actively seek out defectors, instead of merely reacting to defection when it crops up.) The bottom line is this: any assessment of the cost of defection and the likelihood of punishment must take into account the existence of these “squeaky wheels” who are eager to punish.

This insight into human behavior would not delight the marketer contemplating a first foray into engagement with recently empowered consumers in the social media arena. If a broken guitar neck or a baby sling joke can ignite a social media firestorm, how does one go about protecting the brand’s interests in this new space? At what point do the benefits of participation outweigh the costs of non-participation, and how can that even be gauged?

While I’ve given a chapter’s worth of attention to the concept of punishment, social media marketing will ultimately be defined by its many successes rather than its few failures. It may be cold comfort to United Airlines, Bank of America, and Motrin, but these rounds of defection are as integral to long-term cooperation as the success stories; they advance the creation and transmission of the rules of the game. In the next chapter, I’ll explore the next evolutionary step: the success stories that result from the kind of clarity that can only come from failure.

## Chapter 5: Sustaining Marketer-Consumer Cooperation through Coordination Games

*ABSTRACT: Previous chapters have demonstrated that both marketers and consumers would choose more cooperative strategies if they could sustain and coordinate them. Coordination game concepts can be used to analyze this potential. Disarmament agreements in the Cold War exemplify coordination games at work; both sides have a built-in incentive to cooperate, but they must provide reassurances to the other side in order for the agreement to be trusted. The concept of the stag hunt, in which hunters reap richer rewards by coordinating their actions, describes some of the opportunities available in social media. My own concepts of the “exposed flank” strategy and the “neutral ground” strategy describe methods used effectively by marketers to signal and sustain cooperation with consumers in social media marketing.*

In the last chapter, I focused on social media’s grim lessons for marketers, because it is crucial to understanding its evolution to see how a certain amount of score-settling by consumers was and is inevitable. To recap: when seen as an iterative game, the marketer-consumer relationship is inherently unstable, with each side gaining temporary advantages over the other – consumers seeking to minimize their dependence on marketing while marketers try to maximize it. Marketing itself is inherently sub-optimal because both sides – brands and consumers – would prefer to reach their goals without it.

Yet as I showed, there is a seemingly unattainable point of point of equilibrium at which consumers and marketers engage in a meaningful exchange of ideas without the use of advertising – an optimal condition for both. Whether social media can actually provide such conditions, and under what circumstances, is the subject of this chapter. The instances of consumer defection in social media are part of the evolution toward that ideal state of cooperation, because each defection makes a contribution: it raises the cost of future defection by marketers, and it allows a transfer of information by consumers to marketers. Consumers say to marketers, as the Motrin Moms

did: *Here is what I need from you. Coordinate your response so that I feel respected and listened to, and I'll respond in kind.* In this context, consumer bloodletting helps to save the patient.

If you'll accept as a starting point that both consumers and marketers would like to cooperate to achieve that mutually beneficial ideal state, and that social media *might* have the potential to get them there, then we need a new theory – an evolution beyond TIT FOR TAT – to find the way. TIT FOR TAT shows the path toward cooperation, but even after that goal is established, getting there is tricky and dangerous, as Motrin found out. For this, we need to analyze the marketer's options as a *coordination game*.

Simply put, a coordination game involves an effort by each player to synchronize their moves with their opponent's so that both can achieve the optimal outcome. It assumes that both sides want the optimal outcome, but that the transfer of information between the two sides is imperfect or non-existent, so that both sides may settle for the sub-optimal if coordination is not achieved (Dixit & Nalebuff 2008). For the marketer and consumer, for instance, settling for the sub-optimal would involve going back to the old way of doing things, with marketers maximizing exposure while consumers minimize it, in a cycle of mutual defection.

## 5.1 Cold War Coordination Games

As is so often the case, game theory's original subject, the Cold War, provides the best examples of the theory in action. In the example from Chapter 2 of the Cuban Missile Crisis' delicate resolution, both sides ended their stand-off with a careful coordination of mutual cooperation, nearly undone by an errant bomber pilot. Nuclear disarmament followed a similar logic: the doctrine of Mutually Assured Destruction was a kind of Prisoner's Dilemma, with both sides accepting the sub-optimal equilibrium of a prolonged arms race because the optimal solution – mutual disarmament – was so difficult to achieve.

As shown in Table 6, the US and the USSR both faced unstable, dangerous outcomes if one of the players armed while the other disarmed. One side would have an advantage over the other, but unilateral advantages are a dangerous thing, as they tend to make the other side much more provocative. The sub-optimal equilibrium (3-3) in the lower right is the MAD solution; both sides have to spend stratospheric amounts on weapons that will end the world, but as long as the logic holds, no one gets blown up.

**Table 6:** Coordination payoff matrix

	<i>USSR: Disarm</i>	<i>USSR: Arm</i>
USA: Disarm	4-4	1-2
USA: Arm	2-1	3-3

The optimal solution in the upper left is the one that poses the coordination problem. In the absence of perfect knowledge, it is generally safer to remain at the sub-optimal equilibrium. The prisoner can't count on the other prisoner's silence, and Gorbachev can't count on Reagan's word alone. But the sub-optimal is so unsatisfying that an effort toward the optimal can evolve.

Whenever disarmament *was* achieved during the Cold War, it involved a careful signaling of intentions and continuous verification of the other's actions. Were the missile stockpiles properly disclosed? Were they actually destroyed? What other schemes did the Soviets have up their sleeves? Perfect coordination, which would have involved total disarmament, was so elusive that it occurred only in movies, and even then it required a *deus ex machina* like Superman or an alien power to force cooperation.

Marketers and consumers don't have a *deus ex machina*, but like Reagan and Gorbachev, they do have an incentive to avoid the cycle of mutual defection. If social media can provide the means for coordination, there's little question that both sides will want to play.

The first thing to know about the coordination game is that it is often – though not always – an adjunct to the Prisoner's Dilemma, as the above example shows. Multiple rounds of TIT FOR TAT often lead to the coordination game, as both sides decide on cooperation, then set about trying to achieve it. Axelrod recounts examples of German and British infantry in their trenches in WWI trying to restrain the actions of their own artillery so as not to upset the delicate truce that had emerged between foot soldiers staring down each other's rifle barrels across No Man's Land. Whenever one side shelled the other, there was a risk that the cycle of defection would begin again, and so the enemy infantries would provide reassurances to each other, blaming their artillery divisions and reaffirming the truce (Axelrod 2006). Behavior such as this is why the coordination game is also sometimes referred to as an *assurance game*.

## 5.2 The Stag Hunt

As these examples illustrate, the basic ingredients of a coordination game are 1) a scenario in which neither side can achieve the optimal outcome without the help of the other, but 2) both sides have a sub-optimal solution available if they go it alone. There are many variations on the coordination game, but the one that's most useful to analyzing the marketer-consumer opportunity in social media is known as the *stag hunt*. Interestingly, the stag hunt game did not originate with the game theory crowd at the RAND Corporation, but rather with a political philosopher who lived two hundred years before: Jean Jacques Rousseau is credited with developing the stag hunt concept as a way of explaining why individuals should accept the risks involved in entering into a cooperative social contract with others, as democratic systems demand.

In the stag hunt, we are asked to imagine two hunters who each have the choice of hunting stag or rabbit. The stag is a more valuable prey, but it requires that the two hunters coordinate their actions; a hunter who tries to hunt stag alone ends up with nothing. A safe option for each hunter is to accept the sub-optimal solution and hunt rabbit alone, even though the result is less valuable (Skyrms 2004).

This game produces a very simple payoff matrix that demonstrates the trade-offs of the coordination game. If both players "hope for the worst" and don't bother trying to coordinate their actions, they at least enjoy a stable outcome: a reliably productive day of rabbit hunting (3-3 in the lower right quadrant). They're better off than if they showed up for a day of stag hunting and found that the other had gone rabbit hunting, in which case they go home empty-handed (0-3 or 3-0). But if the two players can somehow coordinate their actions, they can both claim the big prize: the stag (4-4).

**Table 7:** Stag hunt payoff matrix

	<i>Hunter 2: Stag</i>	<i>Hunter 2: Rabbit</i>
Hunter 1: Stag	4-4	0-3
Hunter 1: Rabbit	3-0	3-3

At a glance, one can see that this game has two Nash equilibria, and that one is more valuable than the other. The difference between the two is the classic dilemma of safety vs. reward: the 3-3 payoff is a safer bet, but the 4-4 payoff is a richer one. Marketers contemplating whether to take the leap into social media marketing or remain entirely reliant on tried-and-true direct marketing techniques will sympathize with the hunter's dilemma.

As Poundstone points out in his analysis of the stag hunt, there is no question as to which outcome is better. The hunter would *only* choose the rabbit hunt if the hunter had serious doubts as to the other player's commitment to cooperation, i.e., whether they'll show up for the stag hunt (Poundstone 1992). And that is precisely why marketers and consumers come to the social media table only after a long and delicate dance, and why the stag hunt is such a useful game for analyzing social media marketing opportunities. The risk of defection – marketers taking advantage, or consumers not participating or flaying the marketer – still feels real and raw. So how does one choose?

Despite the attractiveness of the stag hunt outcome, a social media truism pertains: in strategies with multiple Nash equilibria, the optimal solution cannot be determined within the payoff table itself. One cannot get there without an analysis of the other player's motives and behaviors. That analysis seeks to answer one simple, critical question: how likely is the other player to cooperate?

When we look at examples of social media marketing efforts through the lens of game theory, we uncover ways to answer this question. What we find is that marketers have tried two main strategies for coordinating the social media stag hunt: 1. Marketers have learned to *signal* their interest in cooperation by openly exposing themselves to risk in social media, thus winning loyalty, in what I'll call *an exposed flank strategy* 2. Marketers have tried to reduce consumer perception of risk and enhance their perception of payoff by offering hands-off forums for feedback, in what I'll call a *neutral ground strategy*. In this chapter, I'll examine the exposed flank and neutral ground strategies in turn, by way of example.

### 5.3 Dell and the Exposed Flank Coordination Strategy

One of the happy outcomes of social media horror stories like Motrin's viral backlash is that companies that are burned by social media often evolve into its best practitioners. Long before the Motrin Moms, in the nascent social media era of 2006, the computer manufacturer Dell found itself in the social media crosshairs. The company had already earned a poor reputation among the highly vocal and influential technical blogger community for turning a deaf ear to complaints about its laptops, particularly the widely reported problem of overheating batteries. Dell's first attempt at corporate blogging, Dell2One, was considered a case study in doing blogging wrong, as it shamelessly touted Dell products while offering no outlet for discussion of customer complaints.

As is so often the case in an image-based era, it took visual evidence to turn grumblings of discontent into a full-blown PR catastrophe that the company was forced to reckon with. After several dramatic videos of Dell laptops catching fire or exploding began circulating online, Dell's public relations problem had finally reached critical mass. The videos of flaming laptops were so morbidly enthralling, and so vindicating to anyone whose lap had ever been overheated by a Dell, that the impact overwhelmed Dell's bland reassurances that the problem was a limited one. Dell was forced to recall 4.1 million batteries and acknowledge its slow response to the problem.

And then a remarkable thing happened. In the process of taking accountability, Dell began to embrace its need to engage in social media with the evangelical fervor of the recently converted. With unprecedented candor, Dell reached out to the very blogs and forums it had assiduously ignored and owned up to its mistakes. Dell built one of the most sophisticated social media marketing teams in the business, with as many as 40 full-time employees dedicated solely to responding to customer conversations.

By Dell's accounting, the process was transformative. The company's online forums reached a membership of 1.5 million, with up to 10,000 customer queries being responded to by other customers each week. Dell representatives respond to all blog posts concerning their products, both positive and negative. The company's Twitter following is in the Top 50, with more than 700,000 followers. And the company created the "IdeaStorm" site to solicit and promote user-generated ideas for making Dell products better. It produced more than 11,500 suggestions, over 300 of which were implemented by Dell (Nelson 2009).

Dell's rationale for its 180-degree turnaround, expressed by Michael Dell himself, is as succinct a description of a cooperative strategy as you'll find anywhere: "These conversations are going to occur whether you like it or not. Do you want to be part of that or not? My argument is you absolutely do." (Jarvis 2007). The argument works because unlike many social media rationales, it fully acknowledges self-interest. Consumers, if left to their own devices, will use social media as a tool of defection more often than they use it as a tool of proactive cooperation (see the discussion of defection cycles in the previous chapter). The marketer can't make that condition go away, but he/she can steer the conversation toward cooperation. And indeed, Dell discovered that when it went out and answered customer objections directly on blogs and forums, the response was overwhelmingly positive.

What I'm calling the "exposed flank" strategy for social media engagement, in which the marketer actually *increases* the opportunities for consumers to provide potentially negative feedback, is a canny form of coordination,

even though it may not appear so at first glance. In the stag hunt, the hopeful stag hunter's best chance to find a cooperative hunting partner is to focus on the territory where cooperation is more likely. The hunter's own territory is a poor candidate; Dell may have thought it was doing a better job controlling the message when its original Dell2One blog was used as a corporate mouthpiece, but that strategy simply moved the conversation elsewhere, outside of Dells' control. When Dell provided a safe place for feedback and signaled its willingness to engage with customers on hostile territory, its risk-taking also signaled cooperation and brought customers to the table.

The consumer's positive response to neutral or friendly territory as a signal for cooperation has to be considered in light of marketer's historical obsession with controlling the message. Outlets for consumer ranting like PissedConsumer.com exist precisely because of marketers' failure to provide such outlets of their own. If one can rant directly to the brand, why rant in a forum the brand might never encounter? Why choose futile grumbling over making a difference? Why hunt rabbit when you can hunt stag? The exposed flank is a way of establishing a meeting place where dialogue – positive or negative, but in the long run, mutually beneficial – can take place.

The exposed flank maneuver is so effective at signaling and gaining mutual cooperation that one wonders why marketers ever bother to fret over the risks of exposure. But in fact the opposite is true: fear of exposure remains the largest barrier to meaningful adoption of social media marketing. Despite deep inroads and interest, most brands remain fundamentally stuck at the sub-optimal 3-3 strategy. This is why the best case studies for the effectiveness of exposing the flank are, as in the case of Dell, the result of brands being backed into a corner. In the next chapter, we'll examine the usefulness of *deliberately* backing oneself into the corner with a "self-command" strategy; for now, let's examine further evidence for the benefits of the exposed flank.

## 5.4 Greenpeace Exposes Its Flank

My favorite story of a successful exposed flank maneuver is slightly outside of the realm of the traditional marketer-consumer game, but its lessons are just as pertinent. In November 2007, the environmental group Greenpeace launched an online competition to name a humpback whale being tracked by Greenpeace to bring attention to the issue of whaling. Among the 11,000 names submitted, most of them somber and dignified, was "Mr Splashy Pants." A site visitor discovered that the contest's safeguard



against multiple votes could be easily overcome, and he or she began voting for the name at a rate of 120 votes per minute, pushing Mr. Splashy Pants to the top of the chart (Nicole 2007).

Greenpeace discovered and removed the duplicate votes, but not before the story had spread to social content aggregator sites like Digg and Reddit and had been picked up by blogs and news outlets. What followed was a spontaneous movement to make Mr. Splashy Pants a legitimate winner, with votes flooding in from across the Web. Mr. Splashy Pants ultimately garnered nearly 80% of the 150,000 votes cast, and he spawned an overnight cottage industry, including a Facebook page, T-shirts, mugs, and bumper stickers.

As Greenpeace's own bloggers later recounted, the name caused some consternation within the walls of the organization; whaling was regarded as a serious issue, and a cartoonish name for its poster whale seemed to diminish the cause. But whatever the internal conflicts, Greenpeace's response has been widely lauded as an exemplary response to the curve balls that social media engagement can throw. The organization came to embrace the name and the publicity that came with it, and today Mr. Splashy Pants is an integral part of Greenpeace's public image and communication on what remains a somber issue (Greenpeace 2007).

Some commentators have chosen to view the Mr. Splashy Pants incident as an example of brand hijacking or "brandjacking," in which social media's openness, e.g., an open voting platform, allows consumers to steer the brand conversation in an unflattering direction. Indeed, brands that seem to take themselves very seriously – and Greenpeace could be accused of that – tend to be more prone to such hijinks. But Greenpeace could have easily defected without much cost, by nullifying the results and choosing a more dignified name. A few voices would have chided them for their stodginess, and then the matter would have been forgotten. Greenpeace's response was clearly in pursuit of the 4-4 payoff; it signaled the organization's willingness to have more fun with a serious cause than its audience might have expected. And the payoff from this risk-taking in terms of exposure to a much wider audience was well worth it.

So both Dell and Greenpeace gained audience loyalty and improved their reputations with an exposed flank strategy. But why exactly does it work? It is, I believe, mainly a matter of transfer of authority – a subject I'll explore in more detail in Chapter 8's examination of the fundamental shift going on in how brands gain dominance. The payoff for consumers involves a far greater say in the brand's marketing and its overall direction, and consumers naturally apply their loyalty where they have the greater sense of ownership.

Neither the marketer nor the consumer can gain the maximum payout by hunting rabbit alone; the marketer can't gain long-term loyalty, and the consumer can't gain power over the brand's content and meaning. By coordinating their moves, both the marketer and the consumer can maximize their long-term gain from the relationship while working less and spending less – an optimal payoff for both.

## 5.5 The Neutral Ground Coordination Strategy

An implicit transfer of authority is also the foundation of the second method of coordination-signaling – the neutral ground strategy. It is no coincidence that the coordination game of nuclear disarmament agreements – or any kind of treaty, for that matter – invariably takes place on neutral ground, never a negotiator's home turf. The choice of neutral ground is a deliberate signal that by voluntarily removing the natural advantages of one's home turf, the player will not defect.

In social media marketing, it has to be said, most marketers give up their home turf kicking and screaming; it may turn out to be the very last aspect of social media that marketers willingly embrace. Why? For as long as there has been a Web, a brand's Web site has been its cathedral – a sacred space filled with iconography and ritual, rich with meaning. Consequently brands exert great control over the experience of visitors to their site. While savvy brands place great emphasis on “user-centered” design – the process of creating a design with the user's needs in mind – ultimately it is the marketer that has designs on the user, leading them through a set of prescribed actions and evaluating the site's success based on those results.

There is nothing inherently unstable in this arrangement; in fact, when brand sites are done well, the encounter between site and user is a mutually understood equilibrium. Consumers have goals when they visit a brand site, and the brand has goals for what it wants from those consumers, and what takes place on the site is essentially an effort to find common ground between those goals.

That common ground nevertheless takes place on the brand's turf – in its cathedral – and that has consequences for how the game is played. Just as with advertising itself, a pair of mutual minimax/maximin strategies are in play: the users will try to find the information they want on the brand site while minimizing their maximum exposure to extraneous marketing messages (e.g., cross-sell), and the marketer will try to maximize that minimum exposure. This is precisely why users develop heuristics – common modes of behavior – that help them to ignore marketing messages on Web sites,

e.g., learning that cross-sell messages appear most often in the site's right channel.

The problem is that external changes can easily disrupt the minimax point, as we saw with the impact of a site like TripAdvisor on hotel marketing, and the marketer begins to lose advantage, increasing the risk of defection. The rapid growth of social networking sites like Facebook constitute exactly such an external change. In July 2009, Nielsen reported that use of Facebook had grown 700% in a single year, with an average 4 hours and 39 minutes of visitation per month per user, making it the world's top Web destination (Nielsen 2009). The time users spend on Facebook is not incremental; it comes at the expense of other activities, like watching commercial television, and visiting other sites, including brand sites. As we saw with the Edelman study, consumers become less reliant on advertisers as information sources and more reliant on other consumers – a direct consequence of more time spent in social networking.

## 5.6 Social Media as Neutral Ground

So a brand faced with a sliding minimax point on their own Web site would naturally look to establish a presence on social networking sites, and this would not be an act of self-sacrifice; the brand would rightly conclude that going where users are is a better strategy than presiding over a cathedral with dwindling attendance at its services. Certainly brands have absorbed this insight enough to have established a beach-head on social networking sites. While Facebook does not track or report on the number of brand pages on its site, the site has noted that 83 of the top 100 brands have established a presence there so far (Morrissey 2009).

But does a brand page on a social network site constitute a neutral ground strategy? It does not. One of the many hindrances to marketer's adaptation of social media is that marketers continue to believe that social network sites are merely a conduit to interactions that take place elsewhere, i.e., the brand's Facebook page is merely a means of driving traffic to the brand site, and not an end in itself.

This damaging notion of a brand social network page as a conduit rather than a destination is actually perpetuated by social media marketing's fiercest advocates, such as former Forrester analyst and blogger Jeremiah Owyang, who persistently argues that traffic from social media to brand sites is a key measure of its worth (Owyang 2009). But these advocates are trapped by a paradox of social media marketing's present state: in order to encourage brands to take the risks involved in reaching for social media marketing's

4-4 promise, one often has to make the case in terms of traditional marketing metrics that brands are conditioned to accept. If you tell a brand that a social network presence will lead to a better long-term relationship with its customers at lower costs, the brand will actually be far less persuaded than if you told them that the social network presence would increase traffic to their Web site by 10%.

This troubling reality was brought home to me in early 2009 when I spoke at a marketing conference about the importance of social media marketing in “de-centering” a brand’s Web presence. I pointed out all the ways that important brand encounters were taking place or could take place online in places other than a brand’s Web site: on Facebook, YouTube, MySpace, etc.

By way of evidence, I pointed to my agency’s own business development efforts. The white papers, webinars, and articles that I prepare as part of my advocacy role for the agency are not squirreled away on the agency’s Web site, but are widely distributed through Scribd, Slideshare, YouTube, LinkedIn, and many other social sharing venues. Not only are the encounters we have off-site quantitatively far greater than anything we could muster on our own site, but more importantly: we’re OK with that. If a prospect reads a white paper and concludes that we are credible on a subject like, say, cross-channel optimization, it shouldn’t matter to us whether the prospect reaches that conclusion on our own Web site or anywhere else. If they decide to hire us, they know where to find us.

In the Q & A session following my talk, an audience member challenged me on my claim that having branded content on sites like Slideshare helped with a brand’s search engine presence by providing more places for people doing Web searches to find the brand’s content. That doesn’t work, he said, because those social sharing sites offer limited ability to link to your own brand site, and inbound links are necessary to improve search engine ranking.

I clarified that I wasn’t speaking of a *site’s* search engine presence, but rather of a *brand’s* search engine presence, because brands needed to embrace the idea of being found in other places than their own site. For my own agency, embracing this idea was a tremendous leap forward: our small Web site stood little chance of making it to the top of search engine rankings on topics like “cross-channel optimization,” but content that we provided to a mammoth site like SlideShare could indeed reach that status. But my questioner remained convinced all paths needed to lead back to the brand site, and what I was proposing wouldn’t advance that goal (it wouldn’t). The encounter left us both frustrated.

The questioner’s mindset, if it remains typical of marketers’ general attitude toward social media marketing, is sufficiently dire in its prospects for

the medium that I have devoted the final chapter to the subject. For now, it's enough to note the creation of neutral ground is a tough nut for most brands to swallow, but that some brands have done it very well, as a few examples will illustrate.

## 5.7 Coca-Cola's Neutral Ground

In April 2009, it was revealed that the second most popular “fan” site on Facebook (after Barack Obama's) was Coca-Cola's, with more than 3 million fans (Gaffney 2009). This by itself is not remarkable, given the ubiquity of one of the world's best-known brands. But the site was not owned or administered by Coca-Cola, but rather by two unaffiliated consumers, Dusty Sorg and Michael Jedrzejewski, an actor and a writer, respectively, living in Los Angeles. The year before, the two had discovered that more than 250 Coca-Cola fan pages already existed on Facebook, but they felt they could build a better one. And so they did.

This particular variety of brandjacking is usually a serious sore spot for most marketers. On most social network sites, any user can register any available name. Because large organizations adapt to trends more slowly than individuals, the last several years have seen a tidal wave of “branded” social network sites that don't belong to the brands – a trend reminiscent of the Oklahoma Land Rush for branded domain names in the mid to late 90's.

Most social network sites take a measured stance on the issue, trying to stay on the right side of trademark law without saddling themselves with a huge policing task: generally if a trademark violation is brought to their attention, they act on it and award the naming rights to the trademark owner. Such a policy is obviously necessary to prevent brands from being hijacked by competitors or by pernicious content that might be wrongly associated with the trademark. But it has also spawned some of the more notorious cases of brand heavy-handedness in social media. When the cable TV network AMC discovered that fans had created Twitter accounts posing as the characters on the hit show “Mad Men” (ironically, a show about traditional advertising), they invoked their trademark rights and forced Twitter to suspend the accounts.

The disastrous consequences of this act of defection will not be surprising to anyone who read the previous chapter; it included widespread negative publicity ranging from the blogosphere to the *New York Times*, a surge of angry backlash on Twitter, and the potential loss of thousands of loyal fans that had been enjoying the imposters' affectionate interpretation of the show's characters. AMC has since relented (Vine 2008).

Coca-Cola's response to the discovery of the unauthorized Facebook fan site was altogether different. They flew the site's creators to the company's headquarters to thank them for their brand evangelism and to ask what the company could do to help. Facebook's rules required that the site be shut down or turned over to Coca-Cola, but the company wanted Sorg and Jedrzejewski to continue their effort, so they agreed to jointly administer without interference. The site still belongs to its founders and remains one of the most popular in the social media space (Gaffney 2009).

Coca-Cola's wildly successful consumer-run social networking site is indisputably a 4-4 Nash equilibrium. At virtually no cost and even less effort, the company gets millions of loyal fans having meaningful interactions in a branded space, and all the positive press attention that accrues to this triumph. Coca-Cola's consumers can interact with the brand, make their opinions heard, and contribute meaningfully to the brand's shape and direction without any unwanted exposure to advertising. Clearly Coca-Cola benefited from having a well-regarded brand in the first place, coupled with the sheer luck of having two enterprising fans at the right place at the right time. But they also played the coordination game just right, recognizing a lucrative neutral ground when they saw it. Consequently, they and their consumers bagged a stag.

I want to be careful not to imply that foregoing corporate ownership is a prerequisite for creating neutral ground, though in Coca-Cola's case it certainly helped. In nearly all cases, large-scale social media marketing programs will be owned by the brands themselves, and this does not at all preclude the use of the neutral ground strategy. Establishing neutral ground is a matter of making consumers feel "safe" to interact on their own terms, and so it can be accomplished in any environment in which consumers reasonably feel a sense of control. Yet it is inherently difficult to impart that sense of control on a brand Web site; since he who pays the piper calls the tunes, consumers may believe – whether it's true or not – that negative comments will be deleted, or that the conversation is being controlled by marketers.

## **5.8 Skittles and the Limits of the Neutral Ground Strategy**

So what does one do to bring "neutral ground" customer conversations into the brand site when so many of those conversations are happening elsewhere? In March 2009, the Mars candy brand Skittles made a bold and short-lived attempt to channel external brand conversations onto its brand Web site. The brand launched a site redesign that entirely replaced the traditional brand material on the home page with live feeds from user-controlled

social media content related to Skittles: YouTube videos, Facebook posts, and most prominently and notoriously, any and all Skittles-related posts on Twitter.

Skittles' experiment was very much in the spirit of the neutral ground strategy, and was indisputably a cooperative move in the best tradition of the stag hunt. The boldness of it set the social media marketing blogosphere ablaze for several days, attracting as many detractors as proponents. One blogger dismissively claimed "Skittles and social media – obviously a company that doesn't get it" (Evans 2009), while another proclaimed "it is a win for the importance of social media and the impact it is having on both commerce and culture" (Chaney 2009). The attraction of the experiment for its proponents was obvious: it demonstrates a brand's willingness to take risks in making cooperative moves, which as I've shown, is essential in prompting reciprocal cooperation.

But the detractors foresaw something that TIT FOR TAT could have also predicted: that the strategy offered no means of responding to defection, which would leave Skittles vulnerable to the "patsy" position. Consumers who wished to defect by posting negative content could do so with impunity, and the content would be prominently displayed on the Skittles home page. And that's exactly what occurred.

Within hours of the new site going live, pranksters began organizing to post Skittles-tagged profanity and negative content on Twitter so that it could be pulled onto the Skittles home page. At least one Web site, "Skittle Fisting," was created expressly for the purpose of organizing the Skittles brandjacking. With its home page littered with profanity, Skittles threw in the towel on its bold experiment 24 hours after it began (Steel 2009).

At first glance, the Skittles incident is as puzzling as it is perturbing. A candy brand should be a safe bet for an experiment of this kind. Skittles isn't exactly a lightning rod for criticism; people either like the candy or they don't. Some commentators noted that Skittles had no business conducting this kind of experiment because they weren't actively using Twitter in the first place. Certainly Skittles would have benefited from more experience, but it's unlikely that a history of Twitter usage by the brand would have prevented the brandjacking, or that the pranksters were in any way motivated by Skittles' historical lack of tweeting.

Skittles' marketers justifiably felt that letting consumers' personal experiences of Skittles stand in for the brand itself was a cooperative move that would be answered in kind. But as I said at the start of the chapter, coordinating cooperation on a massive scale is a daunting task, and it's easy for a few rotten apples to spoil the batch. I included the Skittles story because it illustrates what a successful coordination game is and is not: it is an effort to

signal cooperation by assuming some risk in the interests of greater reward, but it is *not* a subsuming of one's own needs – in this case, the marketers' need to provide reasonable protection to the Skittles brand – in order to achieve cooperation at too great of a cost.

Finally, the incident is also a reminder that the coordination game is not a free-for-all. It is not meant to coordinate *everyone's* actions, but only one's fellow stag hunters, i.e., consumers interested in a brand conversation. Skittles' experiment treated everyone's participation equally – whether you were a lifelong fan of the brand, or a bored teenager who cared nothing for the brand but enjoyed a good prank, you both got a seat at the table. To put it another way, Skittle chose the correct game for increasing brand engagement – the coordination game – but the wrong execution.

But as in any other iterative game, Skittles' experiment provided a set of lessons that others can draw from. As little comfort as the brand may take from this, I believe it represented a vanguard in the movement to tear down the brand cathedral and build outposts in social media venues. It will be followed by other, more advanced and careful efforts to change the sense of “place” that brands occupy online, and consequently, in the world. When marketers acknowledge that cooperation demands a commitment to letting other voices own the conversation, the game has changed for good. In the next chapter, I'll examine this concept of “commitment” and how marketers can use it to make cooperation with consumers more lasting and stable.



## Chapter 6: Crowdsourcing and Schelling's Theory of Self-Command

*ABSTRACT: Coordination games require a degree of self-restraint; marketers face a natural temptation to defect in the interests of short-term gains. Thomas Schelling's concept of self-command provides a potential antidote; a player compels themselves to cooperate by setting conditions that make defection costly or difficult. Successful self-command often involves enlisting a group in the enforcement of norms; brands have done so by opening themselves up to consumer feedback on blogs, but the more dramatic use of self-command occurs in the phenomenon known as crowdsourcing. In crowdsourcing, marketers solicit direct collaboration with consumers in identifying and developing brand assets, customer service features, and even products. Crowdsourcing must operate within certain rules of engagement in order to be successful, but it has the potential to be transformative in marketer-consumer relationships.*

Looking at stag hunt examples like Greenpeace and Coca-Cola from the last chapter might lead one to concoct an easy prescription for success at social media marketing: *just exercise self-restraint*. After all, it's relatively easy (and, let's face it, rather enjoyable) to observe the pratfalls of brands that have blundered into social media and to identify the exact moment when a healthy dose of restraint could have saved them. But as with everything else in social media marketing, it's easier said than done.

There are two problems with a simple self-restraint prescription. The first goes back to the long game of mutual defection that marketers have played with consumers: there is no reason for consumers to trust marketers to exercise self-restraint, because marketers have a poor track record of it. An industry that places advertising messages on airsickness bags on airplanes is not an industry that has built up extra reserves of goodwill, and it is not an industry that is likely to restrain itself from pestering consumers in online forums.

The second problem is a little more subtle but just as important: self-restraint demands nothing of the other player. It is unilateral, which makes it easier to ignore. The consumer won't thank a brand that decides *not* to advertise on airsickness bags, because the act of self-restraint was never recognized to begin with, and it doesn't retrospectively affect the consumer in any way. As Skittles saw, passive self-restraint that asked nothing of consumers put the brand in the patsy position. Pranksters had no compunction about maligning the brand just for fun, and Skittles had established no obligation among its loyal fans to stand up and defend the new territory.

## 6.1 Introducing Self-Command

With such seemingly bleak prospects for self-restraint, it's no wonder that brands so often defect instead; it's sub-optimal but reliable. Fortunately game theory's emphasis on rational self-interest provides a reliable middle path between playing the patsy and defecting altogether. It is the concept of *self-command*, a practice as old as conflict itself, but finally articulated by game theory scholar and economist Thomas Schelling in the 1960's.

Unlike many of his fellow first-generation game theorists, Schelling is an expansive thinker who draws on diverse examples from literature, pop culture, and history to demonstrate his theories. Consequently, these theories lend themselves very naturally to the softer science of social media marketing. His self-command concept is the kind of linchpin idea that, once it is in one's grasp, seems to unlock every kind of social phenomenon around it. To explain the concept, it is best to begin by way of example, as Schelling does.

Schelling draws on the famous scene in *Moby Dick* when a young Ahab, having lost his leg to the whale, is forcibly restrained from avoiding the horrendously painful but life-saving act of cauterizing the stump (Schelling 1982). My own favorite example is less cringe-worthy but just as illustrative: in the *Pink Panther* movies, Inspector Clouseau has ordered his manservant Cato to attack him without warning or mercy whenever he enters the house, in order to keep his combat skills sharp. One of the films' best running gags is that Clouseau always arrives home exhausted and disheveled and attempts to call off the attack, but Cato has been ordered to ignore such pleas.

For those who prefer classical examples, the one that's frequently cited is from Homer's *The Odyssey*. When sailing past the rocks where the sirens lure sailors to their deaths with an enchanting song, Ulysses orders his men to tie him to the ship's mast so that he can hear the siren song without being tempted to dash the ship into the rocks. One can also cite the Roman army's

practice of burning the bridges that would allow them to retreat, thus forcing themselves to fight to the death.

Examining the common features of these examples should offer some clue as to what makes self-command different from self-restraint. In all of these examples, the subject is undergoing a short-term hardship in the interests of a long-term goal: Ahab and Ulysses are both trying to stay alive, Clouseau is trying to sharpen his fighting skills, and the Roman army is trying to win the battle. And significantly, in all cases, *the subject does not trust himself* to act in the interests of that goal if left to his own devices; self-restraint won't cut it. We have no trouble imagining that Clouseau would avoid fighting Cato if he could, because he's actually *trying* to avoid it, and we have no trouble imagining that Ulysses would dash his ship into the rocks, because he's straining at the ropes in order to do exactly that.

Schelling goes so far as to suggest that we think of these scenarios as involving two separate selves: the present self that recognizes that the difficult action will be better for us in the long run, and the future self that is likely to defect from that difficult action in the short run, because it is painful, limiting, or unsatisfying. Because we want to achieve the long-term goal, but we know better than to trust our future selves, we constrain our action in some way that we can't unravel. As Schelling explains:

“What I have in mind is an act or decision that a person takes decisively at some particular point in time, about which the person's preferences differ at the time of action from what they were earlier, when the prospect was contemplated but the decision was still in the future. If the person could make the final decision about that action at the earlier time, precluding a later change in mind, he would make a different choice from what he knows will be his choice on that later occasion.” (Schelling 1982)

Some of the most effective forms of self-command involve a public or social display of commitment, because the social fabric is a big part of what constrains our behavior in the first place, i.e., we care about what our audience thinks of us or how they act toward us. As Schelling explains, “Self-management is not unilateral. It occurs in a social environment.” A person who declares to their circle of friends that they've quit drinking will feel rather awkward about defecting from that commitment and having a drink in front of them. A publicly made New Year's resolution is much more likely to stick than a private one. In such cases, our public declarations work as a kind of verbal contract, with our peers being granted the power to help enforce that contract.

## 6.2 Self-Command in Marketing

Though it's probably fair to speculate that few marketers have ever read Schelling, self-command abounds in marketing. Many loyalty programs are essentially self-command contracts being offered to the customer. An airline mileage program invites the customer to constrain their ticket purchase behavior – even when a given itinerary may be more expensive than other airlines – in order to achieve the long-term goal of achieving status and/or accruing miles.

Premium buyer programs are even better examples, because they exact an upfront cost. My own personal favorite is Amazon Prime. For \$79 a year, Amazon opens its bounty to me, allowing me to obtain any item in a single click with free two-day or flat-rate overnight shipping. But more than that, Amazon is providing me with self-command over my buying habits; I feel compelled to check Amazon's prices for nearly anything I buy, because I am determined to work off that \$79 in equivalent shipping costs. After my spending has surpassed that threshold, I feel even *more* inclined to shop Amazon, because each purchase improves the efficiency of my total annual purchases.

But Amazon Prime is a simple, stated, 1-to-1 contract; for the purposes of social media marketing, I'm much more interested in implied social contracts that have the potential to bind brands and consumers together in stable, long-lasting partnerships. Though these opportunities really came about with the growth of social media, their base ingredient – transparent behavior – is in the DNA of the Web itself. I discussed in Chapter 2 how the Web's penchant for pricing transparency thwarts the traditional zero-sum game of incremental discounts. Transparency works the same way in providing the kind of public exposure that makes self-command effective.

## 6.3 Blogging as Self-Command

When blogging first began to democratize online self-expression in the early oughts, a curious phenomenon arose: the use of the blog as a tool for self-command. In 2003, the *New York Times* noted the rise of the diet blog, in which an individual dieter faithfully records their weight loss progress, or lack thereof, for the world to see. Every perilous encounter with a jelly donut and every pound-for-pound victory are put on display (Harmon 2003). In 2007, the *New York Times* reported a new variant: the debt blog, designed to do much the same as the diet blog, but with dollars of debt instead of pounds. The paper wryly observed, "Public humiliation in the

stocks long ago fell out of fashion. A virtual version is making a comeback” (2007).

Unlike the stocks, though, this humiliation is entirely a matter of self-interest. The blogger openly subjects themselves to scorn (and pleads for support) through a very public commitment to thwart their own short-term preferences. It does beg the question – and it’s an important one for the brand/consumer self-command strategy – what exactly the observer/commenter is getting out of all this. Why read someone’s diet blog, let alone comment on it? There is an element of mutual therapy, to be certain, in the same way that recovery programs like Alcoholics Anonymous encourage open sharing of harmful behavior in a group setting in order to draw on mutual support. But in such cases, the members are all in it together; no one is being singled out. A better answer seems to be the curious but unmistakable satisfaction that people get from enforcing the norms of a group, punishing defectors and rewarding cooperators.

In describing the evolution of cooperation in iterative TIT FOR TAT, I noted the work that’s been done by behavioral scientists at the University of Zurich to uncover a biological basis for the enjoyment we seem to derive from punishing defectors and rewarding cooperators; PET scans show that our pleasure centers are stimulated by this activity, which is a payoff that wasn’t previously accounted for in analyzing cooperative games.

This is also a more satisfying explanation for why people enjoy commenting on blogs than to concede that we’re simply a species of busybodies. But it remains a puzzling phenomenon and heated debate in the behavioral sciences – one that I’ll return to in the next chapter to help analyze why social media participants place so much stock in the opinions of peers. For now I’ll simply stipulate that many individuals appear very eager to lend a helpful or scornful hand to other individuals who wish to engage in self-command online, and that this eagerness is a rich resource that brands will do well to tap.

By now the thrust of my argument concerning social media marketing and self-command has probably become obvious: brands that open themselves up to input from consumers, both good and bad, in social media are engaging in a type of self-command designed to enhance their long-term brand status with these same consumers. The short-term pains, sacrifices, and limitations for brands are myriad; they include loss of brand control, negative associations, and investments of resources and capital without direct return. Such tangible risks are ample cause for brands that are serious about social media marketing to find methods of self-command, i.e., to force themselves to abdicate control in brand conversations that they would instinctively prefer to control.

In fact the simplest and most widely adapted of social media marketing tools, the humble corporate blog, is a prime example of self-command. GM's Fastlane blog, discussed in the first chapter, is a very worthy example precisely because GM was and is such a prime target for scorn; the public has been unabashed in its criticism of GM's cars and its business practices. A company like GM could *only* engage honestly in social media marketing under the auspices of self-command; otherwise the temptation to defect might be too great. One can easily imagine that GM Vice Chairman and Fastlane pioneer Bob Lutz could find a kindred spirit in the bedraggled Inspector Clouseau being mercilessly attacked by Cato whenever he logged onto the blog to answer criticism about the new Camaro. But GM's payoffs have already been demonstrated; the gains in goodwill, customer insight, and the ability to channel scorn that's taking place anyway far outweigh the more painful moments of public excoriation.

Self-command was an especially useful marketing tactic for GM because it made a public demonstration of the company's eagerness to listen. Like other car manufacturers, GM had a reputation for being very top-down with consumers, trying to dictate driving preferences and create trends rather than collaborating. Nothing GM did in social media could be taken seriously without self-command as an opening move, publicly demonstrated by Lutz' willingness to withstand and respond to the good, the bad, and the ugly of consumer feedback. As a cooperative move, it stood the best chance of reciprocal cooperation because it came with the proof that the intent was genuine. That's self-command.

Many companies fear engaging in cooperative activities like open blogging or tweeting because they don't trust the reciprocal aspect; they believe they are merely opening themselves up to abuse. We've already seen how TIT FOR TAT logic makes this outcome unlikely, but self-command goes a step further: *The best self-command contracts don't merely create an obligation in ourselves; they create an obligation in those that enforce the contracts.* In other words, Lutz' openness and honesty obliges his detractors to keep him honest, but also to treat him with respect.

This is a powerful idea at the heart of self-command: it exerts a gravitational pull on those in its orbit to cooperate in kind. As Schelling explains, "The behavior of others depends on what they expect of me; *by restricting my own freedom of choice I gain influence over the choices of others.*" For this reason, I believe the most compelling and effective use of self-command in social media marketing is the concept of crowdsourcing, in which brands consciously restrict their freedom of choice by allowing consumers to shape key decisions and brand attributes. It goes beyond the notion of "conversation," which has relatively low self-command stakes since it merely forces

brands to listen, and adds an element of meaningful change. The brands that allow the conversations to change them in ways that consumers can actually see stand to gain the most in social media marketing, because they have played the cooperative strategy to the hilt *and* they have compelled cooperation from consumers.

It should be acknowledged that, as with all other aspects of social media marketing, crowdsourcing can be done well, and it can be done very poorly, and the distance between those two poles is often a matter of whether the brand is operating a true cooperative strategy or merely hopping on the social media bandwagon in the hope of garnering some free impressions. In the specific case of crowdsourcing, this difference hinges on the how the brand *uses* the content that's been produced. If a brand makes a show of collecting feedback ("Tell us how you'd improve Brand X!"), thanking participants, but never demonstrably using the input to enact meaningful change, no actual self-command has taken place. The logic of self-command is that you compel *yourself* – not the other player – to make tough choices, in order to gain influence over the other player. Enacting those choices with maximum transparency will breed loyalty. The alternative will breed cynicism. Let's look at an example of a brand that's done it particularly well.

## 6.4 Starbucks' Crowdsourcing as Self-Command

As I described, self-command works best when a company needs to achieve a dramatic reversal, because it sends a very public signal that the brand is serious about cooperation. It says something about the shifting fortunes of the Starbucks brand, then, that the company found itself in need of a dramatic reversal in customer perception in 2008. The company has long been a subject of fascination among marketers, growing to \$10 billion in revenue in just two decades with comparatively little advertising and remarkable brand loyalty. It has been able to do so, I will argue, because cooperation is in its DNA.

I've discussed how brands have striven in social media to create "neutral ground" where consumers feel safe; from the beginning, Starbucks literalized this concept in what it still calls its "Third Place" strategy. A Starbucks store is meant to be a third place between work and home, where people can gather and socialize in a space that feels comfortable, familiar, and distinctly non-commercial. By spending little on advertising, Starbucks opts instead to create a very thorough and immersive brand experience in the store, so that the customer feels swaddled in the colors, smells, sounds, and flavors around them.

These multi-sensory experiences are all for sale, of course; you can buy the music, the merchandise, and the coffee. But Starbucks bets on cooperation by deliberately foregoing opportunities to over-monetize its customers; the company has publicly declared that customers should feel free to linger in the stores as long as they like, with no expectation of new purchases (Needleman 2009). (No doubt repeat purchases are helped by the fact that the stores sell an addictive product.) Anyone who has witnessed the habits of Starbucks regulars, many of whom appear to have set up virtual offices in the stores, can appreciate the company's visible commitment to this policy.

But Starbucks' story also illustrates the problem of how far a cooperative strategy can scale. If we posit that a highly cooperative brand is one that excels at maintaining dialogue with its customers in all facets of its business, staying close to that 4-4 equilibrium with its customers, then we also find, not surprisingly, that most brands have a much tougher time doing this once they've grown to a certain size. Indeed, some brands deliberately hold back growth for this reason; cooperation is, in effect, a key commodity that the brand is selling, and this commodity becomes scarce as the brand grows. Organizational complexities make it more difficult to make every customer feel heard, and as new customers are brought into the fold, the experience changes for the original loyalists. Whether the brand is a rock band or a coffee shop, some loyalists will invariably feel that the brand has "sold out" as it grows.

While such deterioration of cooperation is typical, it is not axiomatic. Many large-scale brands – the department store chain Nordstrom's, for instance, or the car rental chain Hertz – maintain a cooperative stance on a large scale. Starbucks did surprisingly well in this regard, given its pace of growth, but by 2007, visible cracks began to appear. It became more difficult to maintain consistency of customer experience across all stores, and the loss of consistency was worsened by Starbucks' continuous experiments in introducing new products into the stores. The introduction of warm breakfast sandwiches provided new revenue opportunities, but it also meant that the multi-sensory customer experience that included the aroma of fresh coffee was now tainted by the smell of fried eggs. Larger, more efficient espresso machines blocked customers' view of the dexterous barista preparing their order.

This is the tightrope that a cooperative brand must walk: these seemingly tiny changes are magnified into a core loss of faith among the faithful, because *the experience itself was the thing being sold*. Starbucks' customers weren't lured to the store by advertising; they were there because of mutual cooperation. Starbucks' internal fretting about this loss of equilibrium was dramatically revealed in a leaked memo from former CEO Howard Schultz



in February 2007. The Schultz memo is a veritable case study on how scale can diminish cooperation, in its cataloguing of the small changes that have eroded the Starbucks experience. Schultz laments “the loss of aroma – perhaps the most powerful non-verbal signal we had in our stores; the loss of our people scooping fresh coffee from the bins and grinding it fresh in front of the customer, and once again stripping the store of tradition and our heritage?” These changes combined with others, Schultz contends, led to the “watering down of the Starbucks experience” (Schultz 2008).

Schultz’s memo and his subsequent return as CEO were the catalyst for widespread changes at Starbucks (alongside the closing of 600 underperforming stores) that were laser-focused on regaining the equilibrium of customer cooperation. Many of these changes are outside the purview of this analysis, but one in particular illustrates the power of self-command. Schultz himself was reputedly the driving force behind the launch of Starbucks’ renowned crowdsourcing site, MyStarbucksIdea.com, in spring 2008. The site’s logic is simple yet profound: any customer can submit an idea for improving the Starbucks experience. The ideas appear on the site, and any customer can comment on them – add, detract, recast, etc. Customers, not Starbucks, decide on the merit of an idea. Ideas fall into categories and 40 “Starbucks Idea Partners” – employees with responsibilities and expertise in those areas – comment on and help shape the ideas. Most significantly, the site features an “Ideas in Action” tab, which documents each and every change wrought by the site’s crowdsourcing engine (Groundswell 2008).

As a self-command strategy, MyStarbucksIdea.com is peerless. It’s a leap beyond open blogging, in which the brand merely constrains itself to respond verbally to customer input, because it holds itself accountable for enactment of ideas that are promoted by the customers themselves. The improvements that Starbucks makes on the basis of the most popular ideas have retroactive justification as true enhancements to the customer experience, as opposed to mere experimentation by Starbucks. But the site also isn’t an exercise in altruism, any more than the Roman army’s burnt bridges are meant to make it easier for the opposing army to slaughter Roman soldiers. Starbucks is still playing to win, because the customer experience *is* the commodity, and that commodity is burnished and enhanced by self-command.

## 6.5 What Crowdsourcing Can and Can't Do

The MyStarbucksIdea site has attracted its share of naysayers in the marketing community, many of whom contend that the exercise, while novel, can’t have much of a material impact on Starbucks’ fortunes. This argument is

worth addressing, since it has bearing on the more general use of self-command as a cooperative marketing tactic. But the logic is flawed for several reasons.

In the first place, those who claim that crowdsourcing should have transformed Starbucks' stock value have a poor understanding of what marketing is and is not. As I've underlined in previous chapters, all marketing tactics are and always have been a means to an end, not an end in itself. In social media, this truism often becomes blurry, because marketing takes place in virtual communities, and communities feel more material than traditional advertising impressions. But ultimately both players are working toward some type of transaction that is external to the marketing itself. The MyStarbucksIdea.com has the short-term impact of signaling cooperation in a virtual space, and it potentially has the long-term impact of driving big changes in the store. But it is *not* the change itself.

Secondly, any single idea that originates from the site may prove transformative, in the same way that the single idea to allow customers to loiter as long as they wished helped to create the stores' culture to begin with. Organizations perpetually struggle with sourcing ideas internally, because the layers of internal politics get in the way. Outside ideas with built-in customer support can go further, especially when the incubator itself has the support of the CEO.

Finally, consider the ratio of cost to payoff, especially in comparison to other marketing tactics. By Starbucks' own accounting, it has a team of six employees to manage all of the company's social media outlets. That's six out of 176,000 employees. Yet the MyStarbucksIdea.com site managed to generate 75,000 ideas in the first six months alone. That volume is manageable because the primary support comes from other customers, which not only saves Starbucks personnel time but also reinforces the self-command message. The site operates on the Salesforce.com CRM platform – a software license cost that's easily within reach of even small businesses. The site's value in earned media alone, in the form of positive press about the initiative, would easily cover its operating cost.

## 6.6 Dealing with Free Riders and Bad Actors

The MyStarbucksIdea.com site has made a powerful impression on marketers in part because it is dramatic in its resoluteness, in exactly the same way that Odysseus' demand to be lashed to the mast makes a powerful impression many centuries after the story was first told. Resoluteness is certainly a requirement; Starbucks' experiment would be less impressive if they only

featured their favorite ideas. But in all instances of self-command through crowdsourcing, the goal is to create a marketplace of ideas, not a bazaar. A cooperative experiment can easily be wrecked by free riders and line-cutters, as previous chapters have shown. For this reason, successful crowdsourcing contains enough constraints to keep the cooperative spirit alive while still disempowering free riders. Generally this is accomplished handily through the community's own self-policing, since as the last chapter showed, punishment of defectors is psychologically rewarding for the other players. But brands have to enable this policing. Starbucks does this by placing the onus of promoting ideas on the community itself. And like other brand communities, it requires participants to signal their willingness to cooperate through registration.

The Web publisher Salon.com enacted its own self-command strategy in 2006 by inviting readers to submit comments on stories directly for publication. The strategy was very effective in building deep and mostly substantive dialogue about Salon's content, which pleased its readers and provided the brand with coveted site "stickiness," in which users spend more time on the site and are exposed to more advertising. A 4-4 win, to be sure, but the elimination of free riders proved necessary to the strategy's success: in April 2007, the site's editor-in-chief announced that the site would require registration in order to "cut down on drive-by insults, off-topic postings and strictly ad hominem attacks." Policing would still be done by the members themselves, but the brand now had the power to act on member complaints to banish bad actors. In such cases, equilibrium is enhanced, not diluted, by a reasonable demand for accountability among members. Salon does not censor its own bad publicity; its forums are still rife with complaints about the site's articles, accompanied by threats to cancel membership. But its detractors are more likely to stick around for the invigorating debate despite these threats.

## 6.7 Schelling's Focal Point

What I've tried to describe in the preceding pages is the potential contribution of Schelling's self-command concept as a method for brands to coordinate cooperation with consumers. Self-command is only one of two important contributions Schelling made to the conduct of coordination games; the other also has some relevance to social media marketing, albeit as a description more than a prescription. Like self-command, it is a light-bulb theory – one that, once it's grasped, seems to illuminate everything around it. Schelling contributed the idea of the *focal point*, popularly referred to

as the “Schelling point” in his honor, to describe the places where non-communicative players trying to coordinate their action naturally converge (Schelling 2006).

In a literal stag hunt, in which neither hunter knows precisely where to meet the other, they might naturally choose a rocky promontory, hoping to find the other there. In Schelling's own research, he found that study participants who were asked to choose a focal point in New York City most often chose Grand Central Station at high noon. In my own informal experiments with focal points in my city of residence, Portland, Oregon, I found that respondents most often chose Pioneer Square, affectionately known as the city's “living room.”

At first glance, the focal point appears to be nothing more than a matter of choosing an obvious meeting spot, but there's more going on here. As with other coordination games, one's main focus must not be on one's own preferences, but on the anticipated preferences and moves of the other player. It requires a certain effacement of self-interest in order to achieve what is ultimately in one's own interests, as we saw with both the stag hunt and with self-command. Moreover, it requires some anticipation of how the other player is thinking about *you*, since they are also trying to coordinate the most efficient outcome. A successful Schelling point is therefore also a Nash equilibrium, because it represents the best chance of success for both sides in the absence of any foreknowledge of the other player's moves.

I am interested in the ways that social networks might function as focal points, because consumers show increasing preference for them both as places to spend time online and as places to interact with brands. Clearly brands like Dell and Coca-Cola that focus on working within existing social networks grasp the importance of this focal point. Yet despite reams of customer data at their fingertips, marketers often struggle with coordinating focal points on behalf of brands and consumers.

Why? Because self-effacement is hard for large organizations to undertake; brands have an innate preference for wanting to get customers onto their home turf, as I noted in the discussion of neutral ground. My argument with the seminar attendee about whether customers ought to be forced to interact with branded content only on one's own Web site is a perfect example of how a misguided coordination strategy will fail to find its focal point. Maintaining branded content in places where prospective customers choose to spend their time is an optimal outcome, once the marketer abandons the illogical proposition that their own Web site should be the consumer's focal point as well.

Try this thought experiment: imagine a brand's online advertising as an attempt to organize a coordination game by moving consumers from one

place to another whenever a consumer responds to an ad. (Based on my previous analysis of click-through-based advertising, you can see where this is heading.) The marketer does her research and identifies where the target consumers are spending their time. The marketer then shows up at that location, waving a sign and shouting ad slogans, and tries to induce as many consumers as possible to follow her to an undisclosed location.

Not surprisingly, out of 1,000 consumers gathered there, only five choose to follow him. The rest go about their business, and the ranting marketer is quickly forgotten. When the marketer and his followers arrive at the new location, four of the five consumers decide it wasn't worth the trip, and they leave. The sole remaining consumer decides to stay and converse with the marketer. Later, the marketer proclaims this dismal failure to be a smashing success.

When viewed in the harsh spotlight of coordination game strategy, click-through based advertising is glaringly ineffective. But that's not the point of this allegory. The point is that the marketer might have accomplished something altogether different had she found a way to stay in the place where consumers gathered. What Schelling's focal point demonstrates is that achieving true equilibrium demands a rethinking of location, in the most literal sense. Advertising has always been based on interrupting people while they're doing one thing – riding the bus, watching TV – and getting them to do or think about something else. The inefficiencies are obvious and unavoidable, but traditionally there have been few alternatives.

Social networks haven't utterly transformed that reality, but they have at least presented more alternatives. They are increasingly where consumers spend their time online; a recent study by Comscore showed that 20% of all online ad impressions occur on social networks. Those ad impressions stand a far greater chance of success if they allow users to remain on those networks and visit focal points set up for them there, but advertisers haven't caught up to this insight; most ad impressions served on social networks take the user off the network.

This practice is quite simply an evolutionary lag, an unconscious settling for a sub-optimal solution. It does not mean that a network like Facebook is any kind of permanent focal point; rather, it means that focal points will continue to shift with consumer preferences and habits. Marketers' own habit of becoming overly invested in maintaining a single solid presence, a brand cathedral, may ultimately cost them the loyalty of consumers who prefer to worship elsewhere.

This chapter has attempted to show the lengths to which marketers may need to go in order to succeed at the coordination game, including the use of self-command – with consumers as willing enforcers – to ensure their

own cooperation. These lengths are the price paid for marketers' historical dominance of the brand conversation; for brands to be allowed into the social media party, where consumers now dominate, they must be willing to check their weapons at the door. Some of the world's most dominant brands – Dell, Starbucks, Coca-Cola, and many others – have absorbed this lesson and reaped the rewards.

But self-command is just one tactic within a larger coordination game in which brands must rapidly adjust to new ways of building consumer relationships. As I'll discuss in the next chapter, traditional advertising is in no danger of disappearing as a means of brand-building, but the strength of its signal has faded. That signal will be replaced by new forms of signaling that are inherently more cooperative, and therefore more rewarding, but also trickier to manage and maintain. Marketers that resist augmenting tried-and-true, albeit sub-optimal, dominance of paid media with the uncertainties of social media may take some cold comfort in the fact that consumers will leave them with no choice in the matter.

## Chapter 7: Content Popularity and Spence's Theory of Costly Signaling

*ABSTRACT: The economist Michael Spence's groundbreaking work on costly signaling in the job market demonstrated how advanced degrees could serve as an accurate signal of candidate ability, because more qualified workers could acquire the costly signal at lower cost than unqualified ones. External forces, like the proliferation of MBA programs, can devalue a costly signal over time. Marketing is undergoing such a shift in its signaling system. In traditional advertising, the high cost of media exposure signals legitimacy, irrespective of content. But the Web itself has introduced disruptions into this traditional costly signal, as entities like Google have made the popularity of content a condition of exposure. Social media marketing extends popularity-based signaling into a systemic form, in which marketers must learn new rules for gaining exposure. This has provided new opportunities for upstart brands, as well as significant disruptions and adjustments for many traditional brands.*

In the last chapter, I discussed the effectiveness of self-command as a means for brands entering the social media space to signal their willingness to cooperate – a critical need, given the long history of mutual defection. I stressed the importance of very public self-command demonstrations because of their potency as a signal. Self-command not only constrains the marketer from following the impulse to defect, it also signals cooperation to the consumer, creating a sense of mutual obligation around the success of the relationship. Starbucks' crowdsourcing experiment demonstrated that the brand was willing to listen to consumers, but it also obliged consumers to take the brand's efforts seriously and to contribute cooperatively to the chain's effort to better itself.

As a form of signaling, self-command is also a strong show of confidence. Consumers might conclude that only a brand that has historically cooperated and/or is serious about cooperating in the future would take the

risks involved in crowdsourcing its path to improvement. So on the basis of the crowdsourcing move alone, *regardless of the actual content of the crowdsourcing experiment*, Starbucks earned dividends for its brand.

Historical examples of self-command also demonstrate its effectiveness as a way of signaling confidence. When the Romans burned bridges, they weren't merely signaling their willingness to fight to the death; they were signaling that they had the strength and confidence to risk a fight to the death and to prevail. The opposing army would reasonably conclude that a weak Roman army would not take the suicidal move of cutting off their own escape routes; these were clearly soldiers with a high degree of certainty about the future outcome of the battle.

For Starbucks and the Roman army, then, we can conclude that they each calculated the *cost* of their risk-taking to be less than the expected return, and this conclusion about their calculations is a type of information that gets transmitted to their opponents; it is a *received signal*. We might further conclude that other armies that don't burn their bridges – that don't publicly pay the risk-taking cost – are weaker than the Roman army. Maybe they're not, in actuality, but the Roman signal is a powerful one, so it colors the opposing army's perception. And we might conclude coffee chains that don't engage their customers in crowdsourcing are less cooperative than Starbucks. All of these information signals are made available to the other player(s) in the simple act of self-command.

## 7.1 The Theory of Costly Signaling

By describing self-command as a signal with a particular cost attached to it, I am seeking to broaden the discussion about coordination games to include the whole science of signaling – a fascinating adjunct to game theory. This science features yet another luminary in our line-up of game theory's Nobel Laureates – Michael Spence, whose concept of job marketing signaling, introduced in 1973, had a great impact on the field of economics.

As in the examples above, Spence's work has been particularly focused on how signals convey information in *asymmetrical* games, i.e., when one party can't directly know everything they need to know about the other party in order to make their best move. If the Romans fail to convey to their enemy their resolve to fight with no chance of surrender, then needless slaughter will ensue.

Spence's work on signaling focused on the knowledge asymmetry between new job seekers and employers – a scenario rich with the kind of conflicting and overlapping interests that game theory thrives on. The



employer wants to find the most productive candidate while *minimizing* their payout, and the candidate wants to find a job while *maximizing* the employer's payout. But the employer can't empirically know how productive the candidate will be prior to the hire, and so an unstable conflict ensues.

In Spence's analysis, education is the stabilizing factor that allows the two parties to achieve equilibrium. In pursuit of a better salary, the candidate goes out and gets an advanced degree as a way of signaling their abilities. The cost of acquiring that advanced degree, Spence reasons, is much higher for a low-ability candidate than a high-ability candidate, because the low-ability candidate struggles and risks either being unable to complete the degree or receiving poor marks. (Note that when we speak of higher cost for lower abilities, it includes the cost of time, commitment, emotion, etc, on top of hard costs). A high-ability candidate stands a better chance of making the advanced degree work for them, and so more high-ability candidates would complete advanced degrees.

The employer will have to pay more to acquire an advanced-degree candidate, but doing so is worthwhile, because hiring and firing low-ability candidates is more costly in the long run. So the advanced degree provides an equilibrium point for the employer. It does the same for the candidate: the high-ability candidate that acquires an advanced degree will be able to pay off the cost of the signal by getting and keeping a higher-paying job. A low-ability candidate will accept less risk in acquiring a costly signal, but they can then accept a lower-paying job in return (Spence 1973).

Let me underscore a few aspects of Spence's model before making the leap to its application to marketing. Most importantly, it is education's value as a *signal*, rather than the content of the education itself, that creates the equilibrium. While it is assumed that the advanced-degree candidate also learns a few things that are relevant to the job, that learning is not at all important to the success of the model. What is important is that high-ability candidates are more likely to consider advanced education to be worth the risk than low-ability candidates, and that probability allows the employer to accept the signal as valid.

The second point is that we are speaking of *probabilities* here, not certainties. The educational signal doesn't guarantee high-ability candidates; it merely creates a pool of candidates with a greater likelihood of being high-ability. A few duds will always slip through the cracks.

The third point is one that may have already occurred to you as you read the description of educational signaling, because we live in an era in which the value of advanced degrees has become highly unstable. The last 20 years have seen a proliferation of MBA programs, which means that less qualified candidates have a better chance of entering the pool at lower risk/cost. The

programs have proliferated precisely because institutions recognized that the MBA-signal had become a common short-hand among recruiters, and they could benefit from candidates' desire to acquire this signal.

Like all such instances of saturation, this proliferation of MBA programs has had a deflationary effect on the value of the signal. In a 2007 survey of corporate recruiters by the Graduate Management Admission Council (GMAC), 37% of recruiters cited the inconsistent quality of MBA candidates as a barrier to hiring, and 36% cited unrealistic salary expectations – two very clear indications that the traditional costly signal for MBAs has begun to deteriorate (GMAC 2007).

Spence's model depends on the ability of both players to assign a stable value to the signal; when external factors (such as a flood of unqualified candidates, or rising salary demands) disrupt the signal, a new equilibrium must be found. This could consist, for instance, of an emphasis on tougher, more exclusive advanced-degree programs, so that the value of the signal rises again for both parties.

## 7.2 Traditional Advertising and Costly Signaling

The leap I wish to make may be obvious by now: traditionally, advertising has functioned as a form of costly signaling. Again, as with education, the signal is clearly not advertising's only function, but it is an important one. In the case of traditional advertising, the signal is *legitimacy*; the mere presence of a given brand in a high-cost media venue signals the brand's prominence within its competitive marketplace. The content of the ad itself is vastly less important. A new luxury car that appears in a high-end, glossy magazine may not be reaching its audience in the most cost-efficient way, but the act of wastefulness itself can create an equilibrium that isn't measured in dollars. The wastefulness signals the car's suitability to its high-end audience, regardless of the quality of the car.

Just as costly signaling is not merely convenient but essential for the recruiter faced with candidates of unknowable productivity, consumers come to rely on costly signals quite heavily in the absence of other information. If a cure for baldness were touted in a quarter-inch text ad in the back of a tabloid newspaper, it would not be taken seriously by most discerning consumers; indeed, one could conclude that it appears where it does precisely because the advertiser is trolling for *non*-discerning consumers. But if this very same product were advertised in a half-page ad in *Smithsonian* magazine, it would be bound to attract some attention and at least initial inquiries, even if the content remained the same.

You might object that such a system would be easy to game, since the baldness cure scam artist could recognize the potency of the costly signal and gamble his entire ad budget on the half-page *Smithsonian* ad. But under Spence's theory, the scammer is unlikely to do this, because the costly signal is going to be *more* costly for him than it would be for a legitimate baldness cure. Why? Because he likely has only one shot at the signal before the scam is uncovered. The magazine, because it wants to preserve the value of the costly signaling that appears in its pages, since it translates directly into ad revenue, will try to screen out the scammer in the first place, and will certainly bar him from subsequent issues after readers complain. In this way, a stable costly signal is preserved for the signaler, the signal recipient, and the media outlet.

But costly signals *do* break down, as in the case of the MBA program, because the signal equilibrium is based on repetition of a pattern, and every pattern can be exploited. Recruiters didn't wake up one day and decide that advanced degrees might correlate to productivity; they observed this correlation over huge data sets and long periods of time before the short-hand signal could be deemed useful. As previous chapters have shown, patterned behavior is always exploitable in non-cooperative games: this was demonstrated in both zero-sum games like "head or tails" and in the iterated Prisoners' Dilemma. Once the signal has been sufficiently exploited, it is no longer *true*, i.e., it no longer stands in for the probable existence of the material reality it's meant to stand for. In other words, these new MBA graduates may not be as productive as they look on paper.

### 7.3 The Erosion of Costly Signaling in Super Bowl Advertising

Signal distortion appears to be eroding one of the great costly signals in modern advertising: the coveted Super Bowl ad spot, the most expensive 30 seconds of commercial advertising. For established brands, the Super Bowl spot is less a signal of legitimacy than of ongoing category dominance; traditionally, we could count on Pepsi and Coca-Cola to both make an appearance each year, because failing to do so might signal a loss of stature and cede the field to the rival brand. Befitting its status as a signal, Super Bowl advertising has never been about helping consumers gain knowledge from its content; all of the knowledge transfer is from the signal itself. A revealing study of consumer responses to Super Bowl advertising by Scott W. Kelley, Professor of Marketing at University of Kentucky, showed that 48% of the Super Bowl ads that consumers liked the least also happened to be the ones that relied on rational appeals (Kelley 2002). Consumers want

to be entertained; they are far less interested in learning from the content of the ads.

The legitimacy signal that the Super Bowl conveys is important to new brands, and this was most evident in the mad scramble among dot-com advertisers mentioned in the first chapter – the so-called Dot-Com Bowl of 2000. In that case, the *primary* reason for the Super Bowl ad among the 17 dot-coms was to convey legitimacy; one could not possibly rationalize the move by any measure of direct monetary return. Since the advertisers were online businesses, it requires little analysis to project that \$2.2 million spent on online advertising would have provided a far better return on investment, if indeed direct return had been the goal (Elliott 2000). But online advertising could not and still cannot convey the sought-after legitimacy; it is not costly enough.

Given that very few of Super Bowl XXXIV's dot-com advertisers actually survived the looming dot-com implosion, one might argue that the advertising did a very poor job of signaling legitimacy. In such cases, we expect the signal's receiver to gradually reject the signal and seek new ones, in the same way that recruiters must now turn a jaundiced eye on some of their MBA candidates. And indeed that seems to be occurring, for reasons that go beyond the dot-coms' distorted legitimacy signal. The Retail Advertising and Marketing Association's *2008 Super Bowl Consumer Intentions and Actions Survey* showed that only 9.2% of consumers who viewed the ads felt more likely to purchase the products.

In the current recessionary climate, a costly signal can backfire. The cost of a Super Bowl ad has risen to an estimated \$3 million, and advertisers that splurge on a spot may find themselves sending another signal entirely. The same survey referenced above showed that nearly one out of five Super Bowl viewers felt that the advertisers should have avoided the expense of the ads and passed the savings on to consumers. These consumers have now spent nearly two years watching bloated financial institutions fail; in this climate, excessive costly signaling may signal a bloated institution, ready to fail or at least worthy of failure. If the cost of the ad spot rises while consumer opinion deteriorates, the costly signal will become ineffective for both players, and the entire Super Bowl advertising system will face a severe reckoning.

That reckoning appears to be underway; Pepsi recently created an uproar in the advertising community when it announced that it will not run commercials during the 2010 Super Bowl, and will instead spend the money in online advertising (Vranica 2009). It's hard to conceive of a starker symbol of the loss of costly signaling in traditional advertising. The effect of Pepsi's radical reversal is sure to reverberate, as its rejection of the costly signal will make it permissible for other dominant brands to do the same. But from

a game theory perspective, Pepsi's decision is entirely rational; the cost of the signal has exceeded its return, and Pepsi has correctly calculated that the same brand effects can be achieved online at greater cost efficiency.

## 7.4 How Web Transparency Disrupts Costly Signaling

Thus far I have been speaking of costly signaling's function in traditional advertising only, in order to convey how it can work when the system is relatively stable. But I am mainly interested in the destabilizing effects of the Web in general and of social media in particular, and what kinds of signals might now be emerging in these new media. In order to get there, I need to start with an explanation of how knowledge transfer works in costly signaling.

Just as in the case of the recruiter and the candidate, costly signaling works best where the information gap is largest, i.e., when the receiver simply can't get tangible information by any more efficient means. If the recruiter could get very reliable information about candidate productivity without the signal, they would seek a more optimal equilibrium by getting rid of the MBA requirement altogether. They could then afford to pay less by choosing from a pool of candidates that didn't come with a costly signal; they could simply choose the brightest and the best using this new method. MBA graduates would be unhappy about this development, because their costly signal would be devalued, and they would suddenly be competing *for lower pay* against candidates that hadn't paid the costly signal but were nonetheless judged to be productive. But the market overall would be open to more players during this readjustment.

Sounds familiar? The democratization of access to content, or to choose a more succinct term, the *transparency* of the Web, plays hell with traditional costly signaling because the information gap between marketers and consumers has narrowed. This does create opportunities for more players – including lesser known brands – even as it creates disruptions in the transfer of information. As we saw in the case of the hotel chain and TripAdvisor in Chapter 3, the consumer seeking a resort vacation is now far less reliant on the costly signal of advertising to convey that a luxury hotel is indeed luxurious; he or she can find substantive proof or disproof just by reading the comments and ratings of other travelers. In the long run, this is a good outcome, a healthier equilibrium, for the hotel chain too, because they can convey information to the consumer at a lower cost. But in the short run, it's bound to be difficult, for reasons described in Chapter 3: the hotelier now has to learn how to play an entirely new game, to cultivate good reviews

through customer service, and to be active and engaged in the places where consumers freely share this formerly precious knowledge. Costly signaling was indeed costly for the hotelier, but it was reliable.

The long-term loser in this system evolution is, of course, the glossy magazine that relied on the hotelier's costly signal for ad revenue, and it is no coincidence that the rise of social media in the last three years has occurred alongside massive downturns among magazine publishers, with high-end magazines taking the worst of it. In October 2009 alone, the magazine publishing giant Conde Nast announced the closure of its *Gourmet*, *Cookie*, *Modern Bride*, and *Elegant Bride* magazines, as well as layoffs and cutbacks at *Glamour*, *Wired*, *Lucky*, *Bon Appetit*, *Details*, and *Architectural Digest*, among many others. The reason for the layoffs: declining ad revenue. The reason for the reason: the diminishing effectiveness of costly signaling in an era in which consumers' focus is increasingly online. In perhaps the supreme irony of the publishing collapse, *Wired* editor Chris Anderson was reportedly absent on the day that his magazine laid off six key staffers, because he was busy promoting his new book, *Free*. The book's subject is the proliferation of free content and services that undermine traditional paid business models (Tate 2009).

## 7.5 The Evolution of Costly Signaling on the Web

But as Spence showed, costly signaling itself is a constant, even as its terms may evolve; it gets disrupted by external factors (like the growth of the Web and, more recently, social media), but ultimately a new equilibrium takes hold. The "costly" part of costly signaling is almost never a matter of pure capital; it is comprised of all of the efforts the signaler must make to convey a particular status to the receiver. And for equilibrium to occur, these costs must be worthwhile, i.e., they must produce a return.

Thus when one speaks of the democratization of access to content on the Web, a certain reality check is needed before the breaking out of the guitars and the strumming of "Kumbaya." The primary currency used in costly signaling is evolving rapidly, but there is a currency involved nonetheless, which means that some players will gain more access than others. The current anxiety among marketers as to how to "monetize" social media marketing springs from this uncertainty about how to make it provide a reasonable rate of return as a costly signal.

I will argue that we are in a highly disrupted period, with a very high signal-to-noise ratio, as marketers and consumers attempt to figure this out. But I also believe the new currency for costly signaling is beginning to stabilize.

But before I can make this case, I need to briefly trace the evolution of Web-based signaling, starting with search engines.

## 7.6 Google Changes the Costly Signaling Currency

While I would prefer to avoid offering up another account of How Google Changed the World, I can't avoid the subject of Google entirely, because it has been one of the prime movers in the shift to a new currency for costly signaling. In fact, Google is a worthy subject both for its *contribution* to the shift in costly signaling and as a *practitioner* of a distinctive new form of costly signaling in the development of its own brand. I'll consider these subjects in turn.

The Google search engine debuted to the public in 1998 and rapidly gained ground on other search providers like Microsoft MSN and Yahoo. Significantly, both Microsoft and Yahoo enjoyed their previous dominance in the burgeoning search engine space primarily because both were and are portals, i.e., destinations that aggregated and privileged a variety of life-style content to provide users with a kind of all-in-one Web experience. In other words, their search engines were popular less because of any inherent qualities than because users were on their sites anyway, so using the search engine was convenient.

Google was and is conspicuously *not* a portal, but rather a stunningly simple and fast-loading search query box, alone on a page. This difference alone is important to the evolution of costly signaling, because it meant that those first users in 1998 who chose to abandon the portals and do their searches on Google were moving away from a reliance on pre-sorted, privileged, ad-supported content on portals – more closely aligned with traditional content publishing models – and toward less privileged, more open-ended access to content.

But of course, search engines *do* privilege some content over others as a matter of practicality, by means of the ranking of content results that are displayed after a user conducts a search. Since users most often choose their content from the first page of search engine results, this ranking is all-important; it is in no small way the user's experience of the Web itself, outside of their regular destinations. But in this respect as well, Google represented a break from the usual way of doing things.

Traditionally (to the extent that a brand-new medium can be said to have a traditional mode), search engines focused mainly on the relationship between terms that users search on and the density with which those terms appear in the page content, as a way of assigning privilege or authority. A

page with a lot of content about alternative fuels was deemed by the engine to be more valuable than a page with very little content, and so on.

This would be a reasonable system if one were dealing with traditional content. Suppose that you take all of the books in a university library, scan their contents, and then make the whole library searchable with a simple content density search engine. The search results would be a reasonably good reflection of a given book's authority on a given subject. Why? Because costly signaling has already taken place.

In order to make it into the university library, the book has passed through at least two filters: first the publisher, then the library. A significant level of authority had been conferred on it before it was ever made searchable. In a Web search, by contrast, the search engine is the first and only authority filter the content passes through before it reaches the user. In a non-hierarchical Web structure, there is no *inherent* distinction between library-worthy content and a random assortment of words generated by a machine. While this greatly democratizes access to content, if the search engine can't provide some level of qualitative filtering, democracy simply leads to chaos.

The obvious problem here is that repetition of keywords has nothing to do with a site's authority on a given subject, or its usefulness to users. There is also nothing costly about a keyword-based ranking; nonsense pages with the right keywords would be given unmerited authority at no cost to the signaler and no value to the receiver.

On a more pragmatic level, a mostly keyword-based search algorithm would make it relatively easy for competing players to introduce noise to disrupt the signal. An oil company, for instance, that wanted to dampen the discussion of alternative fuels could create a page saturated with alternative fuel-related terms solely for the purposes of discrediting the subject; they would effectively dominate the user's access to information on the subject, even though they are more interested in disrupting a signal than sending one, by *preventing* users from learning more about alternative fuels from other players.

The oil company arguably *does* deserve a seat at the table in the discussion of alternative fuels, if they have worthwhile content to contribute, so long as other players are also able to signal their authority on the subject in a way that allows the receiver – the end-user conducting the search – to uncover these perspectives. But how is such assigning of authority even possible in a system that consists of some 112 million Web sites – each with multiple signals – and hundreds of millions of searches each day?

Google's answer to this infinite-monkey problem is to introduce *popularity* as a major currency in its costly signaling requirements. In an effort to build an algorithm that more closely aligned with what a given user is



actually looking for, Google assigned weight or authority to the content's popularity, i.e., how important other Web users judged the content to be.

Google discovered that the number of human-generated links to a given page is a good indicator of how much authority the page had on a subject, and so it weighted those inbound links in assigning a rank to the page. In this way, for instance, an oil company couldn't drown out the discussion of alternative fuels, because users with first-hand authority on the subject would link to alternative fuel providers and/or forums, and those sites would rise in search engine results ranking accordingly.

I should note that Google's actual search algorithm, known as PageRank, is a closely guarded secret, precisely in order to dampen attempts to game the Google system and introduce noise into the signaling system that is Google's reason for existing. Inbound linking is ostensibly one of many factors that Google uses to assign popularity. But Google's own description of its technology acknowledges this much:

PageRank also considers the importance of each page that casts a vote, as votes from some pages are considered to have greater value, thus giving the linked page greater value. We have always taken a pragmatic approach to help improve search quality and create useful products, and our technology uses the collective intelligence of the web to determine a page's importance.

Google's description is even more revealing as a kind of credo for a popularity-based currency of costly signaling than it is as a description of search technology. A system that "votes" for content on the basis of its popularity with end-users has vast potential for upending traditional marketing systems for costly signaling. To return to my seemingly inexhaustible example of the luxury hotel, it finds itself, under this system, unable to drown out the signals being sent by the review site TripAdvisor, since the site is enormously popular and rich with user-generated content; it has no choice but to participate in the system by working to improve its stature within TripAdvisor.

Well, *almost* no choice. Every system can be manipulated simply by virtue of its "systemness," i.e., systems have rules, and rules can be gamed. In the early days of Google's rapid rise, a common trick for gaming the system was to sign up with "link schemes" whose sole function was to create heavily weighted inbound links to boost ranking – a simulacrum of popularity that had nothing to do with the judgment of actual users.

But in order to protect its stake in the stability of popularity as a costly-signaling currency, Google manually seeks out and de-ranks link schemes, urging sites to follow the virtuous but arduous path of simply making better

content that can, in turn, become more popular: "The best way to get other sites to create relevant links to yours is to create unique, relevant content that can quickly gain popularity in the Internet community." In other words, pay for the costly signal.

## 7.7 Paid Search and Popularity-Based Signaling

Easier said than done, but Google offers another antidote, in perhaps the most lucrative monetization scheme in the Web's history. Signalers who don't wish to pay the popularity cost can cut to the front of the line by paying a monetary cost. Google's AdWords program allows marketers to bid on keywords in an elaborate pay-per-click (PPC) auction system. The top ranking paid keywords then appear alongside the "natural" (popularity-based) search engine results. The system thus allows advertisers to pay to acquire relevance or popularity rather than to build it organically.

Such a system could easily undermine popularity as the coin of the realm for costly signaling, were it not for certain safeguards. Most importantly, users can visually distinguish paid results from organic results quite easily, which allows users to keep faith in the costly signal being offered by the organic results. As one would expect, users prefer organic results, clicking on them over paid results at a rate of 9 to 1; this behavior is consistent with Web user's overall privileging of popular over commercial content, whenever the two types are in competition.

Google's other safeguard is quite revealing when viewed through the lens of costly signaling: Google actually demands relevance from its *paid* links as well. Google Adwords participants must demonstrate the relevance of their site content to the keywords they're bidding on; linking to tangentially related or unrelated content (e.g., an oil company grabbing up alternative fuel-related keywords) is not allowed. And Google layers on a popularity standard as well: a marketer receives a "quality score" for a campaign based on the campaign's popularity (measured in clicks), so that an advertiser with popular content achieves a higher ranking at a lower cost.

It is worth taking a moment to ponder the significance of this seemingly mundane detail in Google's paid search program: for the first time in the history of advertising, marketers gain access to consumers based in part on the *popularity* of their content with consumers. This is nothing short of a sea-change in the marketer-consumer relationship.

It would be unthinkable for TV networks to allow access to coveted Super Bowl slots based on the popularity of the advertiser's previous ads, or to give an advertiser a lower rate because consumers enjoyed their ads, but

that is in effect what the Google model does. Doing so is, in fact, essential to Google's success: if users decide that top paid search results are irrelevant, they'll desist from clicking on them, and the whole system will go to pieces. By developing and stabilizing popularity as a currency in costly signaling, Google and other paid search engines programs served as a catalyst for the even bigger sea change now taking place, in which marketers' participation in social media is a form of costly signaling that must increasingly be paid in order to have access to consumers.

So it is clear that search engines offer a form of costly signaling, but does it work? It does. Achieving a high search ranking is not merely a matter of access to consumers for the purposes of persuading them to click; it also sends a signal about the brand. A 2006 study by Jupiter Research and search engine marketing firm iProspect showed that 36% of consumers regard the brands that appear at the top of search engine rankings to be the top brands in their field (2006). And that by itself is not surprising or new: we think of Coca-Cola as a top brand in its field in part because we see it everywhere. The difference is that in the new model, the brand had to, in part, *earn* its top exposure by being popular in the first place.

## 7.8 Noise in the Signaling System

Before setting aside the subject of search engines and taking up the prevalence of costly signaling in social media, I need to acknowledge that, just as in all other forms of costly signaling, there is noise in the system. If popularity can be manufactured rather than earned, then less costly payers, like candidates with mail-order MBAs, can slip in and disrupt the system. And search engines have their share of mail-order MBAs. In certain industries, particularly travel and mortgage, "lead aggregators" occupy many of the top search results for the most common keywords. These companies do not provide mortgages or trips to Paris themselves; instead they capture contact information from consumers interested in these things. They then sell the leads to the actual providers, e.g., banks or travel agencies, who would otherwise have dominated the search engine results, were it not for the existence of the lead aggregators.

It is a fit subject for debate – a debate that can't be fully explored or resolved within these pages – whether a lead aggregator truly pays a costly signal. Unquestionably they are providing a relevant service to the consumer. If the consumer is seeking a mortgage, and the lead aggregator introduces the consumer to three lenders in response to a single submission by the consumer, it is arguably a better deal for the consumer than having to

gather multiple competing offers on their own. But the lead aggregator gains their status not on the basis of *actual* popularity, e.g., consumers “voting” by linking to their site, but rather by virtue of thousands of paid links (affiliates) that drive consumers to the aggregator site, thus boosting the appearance of popularity to other consumers. It’s a rather perfect self-perpetuating system.

The aggregator also has the advantage of having little or no brand equity to protect. They are interested in a one-off transaction with consumers, not a relationship, and it hardly matters whether the consumer remembers the name of the aggregator afterwards; the actual service is being provided by someone else. When a system like this works for the consumer – who gets multiple competing offers – and it works at least begrudgingly for the marketer, who can pay for the lead at predictable cost, then it has the long-term effect of diminishing the value of the marketer’s efforts toward building a brand. Branding still matters, but it matters less if the marketer is buying the lead than it would if the marketer were trying to attract the lead on their own. The net effect is to change the playing field for traditional brand-building as a form of costly signaling; simply put, being Wells Fargo means slightly less than it did before LendingTree.com learned how to play the search engine game.

The aggregator scenario is just one example of the headaches that large brands face in trying to translate the ubiquity that they’ve bought and paid for in the traditional media space into a medium where ubiquity is far less straightforward. Wells Fargo could not effectively advertise their way out of the scenario above, in which a certain volume of their new customer leads is going to have to be purchased from a much, much smaller brand – the lead aggregator – that has learned to play the costly signaling game of popularity. Since the aggregator is not *truly* a more popular brand than Wells Fargo, but simply better at simulating popularity in a way that tricks the search engine, one could conclude that this scenario constitutes noise in the signaling system.

Nevertheless the noise does not threaten the system as a whole, because a shaky equilibrium still exists: the big bank brands are still big, the lead aggregators are willing to “sell” their popularity (in the form of leads) rather than keeping it for themselves, and consumers have no reason to defect from a system that successfully pairs them with lenders. The game is sub-optimal for the big bank brands, because they’d prefer to acquire the leads directly at lower cost, so the costly signaling challenge is successfully managed rather than overcome.

## 7.9 Popularity-Based Signaling in Social Media

I've established that search engines like Google compel marketers to manage costly signaling on very different terms than they are accustomed to, i.e., they must learn to pay a costly signal with popularity as the coin of the realm. But search engines themselves are far less of a costly signaling challenge than the content results that they return, and the nature of that content has changed with the advent of social media. Brands now contend not merely with competitor content available online in a head-to-head setting, but now with the vast brand-related content of the social media realm.

Some perspective on this vastness: the British public relations firm Immediate Future released a study in 2008 identifying the top 100 brands in social media on the basis of raw number of brand mentions across major social media types. The top 5 brands all had 100 million or more mentions. This is vastly more content than a single brand can fully absorb, let alone control, but it hints at the extent to which the content of those tens of millions of conversations, good or bad, will shape the brand's reputation and its marketing success. For social popularity to gain a foothold as the emergent form of costly signaling, it had to be too big to ignore. Otherwise it's just noise in the system.

On a meta-level, the exponential growth of the social sphere has been essential to its emergence as the playing field for this new form of costly signaling, but on an individual basis, its impact is not a matter of volume but simply a matter of access. In other words, a single piece of content, sufficiently popular, can trump any other signal sent by a brand; the "United Breaks Guitars" incident is a prime example. But the anti-United video is popular because it's clever and taps into consumer desire to see brands taken down a notch, which made it go viral; consumers were not habitually doing Google searches for "United breakage incidents."

When it comes to costly signaling, I'm more interested in incidents where social content aligns with consumer attempts to gain information that brands are reluctant to give. Traditionally brands could rely on costly signaling to control the message; the financial success of a brand gave it access to a costly signal like television advertising, which in turn conveyed the brand's legitimacy.

As a result, large, established brands had nearly exclusive access to mass media, and negative information about the brand could typically not afford the costly signal such media demanded. But as we've seen, costly signaling only works if its participants believe in the materiality behind the signal; thus the growth of social media occurs in concert with the deterioration

of traditional advertising, as one system breaks down and the other gains strength.

## 7.10 The Disruptive Effects of Popularity-Based Signaling

To illustrate how access to social content can overturn traditional costly signaling, irrespective of volume, let's look at the hypothetical example of a national chain of high-quality assisted living facilities for seniors, which I'll call Cuesta Verde. Suppose that Cuesta Verde is the dominant brand in its space, with more than double the number of facilities of its next closest competitor. Clearly this is a brand that can afford the costly signal more readily than its competitors, and so it cements its dominance with national advertising in print, television, radio, and Web.

It would be reasonable to assume that a brand of this level of prominence, offering a service for which adult children, who are guiding this important and emotional decision, are heavily reliant on the recommendations of friends and family, would be much discussed in the social realm. And so one finds that adult children are indeed comparing notes on assisted living facilities on dozens upon dozens of forums and blogs, and Cuesta Verde is a frequent subject of conversation. Some of the content is positive, and some is not, but the brand thus far has seen fit to ignore these conversations and focus on its traditional costly signals in paid media.

Purely from a cost-signaling perspective (i.e., ignoring social media's usefulness for brand-building, cooperative marketing, or simply early detection of consumer defection), this calculation is reasonable, since the essence of signaling is that it must be worth the cost paid. So long as Cuesta Verde is acquiring new residents at an acceptable volume and cost through paid media and not experiencing the equivalent of an exploding laptop incident in social media, there is no imminent need to rock the boat.

But then the recession comes. Allow me to posit that the choice of an assisted living facility is highly cost-sensitive, so that even small economic fluctuations have a big impact on spend thresholds. Suddenly the content of *conversations* about assisted living changes, and by no coincidence, the content of related *searches* changes. The adult children and seniors exposed to Cuesta Verde's costly signaling still properly receive the signal that Cuesta Verde is a top brand in the space, but that is no longer their primary consideration. The top Cuesta Verde-related search is no longer "Cuesta Verde," which would naturally take the user right to the Cuesta Verde site, but rather "Cuesta Verde pricing."

And here lies the problem. Cuesta Verde is expensive. Because Cuesta Verde is expensive, they have historically chosen not to display their pricing on their Web site, but rather to address the delicate matter of cost in person, after the family has taken a tour and fallen in love with the place. This *had* been a wise strategy, not at all unlike the tried-and-true car-dealer method of getting the prospect out for a nice test drive before any discussion of cost takes place. This by itself is a form of costly signaling; it says to the prospect, “I have borne the cost in time and trouble to show you everything that this car/facility has to offer, because I am that convinced that this car/facility is right for you.”

But suddenly Cuesta Verde’s prospects are no longer learning about Cuesta Verde’s cost for the first time on the tour, after the costly signal has been received. Now they learn about it when they do a search for “Cuesta Verde pricing.” Because Cuesta Verde has chosen not to display pricing information on their Web site, per the strategy outlined above, their site is not the most relevant or popular result for this search. Despite their substantial investment in costly signaling, they have been thoroughly trumped by a single blogger who chose to detail her monthly Cuesta Verde costs in a blog post on the subject. The blogger’s cost is asymptotic to zero; the tools were free, and the post took 10 minutes to compose.

Cuesta Verde experiences a drop in enrollment. Many prospects are still driven to the site by paid advertising and complete the lead capture form, but a large portion of these prospects conduct price searches afterward and determine that the cost is too dear. They never receive the second costly signal on the tour. Some prospects do take the tour, but a portion of these do comparative pricing searches afterward and determine that other providers offer similar services at a lower cost.

Cuesta Verde’s response options are limited. Ignoring the pricing issue is a non-starter; it is clearly responsible for the drop in enrollment. The company could discount, but this presents several problems: it hurts the bottom line, and it potentially subverts the costly signal that has been paid to establish the company as a premium brand, worthy of the cost. And most importantly, a discount has to be promoted, at some expense; there is no guarantee that it would trump the alternate information being proffered by the blogger.

Cuesta Verde’s best option is to cooperate – to engage these new terms for costly signaling head-on. In the consumer’s mind, a company that stays silent to a prominent pricing issue is a company that can afford not to care. In other words, Cuesta Verde has inadvertently sent an adverse signal that excludes on-the-fence prospects who might be persuaded once they’re on the tour. That prospect is looking for some signal that, discounts aside, the

company recognizes that cost is an issue and that families are seeking value for their money.

A cooperative strategy for Cuesta Verde could include direct engagement with the blogger. Cuesta Verde could post a comment to the post that says, in effect, "If cost is a concern, at least come talk to us." As outlined above, a vague but cooperative response is preferable to silence, and preferable to a discount; per the randomization strategy outlined in Chapter 2, Cuesta Verde is better off working through pricing issues on a case-by-case basis. Cuesta Verde could take a similar stance on their own Web site; addressing pricing would give the site greater relevance and ranking in searches, potentially trumping the blogger. And as an extension of this strategy, for Cuesta Verde to signal its cooperative stance on other blogs and forums that raised this issue would be a significant step forward in accruing popularity capital in this new system.

Thought it may be obvious, it should be pointed out for the sake of avoiding lapsing into credulity that engaging the pricing issue in social media doesn't make it go away for Cuesta Verde; their services will still be materially too expensive for some prospects in a recession. It is simply a good opening move in a highly complex game. Despite the hype, social media marketing is almost never instantly transformative. For many companies, their first foray into social media occurs when the value of doing *something* exceeds the cost of doing *nothing* – hardly a prescription for changing the world. Starbucks' innovative decision to crowdsource its path to improvement didn't magically spare the company the need to close 600 stores, but it enhanced loyalty with its participating customers, produced good ideas that may win over even more customers, and it set the tone for long-term engagement. When was the last time advertising accomplished all of that?

The blogger with the sought-after pricing information is a challenge for Cuesta Verde because it subverts the traditional costly signal of paid advertising and a branded website, forcing the brand to reckon with a new system based on content popularity. But advertising is not the only system of costly signaling that's been turned on its ear in this new era; public relations changes too. The proliferation of bloggers challenges brands' ability to control information, but it also changes traditional publishers' ability to serve as the conduit for brand information.

I have already noted that the recent demise of many magazines can be attributed to the weakening of the costly signal of paid advertising in those publications, but editorial competition plays a role too. Simply put, publishers face a disruption in their own costly signal brought on by the proliferation of other publishing sources, especially blogs.



The blogosphere is so vast, and its rate of growth so rapid, that quantifying its dimensions is impossible; the blog tracking service Technorati was tracking approximately 112 million blogs when last reported, in 2008. While I have been wary throughout this study of ascribing too great an importance to sheer volume, it is beyond question that the number of blogs and, more importantly, the growing importance of individual blogs in specific areas of specialization, constitute a significant challenge to traditional newspaper and periodical publishing.

## **7.11 The Perils of Negativity in Popularity-Based Signaling**

This shift toward emerging media like blogs has an impact on marketing not only because of the diminished value of advertising in print publications, as described earlier, but also because it disrupts the traditional symbiosis of exclusivity between brands and publications. Brands are skittish about social media not merely because it's easy to get wrong but because negative publicity in social media can germinate and endure like a noxious weed in a way that was never possible in traditional media. Prior to the advent of social media, the average marketer could count on one hand the number of veritable "brand scandals" that sustained any lasting media coverage; now entire blogs are devoted to tracking the proliferation of such scandals in social media.

In traditional media, brands could count on costly signaling to provide a natural delimiter on negative coverage. For a print publication, the inherent cost of producing an investigative piece meant several things. First, it meant that a finite number of stories could be covered, and that those few that made it to publication had to pass through filters of veracity, legal compliance, and reader interest. Secondly, it meant that the publication's competitors were limited to those who could also pay the costly signal to conduct such investigations. A scrappy, self-published periodical might indeed scoop a major publication with an investigative piece, but this is not the same as paying the costly signal: without big sponsorships, the scrappy publication would lack the readership, and therefore the legitimacy, necessary to make hay out of the story. By the logic of costly signaling, it would cost the small and scrappy publication much more – for instance, they might have to give out the publication for free – to transmit their signal as successfully as the large publication could do.

Thirdly, and consequently, most brands are inured from negative coverage by these publications, except in the most egregious cases. For instance, Coca-Cola was covered in the 80's because their "New Coke" was

a spectacular failure, but in that same era, an instance of United Airlines breaking a passenger's guitar could not have merited coverage even in the local news.

It's important to point out that this traditional system neither required nor necessarily involved any actual collusion between publications and brands to hold back negative coverage; the general lack of such coverage was simply a natural consequence of costly signaling. But in the new era of the blogosphere, the traditional system is disrupted by the emergence of the potential for negative stories to achieve wide circulation largely irrespective of ad revenue and investigative budget. If you'll allow me to treat presidential candidates as brands, I will illustrate this disruption by way of its impact on the 2008 U.S. presidential campaign.

## 7.12 Popularity-based Signaling in Presidential Politics

Indisputably Barack Obama emerged as an early master of the new costly signaling system of Web-based popularity. Since entire books will be devoted to this subject, I will not endeavor to make mine one of them. But any brand that takes the time to study Obama's use of social media for clues on brand engagement would probably find the time well spent, and so I will offer some initial analysis here.

Obama's online fundraising has garnered the most attention; according to the *Washington Post*, the campaign added half a billion dollars in online donations to its record-shattering total during its 21-month run. But the fundraising total is simply the index of the campaign's broader success in social networking. The campaign constituted its own Web site as a social network, allowing each visitor to create their own profile; more than 2 million were created. More than 5 million supporters connected with the campaign through other social networks like Facebook. Those networks produced over 400,000 individual blog posts – that's in addition to the extensive blog coverage given to Obama outside of his network of supporters (Vargas 2008).

Lest this be mistaken for history's most successful grass-roots-only campaign, it should also be pointed out that the Obama campaign also spent more on mass media than any other campaign in history. But that expenditure would not have been possible without the social network effect. Social media didn't *replace* the costly signal of traditional media; it merely proved to be an equally viable force.

Consistent with my theme throughout this study, I am less interested in the raw quantitative aspects of Obama's savvy use of social media than with the qualitative aspects of his natural mastery of its peculiar demands. The

biggest impact of the sudden eruption of 112 million+ blog posts on a presidential campaign is not the volume of raw coverage; after all, the same story repeated 112 million times is still the same story. Rather, the impact is felt mostly keenly in the exhaustive parsing, analysis, and meta-analysis of a candidate's every word, nuance, and verbal and facial tic. In this respect, Obama was a candidate remarkably well-suited to the YouTube era, in which no moment of the campaign caught on camera would be free from scrutiny.

Obama's campaign team displayed an astonishing level of message discipline, and Obama himself rarely strayed from carefully chosen talking points. But as a brand engaged in collaborative marketing, Obama showed the most aplomb in his ability to appear unscripted and relatable while burnishing the brand. Widely criticized for his lack of foreign policy experience, Obama made a trip to Kuwait to visit U.S. troops at the height of the campaign. During a speech to troops in a basketball gymnasium, someone in the crowd tossed a basketball to Obama. With cell phone cameras recording his every movement, Obama turned to the basket and drained a perfect three point shot – nothing but net – on his first throw. The troops went wild. The amateur videos of the incident have been viewed on YouTube more than 1 million times. History will not remember what Obama said to the troops, but the three-point shot belongs to the ages.

Possessing the ineffable qualities of a YouTube star, Obama had major advantages over his primary opponent Hillary Clinton and his general election opponent John McCain, both of whom honed their political skills in an era of traditional media. McCain was famous from his quixotic 2000 campaign for his "Straight Talk Express," in which he allowed unfettered access to journalists at a table in the back of his campaign bus. But the journalists, immortalized in David Foster Wallace's account of the campaign as the "12 Monkeys," were all from major media outlets; McCain was still playing the costly signaling game (Wallace 2000). He was derided in the 2008 campaign for his lack of Web savviness, especially after he avowed in an interview that he was finally learning how to use the Internet, thanks to his wife.

In the Clinton camp, the shift in the signaling game was most pronounced in the performance of former President Bill Clinton, whose prowess as a campaigner was legendary. Considered a natural asset to any candidate he stumped for, Clinton nevertheless stumbled in his efforts on behalf of his wife, precisely because the 24/7 scrutiny of the new media era outpaced his traditional methods. In Clinton's own presidential campaigns, his long-rumored temper was never on display and was largely unknown to most voters; in the 2008 campaign, every intemperate moment was captured on video and viewed online hundreds of thousands of times – a search on

YouTube for “Clinton tantrum” yields a wide selection of choice moments. A profile piece in the *New Yorker* quoted a Clinton campaign official lamenting that Clinton appeared to have been plucked from a previous era and dropped into one in which he could not adjust to the constant scrutiny (Lizza 2009). In costly signaling terms, Clinton's problem was obvious: he expected journalists to pay a costly signal for access to his remarks, which would then allow him to control coverage much more readily. In the absence of that costly signal, any and all of his remarks were fair game.

And so we find ourselves, as Bill Clinton did, in an anxiety-provoking interregnum period. The traditional methods of costly signaling, which gave marketers generally predictable exposure in major media outlets at a reasonable return on cost, have been upset by the emergence of a new system of costly signaling in which an amorphous standard of popularity allows competing brands and competing points of view to enter and sometimes dominate the conversation. As with my previous examples of marketers reacting to changes in the game structure, the changes in costly signaling invariably begin with a certain amount of bad behavior, as marketers struggle to get it right. A rational view of this bad behavior would hold that marketers *will* get it right, i.e., they'll learn to cooperate rather than defect, as a simple matter of self interest: the benefits of cooperation outweigh the benefits of defection.

### 7.13 Noka and the Disruption of Costly Signaling

In the meantime, though, there is nothing more instructive than brands getting it wrong, since their actions allow us to detect the emerging rules of cooperation in a highly iterative game. My first example involves not a large, established brand but rather a small, up-and-coming brand that has been particularly reliant on a traditional form of costly signaling. The brand in question is an ultra-premium chocolatier, Noka, that found itself at the center of not one but two controversies involving the role of blogs in establishing and maintaining brand reputation. The brand's travails are a dramatic example of the chaos theory concept of the “butterfly effect” – the notion that a butterfly flapping its wings in the Amazon basin could ultimately cause a hurricane half a world away. But as I've shown, the large disruptions caused by small social media eruptions are not chaotic but inevitable. Such was the case here.

Noka is a Texas-based chocolatier founded by a husband-and-wife team in 2004. According to its own Web site, the company specializes in “single estate chocolate” which means that the chocolate originates with beans

grown in a single place, producing a distinctive quality that comes at a premium. A *Forbes* magazine feature lists the chocolate as one of the world's most expensive, at an average of \$854 a pound; most other chocolates in the same exclusive category on Forbes' list sell for less than \$100 a pound.

How does a brand that sells for considerably more than even many of its ultra-premium competitors convey its value? Costly signaling plays an important role. The brand conveys its exclusivity by limiting its distribution; it is sold directly through the company, through only two retail locations, and through high-end retailers like Niemen Marcus. Its distribution alone is a costly signal: a consumer shops at Niemen Marcus in part because the consumer can afford to pay more for things, and in exchange the consumer expects that anything purchased at Niemen Marcus will be of premium quality.

The company also relies on coverage in premium print publications as part of the signaling game. Significantly, the company's press page does not contain the typical chronological listing of press releases; it consists only of reprinted articles and news items from the company's appearance in publications like *Entrée*, *Exquisite*, and *Level Maldiva*. When viewed in isolation, these articles constitute an effective costly signal in much the same vein as the chocolate's availability at Niemen Marcus; the consumer who can afford the items advertised in this magazine can also afford to pay more for chocolate. The clever gambit here is to use *screening* as a form of signaling; the prospect wants the item more, and is willing to pay more for it, because it is exclusive. This is reminiscent of the old chestnut, "If you have to *ask* how much it costs, you probably can't afford it."

In this way, *cost itself* acts as a costly signal, in a neat bit of recursive logic: the product is expensive because it is premium. How do we know it is premium? Because it is expensive. This seemingly bizarre tautology is actually quite effective in creating costly signals for luxury items, especially where highly subjective matters of taste are concerned: we expect good wine to cost us more, and if a very good wine were to appear in the discount bin in the grocery store, we would probably be suspicious and refuse to buy it. My father, who restores and sells antique furniture for a living, explains the logic in this way: "Sometimes people just *need* to pay more." Indeed, consumers' enjoyment of a luxury item may actually be *enhanced* by the act of paying more.

Finally and most distinctively, Noka issued a costly signal justifying the cost of its chocolate by opting *not* to send a costly signal. Bear with me a moment as I unwrap this paradox. What distinguishes Noka the most from other premium chocolatiers, besides the exponential price, is the understatement of the product itself, its design, and its packaging. Premium

chocolatiers generally try to enhance the multi-sensory experience of eating premium chocolate with elaborate creations that appeal to both the eye and the palette. Noka's signature chocolates are small, stark, unadorned rectangles of pure chocolate and simple truffles. The packaging is a plain box stamped with the Noka logo in the center. The simplicity and understatement are instantly compelling.

And they are compelling by design. Some fascinating follow-up work to Spence's theory of costly signaling has been published in the *RAND Journal of Economics*, which posits that signaling through counter-signaling, i.e., by consciously refusing an ostentatious display of status, may be very effective. The study noted the tendency of mediocre students to eagerly answer questions posed by the teacher in the hopes of signaling a higher status, while the very best students tended not to answer, because such obvious displays were beneath them. In this model, the top signaler avoids contributing to all of the noise created by competitor signals, demonstrating that they are above the need for such things. Since signaling is a way of conveying information, refusing to signal can create a vacuum that the consumer is compelled to fill with their own assumptions. In the case of premium goods, the consumer may fill that vacuum with an assumption that the missing information is highly favorable. In other words, only a truly great chocolate wouldn't bother to show off its greatness.

What I've just described is an example of effective cost signaling in its traditional terms, in which exclusivity is a powerful part of the signal, even extending so far as to include under-signaling. Such a brand might actively eschew a popularity-based signaling model, since popularity might actually undermine exclusivity. Popularity-centered social media tactics like participation in Facebook might be ill-advised, since the goal of signing up lots of fans run counter to the goal of appealing to a select few.

The problem is that brands cannot simply opt out of the shift toward popularity-based signaling; marketers have choices as to whether or not to actively participate, but they do not have a choice about popularity-based signaling's impact on their marketing environment, as we saw in the case of Cuesta Verde. So as you might have anticipated, the next chapter in Noka's story is the emergence of a detractor in the popularity-based system.

In December 2006, an amateur blogger published a 10-part expose on Noka on his foodie blog, DallasFood.org. As a piece of amateur investigative journalism, the series is remarkable for its thoroughness and rigor. It began with the simple premise of whether Noka chocolate was worth the price, then went on to establish a set of claims that challenge the chocolate's price. The blogger offered detailed price comparisons among other chocolatiers and showed how others used the same practices that ostensibly set

Noka apart. He showed that Noka sources its chocolate from a French supplier that also supplies other premium chocolatiers, and that Noka's mark-up was 1300% of the supplier's retail price for their own line of chocolates.

Within a month, the *Dallas Morning News* reported, the DallasFood story had been picked up by 10 blogs, but that number alone doesn't tell the tale (Robinson-Jacobs 2007). The story appeared in highly prominent, popular food and consumer blogs, including *ChowHound* and *Consumerist*, which boasts a monthly readership of 1.8 million. The story was plucked from obscurity almost before it had time to be obscure, and Noka's system of costly signaling was effectively compromised.

Because of the exclusivity of Noka's sales channels and its limited press coverage – both of which are essential to its costly signaling strategy – the chocolatier had little positive coverage to fall back on in a popularity-based signaling system, i.e., in a search on Google. As of this writing, a search for “Noka chocolate” on Google returns the company's own Web site as the first result, but the results that immediately follow on the first page consist almost entirely of negative coverage through DallasFood.org, the blogs that picked up the story, or user reviews. There is one exception: a video posted to YouTube titled, “Noka Chocolate – How Luxury Chocolate is Made.” It's a positive local news story done prior to the DallasFood expose, and it is posted by “KeeneyPR,” a Texas-based PR specialist named Dan Keeney. And thus we begin the third act of this saga.

As soon as the Noka chocolate story began getting picked up on various foodie blogs, a very persistent commenter known only as “Dan” began posting passionate defenses of Noka in the blogs' comment sections. In a comment on the blog Crypticide, he noted that the Dallas Food blogger's “previous claim to fame appears to be a multi-part series on chicken fried steak” (Crypticide 2006). Bloggers and their readers were instantly suspicious; in the comment thread on the foodie blog “Kitchen Mage,” for instance, the blogger immediately responded to the comment by asking whether “Dan” worked for Noka. Two days later, Dan revealed himself as PR man Dan Keeney, and claimed that while he had not been employed by Noka at the time of his original posting, but was merely a concerned chocolate lover, he was, in fact, now retained by Noka as their PR representative, a mere two days later (Kitchen Mage 2006). The blogosphere howled in outrage.

To bloggers and readers, Keeney's sin was not his participation in the debate, but his failure to disclose the nature of his interest in the story; his claims to have been unattached to the company 48 hours earlier were not viewed as credible, to say the least. To Keeney's credit, he took up the topic on his own Web site in a post titled “Ethical Considerations In Posting Comments to Blogs” and endured considerable tongue-lashing from

commenters, though he remained steadfast in his claim that he had not deliberately obscured his identity (Keeney 2006).

### **7.14 Sock Puppetry as Noise in the System**

This practice has earned the name “sock puppetry,” referring to any attempt to obscure one’s online identity whenever that identity is material to one’s vested interest in the subject. Scorn seems to be heaped more gleefully on revealed cases of sock puppetry than almost any other social media transgression, and it is worth asking why this is. I believe the level of outrage is a matter of perceived defection. Previous chapters have traced the evolution of social media as one that produces a very fragile cooperation between marketers and consumers after many, many rounds of mutual defection. Social media has emerged as a playing field for consumer empowerment, but one in which brands that play by the rules can not only succeed but achieve the coveted 4-4 equilibrium.

When a brand or its representative masquerades as a consumer rather than a marketer, it utterly shatters the delicate equilibrium; it is perceived as the worst kind of defection. Why? Because identity in social media is not filtered by traditional costly signaling; a blogger or a commenter on blogs does not typically offer a pedigree, but establishes their reputation through the popular acclaim given to or withheld from their remarks. In other words, they must pay a popularity-based costly signal.

When a blogger or commenter has a hidden vested interest, especially in the way of a traditional paid relationship, they introduce noise into the new system of costly signaling. Should all commenters or bloggers who defend a brand be automatically treated as some sort of fifth columnist? Clearly not. But that is just as clearly the danger that is posed by such acts of defection; if consumers come to believe that any cheerleading is automatically suspicious, then marketer/consumer engagement in social media will fail.

While the system can’t protect itself completely from bad actors, the emergence of such implicit rules as “Thou shalt not sock-puppet” is a step in the right direction. Since disclosures of conflicts of interest are entirely common and a matter of basic business ethics in other arenas, such as journalism and law, it should be a surprise to no one that they should apply in this new arena as well. If all else fails, the simple rationalism of the iterative game must prevail: when the cost of defecting outweighs the cost of cooperating, it is always better to cooperate.

To use my agency’s own social media participation as an example, I can acknowledge that disclosure comes with a cost. We frequently participate



in forums and blogs on behalf of our clients, as Dan Keeney (perhaps) did, though it is generally to make an announcement that might interest the forum's participants rather than to engage in brand defense. These posts are always accompanied by a "full disclosure" statement that acknowledges the paid relationship. In nearly all cases, such posts are welcomed, provided they adhere to the forum's bylines for commercial posting. Occasionally, they provoke ire from participants who would prefer not to be exposed to commercial postings of any kind, but that is simply the cost of participation. We might avoid these costs by not disclosing the relationship, but then we would be bad actors, not only threatening our client's position when the sock-puppeting is exposed, but undermining the system as a whole. Our long-term success depends on cooperation, and so we pay its cost.

One final aspect of the Noka saga is worth highlighting: in an unfortunate effort to disparage the Dallas Food blogger as lacking credibility in his critique of Noka, Dan Keeney noted, as described above, that the site's sole claim to fame was an article on chicken-fried steak. This may be regarded as the last defense of the traditional costly signal. Keeney's remark assumes that the seeming triviality of the blog's past topics should diminish its credibility on the subject of Noka. This signaling cost may have indeed been necessary in a traditional model, but it is largely irrelevant here. The Dallas Food blog presented facts that were independently verifiable (in many cases with links), and it covered a story that interested a great many people. In the new system of popularity-based signaling, nothing more is required.

## **7.15 Target Learns the Rules of Popularity-Based Signaling**

My second case study is an extension of the Dan Keeney aspect of the Noka saga, i.e., what happens when a company embraces the new system of costly signaling but then introduces noise into the system. The mega retailer Target has received acclaim in recent years for achieving a turnaround in its brand identity with a hip, youthful appeal that achieves the remarkable feat of making shoppers at the discount retailer feel hip and smart for shopping there. In concert with its considerable inroads with younger demographics, Target began a program in 2007 called "Target Rounders," which invited college-age students to promote Target products on Facebook in exchange for discounts and prizes.

At the outset, such an arrangement presents difficulties for Target. Providing a direct quid-pro-quo for endorsements runs the risk of creating perverse incentives, i.e., Target Rounders may endorse a Target product simply in order to earn points and not because they actually like the

product. And what is the matter with that? The matter is that such claims are less credible than those based on genuine preference, and the two types become indistinguishable in a scenario like the Target Rounders program. A consumer-reviewer does in fact pay for a costly signal: their first-hand experience of having tried the product and liked or disliked it, coupled with their ability to articulate the reasons why, is their mark of credibility in the new system of costly signaling.

The result of this diminished credibility from not having paid for the costly signal is that the endorsement is *a priori* discounted by consumers on the basis of the paid relationship. How do we know that the Target Rounder *really* liked the lamp? We don't, and so we ignore the endorsement. This action, in turn, diminishes the value of the costly signal; Target no longer has extra credibility for its products based on their popularity with the Rounders. The resultant signal may, in fact, be less valuable than a paid advertising signal, because wrangling a group of college students into paid endorsement scenarios takes considerably more effort than running ads, with no greater hope of credibility.

Faced with the prospect of a diminished signal due to perverse incentives, Target had two options. The first option was to stop paying for endorsements from the Target Rounders. It is my sad duty to report that Target did not choose the first option. Clearly the cooperation game still has a long way to go. The second option is to cover up the nature of the arrangement with the Rounders, and that is, unfortunately, the route that Target chose to take.

In fairness, there is some confusion as to whether the effort to obscure the quid-pro-quo relationship came from Target or from its agency, acting on its own, as Target later claimed. But what is known is this: according to the *Minneapolis-St. Paul Star Tribune*, Target Rounders received an email newsletter from the program that stated, "Your Mission: Try not to let on in the Facebook group that you are a Rounder." The newsletter went on to rationalize the request as an effort to keep the Rounders program from "stealing the show" from the real Facebook star, Target (Crosby 2007).

In yet another incident of a flapping butterfly creating a hurricane, one of the email's recipients was a University of Georgia student whose journalism professor maintained an active blog on the role of PR and social media. The student was alert to the ethical considerations involved, and posted a Facebook message decrying the new Target Rounders policy. According to the student, her posts were then deleted. Soon after, the student received an apology both from Target and the agency responsible for the program, and the policy was abandoned. A Target spokesperson later declared, "Target is not interested in feeding guest feedback or public opinion. Negative feedback is as valuable as positive." This would appear to be a deviation

from the company's original intent for the program; the program's "Code of Conduct" states that "the message board is a forum for the members of Target Rounders to communicate and share ideas about Target and Target products, *in a positive manner.*" (emphasis mine)

Setting aside for the moment the spokesperson's vested interest in spinning the story, her claim is an important one. Paid-for positive input not only fails to meet the costly signaling standard, it pollutes the data that is arguably the more valuable by-product of social media participation: real-world insights on consumers' experience of the brand. Creation of an artificial environment of positivity is sub-optimal for the both the brand and the consumer, because neither can trust the endorsements being provided.

Target has since abandoned the Rounders program, and it is unknown whether its diminished costly signal or the controversy over disclosure paid a role. Since Target suffered the slings and arrows of being an early adopter in the new system of popularity-based costly signaling, it is ironic that the company subsequently received another round of bad publicity in social media circles for *not* recognizing that they needed to participate in this new system.

In January 2008, the *New York Times* reported on a blogger whose blog, ShapingYouth.org, concerned itself with the way that marketing shapes children's self-perception. The blogger, Amy Jussel had taken umbrage with a Target ad that showed a woman lying across a target with her crotch at the bullseye. She complained to Target, and received an email response that declared, "Unfortunately we are unable to respond to your inquiry because Target does not participate with nontraditional media outlets... This practice is in place to allow us to focus on publications that reach our core guest" (Barbaro 2008).

The fallacy of this position is obvious enough that I don't need to plumb its depths; there are any number of "nontraditional media outlets" that would reach Target's audience as well as or better than traditional media outlets. The claim also assumes that the sole purpose of answering PR inquiries is to get free access to one's audience; no PR professional would accept such a claim. And perhaps most significantly, the claim belies the fact that Target participates very heavily in nontraditional media outlets, with the Target Rounders program being an obvious example. It would be far easier to excuse a very traditional company for which social media remains a vast, uncharted territory; but Target clearly knows its way around the place.

Rather, I raise this incident because it illustrates two points. The first is that companies like Target that have developed some reputation for being customer-centric, particularly with a focus on the younger demographics that comprise the core of the most active social media participants, have

greater opportunities in social media, and therefore greater responsibility. Recalling Spence's axiom that a highly qualified player will be able to acquire the costly signal for less cost than an unqualified player, we can posit that highly customer-centric brands will have greater initial success in social media than non-customer-centric brands – all other factors being equal. For example, if both Target and Costco put up Facebook pages at the same time, and did nothing else to attract fans to the pages, we can rather easily assume that Target's fan base would grow organically at a much faster rate than Costco's.

Consequently, Target will find itself held to a high standard of accountability, because its customers *care* whether the brand cooperates rather than defects in the social space, while defection by other brands may go entirely unnoticed. One of the cardinal (and common) sins of the early social media era is that many marketers seem to believe they can turn participation on and off, as though it were a paid media campaign, and can cherry-pick the venues that cast the brand in the most favorable light. Target may have believed that it could and should take full advantage of Facebook, where millions of its young shopper congregate, but that it could ignore blogs in its PR strategy, because one tactic is better at getting to the “core guest” than the other.

Customer-centric brands like Target *can* choose not to participate in some respects (one can easily avoid the burdens of Twitter, for instance, by not signing up for a Twitter account), but never entirely: one of the challenges of a popularity-based system, as opposed to one in which the campaign can be turned on and off, is that the conversation doesn't necessarily start with the brand; it can just as easily start with the customer. Coca-Cola discovered this and used it to their advantage; AMC discovered this in regard to *Mad Men* tweeting, and made a serious misstep. The trend itself is inevitable; a natural consequence of new opportunities for cooperative games is that consumers will sometimes use this new empowerment to make the opening move. As previous analysis has shown, the correct counter-move for the marketer is *never* defection.

## 7.16 Target and the Problem of Capacity

My second point is a bit more sympathetic to Target's dilemma in setting rules for engaging with “nontraditional media outlets” like bloggers. Clearly Target's non-cooperative stance is the wrong one, and they paid a price for it. But the shift to a popularity-based costly signaling system raises some non-trivial issues of capacity. In a traditional costly signaling system, capacity is not an issue; before the advent of cable television, if an advertiser wished to

cover all of network television in the U.S., it required three media buys and very little follow-through. Assuming the vast and uncharted blogosphere is somewhere between 100 and 200 million individual blogs, it would be reasonable to assume that a broadly appealing, ubiquitous brand like Target would receive a fair number of blogger inquiries.

So the capacity question is this: how exactly should Target's PR department decide when to engage with bloggers? The simplest rule is also the most cooperative: respond to all bloggers. Doing so would acknowledge that under this new costly signaling system, all conversations are important. There is also the practical matter that it would be impossible to make a fair judgment as to which bloggers to respond to and which to ignore, and so the safe course of action is to respond to them all.

Safe, but practical? Many large companies have excellent customer service channels, but addressing issues like shipping errors is qualitatively a much different task than responding to blogger outrage. The latter cannot be resolved by waiving fees; it demands specialized skills. The blogger raising the question of the advertisement with the unfortunate crotch placement could not be set aside easily; the respondent would be compelled to address the tricky question of whether the advertising should be pulled, and why or why not. Should the brand be hijacked by any blogger who raises hell about ad content they believe to be salacious? Clearly not.

The middle path, then, would seem to involve judgment calls about which bloggers to respond to, but as noted above, this is problematic. One could conclude, for instance, that size is important; the brand's marketers could set a threshold for dealing with blogs based on their readership, in the same way that marketers might choose print publications for paid advertising. Larger blogs would get access; smaller ones would not. But the popularity-based system eradicates this kind of hierarchy, because a blog's actual readership is far less important than how often it gets picked up or linked to by other blogs with larger readerships. Take the two examples already under consideration: DallasFoods.org was a small, locally focused blog, but its story got picked up by Consumerist, a massively popular blog with a monthly readership of 1.8 million. The Target blogger story got picked up by the *New York Times* blog; it piggybacked on the costly signal of the newspaper of record.

Setting aside the question of size, then, one could conclude that the current mania for measuring *influence* is the way to go. We can actually screen blogs based on how often their content gets picked up by other blogs, and then by yet more blogs, and so on. This would seem to get us closer to a standard based on popularity.

But influence is a complicated standard. Measuring the influence of a blogger or other social media participant on a purely quantitative basis, i.e.,

how often a story gets picked up, can easily lead to false positives. An individual may be highly influential within a small, reciprocal circle, and influence could wax or wane depending on the topic. DallasFood.org apparently did not make waves with its chicken-fried steak series, but its Noka series made it a short-term celebrity.

All of this may add up to a massive PR headache, but once again, it's only a headache by traditional standards. Despite all the uncertainties about the rules of engagement for the new system of costly signaling, one change is already clear: Marketing organizations will need a serious reappraisal of roles in order to keep pace with the change in systems. Social engagement doesn't require a media spend, but it may require a proportional spend in personnel to manage its vagaries and pitfalls. Social engagement and crisis management roles would exist side-by-side with traditional media roles, and reduced spending in one area could fund the other. This shift may be hard to swallow for organizations that persist in thinking of social media as a free ride that simply replaces a portion of paid media; while its cost may not be borne out in paid media, it remains a costly signal.

Target's capacity problem in its media relations points to an overarching problem in popularity-based costly signaling: if brands all now have to compete on the basis of their ability to command attention, how is such a system even remotely sustainable? We are constantly reminded that consumers' lives are ever-busier and unable to keep pace with the rapid growth of demands on their attention. If a brand's success in social media demands 500,000 hits on its viral video in YouTube, what happens to both the competitive space and to social media itself when every consumer brand aims for that standard?

Once again, I must defer the detailed examination of this issue to the final chapter; I mention it here because doing so allows me to address how the new system of costly signaling is evolving to meet this demand. The great challenge in social networking is no longer how to connect people and content – there are seemingly endless variations on ways to do that – but rather how to make sense out of all of the new content. Enter the content aggregator.

## **7.17 Content Aggregators and the Evolution of Costly Signaling**

Content aggregator tools allow users to identify and organize Web content that interests them. In this respect, they are no different than the indexing tools – for instance, the “Favorites” function on a Web browser – that have

provided added convenience to the Web browsing experience since the Web began. What interests me about the rise of content aggregators like Digg and StumbleUpon is the extra dimension they add to popularity-based costly signaling. These tools allow users to identify content that interests them based on a combination of keywords and popularity; measured in the case of Digg, for instance, on the basis of votes or “Diggs” that a given piece of content receives.

What is remarkable about the content aggregators as an evolution of the popularity-based costly signaling provided by search engines is that they move beyond the implied voting used in the Google algorithm, which is really just a way of treating linking as voting, to a literal vote-based signaling system. Users can search for content on a keyword basis, as they would with a search engine, and then choose content based on its popularity, or they can simply browse content based on popularity. Much as Google became its users’ de facto experience of the Web, now content aggregators can augment or replace that experience with one of their own, with a far greater degree of user control. Since content aggregators now collectively boast hundreds of millions of regular users, it is clear that this shift is well underway.

As Web content proliferates, this second evolution in popularity-based costly signaling is no mere convenience; it is in effect the only reasonable means by which users can make sense of the Web, short of confining themselves to the tiny sliver that search engine results can provide. It is axiomatic that these models will proliferate alongside Web content itself; as the universe of available content becomes wider, the individual user must necessarily improve the lens they use in their telescope. Consistent with users’ increasing reliance on connections among people they trust to help them apprehend this content, content aggregator users will increasingly focus not simply on raw votes in order to select content, but on the specific opinions of those whom they trust.

Signaling models will increasingly account for this need. In selecting books on Amazon, one can focus only on trusted reviewers, some of whom have been voted into special reviewer status by other Amazon shoppers. On Digg, you can focus your preferences on content that’s been voted on by others whom you trust, and even integrate your Facebook circle so that you confine your attention to material that friends recommend.

## **7.18 The Three Rules of Popularity-Based Costly Signaling**

It hardly needs to be pointed out that this more refined signaling system presents even greater challenges to marketers than the initial shift to a

popularity-based model. If users increasingly confine their experience of the Web to peer-endorsed content, and if users increasingly use the Web to the neglect of other media outlets, where, exactly, is the marketer supposed to break into the conversation? There are no easy answers, but I can offer a complex answer, in three parts:

The first part is to point out that the difficulty of this model is precisely the point of costly signaling. In Spence's original concept, the signal is meant to separate better candidates from lesser candidates, because the costly signal would be too expensive for the lesser candidate. In traditional marketing, this *separating equilibrium* simply distinguished brands that could afford marquee advertising from those who could not. In the popularity-based model, customer-centric brands will pay less for the signal than other brands, and such a system works very well for consumer: it means they'll be able to focus their interactions with brands in social media on the ones that that were more interactive and engaged in the first place. Brands that don't meet the criteria for being customer-centric will need to earn that reputation in order to participate successfully in that model; the Dell example from an earlier chapter is proof that brands can do so. In any case, the costly signal is meant to provide a separating equilibrium, so the fact that some brands will get left out in the cold is precisely the point.

The second part of my answer is that popularity-based costly signaling requires a change in mindset in order to obliterate the increasingly outdated notion that the goal for marketers is to "break into" consumer conversations taking place in social media; rather, the goal must be to start conversations and see them through to their conclusion. Marketers remain mired in the traditional concept of marketing as something that interrupts what consumers are doing in order to try to get them to do something else, in much the same way that salmon fishing involves distracting the salmon as it makes its way upstream to spawn. Unless marketers wish to experience the same sense of futility known to salmon fishermen, they must reform these practices.

What would that entail? To return to my persistent theme, it first entails moving away from primarily quantitative means of measuring success. If the marketer is focused on a quantitative goal for the number of "diggs" on Digg, the temptation to defect – by paying for diggs, for instance – becomes overwhelming, and such perverse incentives will ruin the system for both marketers and consumers. Target's Facebook program imbroglio is proof of that.

Instead, the marketer would need to focus on cultivating and then empowering their brand evangelists – something I'll discuss further in the next chapter so that their most passionate advocates are inspired to promote them within the popularity-based system in a way that is authentic and



sincere. The data presented earlier on consumer trust makes it quite clear that consumers will readily accept such advice, and the result of this peer-to-peer brand transfer is likely to be far more sustainable, while being far, far less measurable, than anything that traditional media can provide.

But the third part of my answer is that there is *still* traditional media. Ad-supported content models are in no danger of going away; they are merely waning in influence. As I've described, they will wane to a point of equilibrium, where the cost of participation is deemed proportional to the return, but they will endure in this new form. In meantime, popularity-based models, many of which also incorporate advertising, will continue to grow, insofar as they remain profitable for their sponsors. The bottom line is that marketers have choices. Many marketers may no longer be able to sit on the sidelines of the popularity-based system, given that conversations are already taking place about their brand, but they can choose their level of investment based on the proportional return offered by both systems.

The bottom line is that in order for cooperation to occur between consumers and marketers in this new system of costly signaling, marketers must avoid the temptation to defect, and correspondingly, the cost of defection must outweigh its rewards. A search engine-based model, unfortunately, offers both opportunities and incentives for defection. Anyone who can game the system of inbound linking, thereby simulating popularity, can send a costly signal through a search engine without actually paying the cost. In fact, subverting Google's ranking system has become a popular game with pranksters, because it unveils the mighty search engine's core weaknesses in its method of assigning popularity. In one notorious example, pranksters ensured that the top result for the phrase "miserable failure" was the official White House biography of then-President George W. Bush. The method was simple: convince enough people to create inbound links to the page containing the phrase, and the search engine would infer both relevance and popularity, and return the result. The phrase may have been a popular sentiment about the president, but it was not a phrase that the official White House biography would use to describe him.

Social media venues for popularity-based costly signaling make defection more difficult, because the high degree of user control renders them largely self-correcting. A brand on Twitter that hasn't paid the costly signal of high engagement with consumers simply won't get followers. A Facebook page for a socially disengaged brand will get fewer fans than one for a highly engaged brand. An unpopular viral video on YouTube simply won't go viral. And Digg content that isn't popular won't get "dugg."

In traditional paid media, a marketer can at least count on the paid-for level of exposure; the arena for cooperation is limited to whether or not the

consumer liked the ad and would respond to it. In these new costly signaling systems, paying the high social signaling cost is necessary even to get the *exposure*, i.e., to get consumers to participate and/or pass along the content. Getting consumers to respond is another matter entirely. Both require a high degree of cooperation that build reciprocal behavior over time. This daunting two-step filter is what has marketers pulling their hair out about social media, but it is also what will preserve the integrity of social media as a venue *only for marketers who pay the costly signal*. When the dust settles on these new forms of costly signaling, a separating equilibrium will convince some brands to stop trying to launch viral videos until they've done the hard work of paying the costly signal.

### 7.19 Paying the Costly Signal to Go Viral

I will close out my discussion of costly signaling by taking up this question of “going viral” because I believe it represents the largest gap between what most marketers would like popularity-based costly signaling to be – a form of free media – and what it actually is: a separating equilibrium that relatively few marketers will master, if mastery is measured in raw numbers. I am referring to the current mania for creating marketing videos for the expressed purpose of making them “go viral,” which generally means to generate high impressions at no cost.

One can uncover numerous marketer perspectives on what makes a video go viral, and I have been asked the question many times myself – occasionally by clients that would like to jump on the bandwagon. I struggle with how to answer, but the concept of costly signaling at least gives me a place to start: the popularity signal is costly because it requires mastery, but not mastery of a set of Pavlovian-style tactics that make videos magically popular. Rather, it requires mastery of cooperative signals: the brand that succeeds at viral marketing has mastered the relationship between how its product or brand is presented and what the audience would like to see. It is “authentic,” in the sense that it was developed to achieve the goal of meeting that desire, rather than the more mercenary desire to go viral. This is, I'm afraid, the same tedious but correct solution that states that marketers must do their homework on cooperation rather than defecting in the hopes of a free ride. As always, allow me to make this point clearer by way of example.

When we analyze successful viral marketing videos, we find that content and style don't at all fall neatly into a set of guidelines. In fact, a complete meta-analysis of viral video content would inevitably reach haphazard conclusions, because the content isn't the thing. The relationship of content to

audience is the thing. Nike, for instance, tends to produce viral videos that become popular for roughly the same reason that their paid advertising is viewed with some affection; their ads are heavily stylized, with high production value, and they often feature popular athletes. Does it follow that these are ingredients for a successful viral video? It does not. It merely follows that the Nike videos contain the ingredients that Nike's audience is looking for, and Nike has a big audience. Their videos are a form of cooperation, and the audience reciprocates the cooperation.

In order to develop a point of contrast, we can return to the failure of Motrin's infamous "sling mom" viral video. To the same degree that Nike succeeds by giving their audience more of what they want, Motrin failed by giving their audience something decidedly out of sync with what they want. Motrin moms want to be taken seriously; the video poked fun at them. It is an altogether too common example of a viral video that derives a set of false positives about success from other viral videos, where the lesson seems to be, "Be edgy." But this prescription has never been accurate; the edge moves with each audience, and in Motrin's case, they stepped over it.

## **7.20 A Popularity-Based Success Story: Blendtec**

One of the most successful viral video campaigns – very popular among social media marketers for its Cinderella quality, which I am equally unable to resist – is the one launched by by a small commercial blender manufacturer named Blendtec. The company's blenders were well-known within a small commercial sector of the appliance market for their astonishing toughness. As part of its quality assurance process, the company routinely blended household objects like hockey pucks to validate the blenders' power; the company's marketing director had the idea of creating a kind of mock science experiment segment called "Will it Blend," in which the company's CEO, Tom Dickson, would blend household objects of increasing improbability: action figures, CDs, golf balls – even an iPhone. As the story goes, the company invested only about \$100 in its first video, which garnered 23,000 hits on its first day on YouTube. More videos followed, all repeating the same formula with new blended objects (Briggs 2009). As of this writing, Blendtec's YouTube channel has been viewed more than 4.1 million times. The company reports that retail sales are up 700 percent, and the story has been featured in many major media outlets, not to mention countless marketing blogs.

The various analyses of Blendtec's success do tend to emphasize the cooperative elements of the videos' appeal: they feature an authentic

and believable company spokesperson, they showcase genuine product attributes, and they're fun without being pushy. And since Blendtec didn't actually have a consumer following prior to the launch of the viral video program, one has to acknowledge that the videos succeeded purely on their own merits, and not because Blendtec was a beloved brand that gave its audience what they wanted.

So how, then, to explain the costly signal being paid here? Beyond its initial success, Blendtec focused on building a channel – a YouTube video channel, specifically – over time, based on the specific merits of the channel and the expressed desires of the audience there. Its audience helped to decide what to blend, egging the company on to ever-greater challenges, which the company happily indulged. In a very real sense, consumers collaborated on the marketing strategy itself, setting its own terms for participation. While other companies have succeeded at viral video, few have taken collaboration to this degree.

I am claiming that Blendtec succeeded not because they figured out how viral video works but because they figured out how collaborative marketing works. If the company's focus had shifted to, "How can we go even *more* viral?" instead of "What does this audience want to see next?" they could not have achieved the same success. Because they came from obscurity, they had the advantage of learning from their audience exactly how they ought to behave in a popularity-based system, and they were rewarded with reciprocal cooperation. And lots of blender sales.

Ultimately the prescription for success in viral video comes with the same warning label as we've seen in other forms of popularity-based signaling: Marketers must not hang their hats on quantification. Doing so creates perverse incentives for bad behavior – like paying for Facebook endorsements – and sows the seeds of consumer defection. As viral video marketing matures, we can expect brands to focus more on creating content that inspires loyalty among the brand's core, in an effort to sow evangelism from within. Blendtec's follower base is big, but ultimately the company would have enjoyed the same success had it uncovered a way to market virally only to its core "foodie" constituents – the ones willing to pay \$800 for a blender. There is certainly no harm in its wider following, but it is also not the only viral video model for other marketers to emulate. Others could enjoy proportional success simply by making good videos for their core, who will reward their costly signal.

In the next chapter, I will take on some of the questions raised by Blendtec's success as it pertains to brand identity. My focus thus far has been on marketing, which I will posit is a *tool* for brand identity but is not the identity itself, or even its expression. But the implications of social media

for the marketer-consumer relationship penetrate all the way to the brand core, and so they merit consideration. We've already seen how cooperative games bring the consumer into the marketing laboratory, where media and marketing strategies, creative content, and even the structure of the marketing organization end up on the table. Is the brand itself next?

## Chapter 8: Kapferer's Prism and the Shifting Ground of Brand Identity

*ABSTRACT: Since the concept of branding first came into vogue in the 1990's, its practitioners have insisted that brand relationships are inherently reciprocal, and that the brand identity itself exists in the collaboration of marketers and consumers. But in practice, brand definition has largely been in the hands of marketers and has been transmitted through one-way vehicles like advertising. The social media arena provides the first practical means of true collaboration between marketers and consumers on brand definition. Just as human identity is increasingly defined by social relationships played out in virtual space, brand identity is increasingly defined by a decentralized set of networked perceptions and feedback mechanisms. Jean-Noel Kapferer's Brand Identity Prism, which holds that brand identity occurs in a nexus between corporate image and consumer perception, forms the basis for examining the profound shift in the power dynamic toward crowd-based brand identity. Consumers now exert far greater authority over brand perception in the myriad brand conversations taking place in social media, challenging the predominance of well-financed and distributed brand campaigns.*

In conference presentations and webinars on the impact of social media on branding, I have titled my speech "Social Media Killed the Branding Rockstar." The title is meant as a tribute to the Buggles' pop classic "Video Killed the Radio Star," which ushered in a sea change in the music industry as the first video to appear on MTV. But it is also meant as provocation to the field of branding, which faces a sea change of its own with the growth of social media.

The notion of a "branding rockstar" refers to the emergence of branding as a dedicated and specialized marketing practice over the last two decades, which has in turn given considerable clout to agencies that claim ownership of the esoteric business of defining a company's brand identity. While

the current recession has slowed investment in branding, no branding rock-stars have died as a result, and most of the implications of social media for branding are only just now starting to be realized. This chapter will examine those implications. In doing so, the question I wish to raise is not whether branding is still relevant – this would be, as I'll explain, like asking whether competition is still relevant – but rather how branding itself and branding practices will be fundamentally altered by the changes taking place in the media landscape.

## 8.1 Branding by Definition

I begin with an apology regarding how I need to begin: sorting through definitions in order to get at a topic can be a tedious business, but branding makes it essential. Perhaps no other term in the marketing world has suffered so much impoverishment of meaning from overly loose usage, though “ROI” is quickly gaining ground. To confess my own culpability in this degeneration: I must disclaim that my use of “brand” in past chapters is meant to refer to a company's public representation of itself as a player in the marketing game, and nothing more. In using the term, I make no claim regarding the success or failure of the company in achieving brand recognition.

Indeed, the marketing industry as a whole can't agree on what constitutes success or failure in achieving brand recognition, but examining those disagreements should get us closer to a usable definition. Two organizations – Interbrand and Millward Brown's BrandZ – run influential, competing annual rankings of top brands. BrandZ's methodology examines the company's financial worth combined with consumer brand loyalty data gathered with the company's own proprietary methods. Interbrand's methodology is more strictly focused on financial outcomes, even assigning a dollar value to key measures of brand loyalty.

The results are revealing of how we think about brands. Brandz and Interbrand rank Microsoft 2<sup>nd</sup> and 3<sup>rd</sup>, respectively. But the Brandz list puts Microsoft's much smaller arch-rival Apple close on its heels, at 6<sup>th</sup> place, while Interbrand ranks Apple all the way down at 20<sup>th</sup>. There is no question that Microsoft is much more ubiquitous; its software powers most of the world's personal computers, making its presence nearly inescapable. But “inescapable” may be exactly the right word; Microsoft has a virtual monopoly on PC operating systems, so for many consumers, there may be a very wide gap between using Microsoft products and feeling loyal toward Microsoft.

Apple, by contrast, appears to inspire much greater loyalty in its smaller customer base, as anyone who has been waylaid by a Mac fanatic and made to hear about its virtues can attest. Individual Apple devices like the iPod have inspired dozens of fan sites, while Microsoft's competitive Zune product receives no such comparable attention. I should acknowledge that neither ranking measures brand loyalty alone, but the greater emphasis on consumer attitude in the BrandZ formulation undoubtedly accounts for Apple's better showing on that list.

The Apple vs. Microsoft brand question illustrates the more fundamental question of what branding is all about: is it a matter of loyalty or ubiquity – or in the case of the two rankings, some combination of the two? Building ubiquity vs. building loyalty not only demands two different kinds of effort, they may at times be in inverse relationship to each other: Apple can command more loyalty by focusing on the needs of a more niche audience, and a previous chapter noted how Starbucks had a harder time maintaining loyalty as it gained ubiquity.

The question is essential, because the answer will to no small degree determine the future course of brands' use of social media. Based on the analysis offered in the previous chapters, it will come as no surprise that I come down on the side of defining brand primarily in terms of loyalty, and more particularly that brands' use of social media must focus on building loyalty rather than ubiquity. This is not to make an absolute virtue out of loyalty – focusing on ubiquity, even at the expense of loyalty, may at times be in a brand's self-interest – but rather to make a case for social media's role in branding based on what the medium can actually do. In other words, I intend to establish that brands that use social media to build brand loyalty will find the medium well-suited to their purposes, but the same cannot be said for brands seeking to use the medium to build ubiquity.

As evidence, I offer the many examples from previous chapters of brands engaged in accidental defection, prompting reciprocal defection from consumers, because they treated increased ubiquity in social media as an absolute good. Motrin was willing to be controversial, even at the risk of alienating its audience, in order to gain viral views. Target was willing – and then not willing – to pay for endorsements on Facebook in order to increase their exposure there. While the threat of consumer defection created by this overstepping is primarily a subject for the next chapter, my point here is to highlight how mixed definitions of brand success can lead to ill-considered use of what might otherwise prove to be a very powerful branding medium.

I believe most brands realize they'll get better results from Facebook, for instance, if they focus on nurturing the relationships that form there rather than signing up as many fans as possible. But these same brands



face a perverse incentive to quantify; a *Marketing Sherpa* study showed that the inability to prove ROI is the second most common reason brands resist investing in social media. The burgeoning social media marketing industry has, with the best of intentions, tried to help marketers overcome that obstacle by formulating common success metrics for social media engagement, and these metrics are, obviously and unfortunately, quantitative in nature.

Thus we find a new contender lining up beside BrandZ and InterBrand in producing a brand ranking index. *Ad Age* commissioned the research firm Infegy to analyze the publication's top 100 brand list and produce a new ranking based on the number of positive/negative mentions for each across all social media formats – an emerging area of analytics known as “sentiment analysis.” In this ranking, Apple shoots to the top of the list with 920,000 mentions in a month; Microsoft had 40 percent less in the same month. However, Microsoft had 79 percent positive mentions vs. 75 percent for Apple (Neff 2009).

Setting aside for the moment what is inherently specious about positive/negative sentiment analysis, the obvious problem is that any ranking based on quantity of mentions provides no insight at all into brand loyalty; Apple could top the list one month purely on the basis of having released a new operating system, and Microsoft could do the same the following month. It gets us no closer to understanding consumer cooperation in social media and its role in building loyalty. In fact, it is rife with perverse incentives to take the opposite course. In a given month, McDonalds could potentially create a huge surge in its social media mentions, all of them positive, if it released a coupon for a free Egg McMuffin through Twitter, and the marketer responsible could reap accolades for having taken McDonalds to the very top of the social media brand chart. But the impact of the free Egg McMuffin on brand loyalty, while it is not zero, is also nowhere close to what is implied by the #1 ranking. It is simply a way to game the system, because the system has easily exploitable flaws.

## 8.2 Branding as Reciprocal Relationship

To get to a more useful definition of branding – one that is consistent with what social media marketing can best accomplish – we have to go back to the way the discipline has been defined as a discipline by its main practitioners: branding agencies and corporate marketers in branding roles. Since branding first came into vogue in the early 90's, its practitioners have consciously sought to define it as something apart from advertising. Advertising

could serve as a vehicle for conveying ideas about the brand, but it was not the brand itself. In hiring a branding agency, the brand is implicitly asking the agency to tell them who they are. The identity that emerges from that process is indeed expressed in the advertising, but it is also expressed in the corporate culture, customer service, signage, attitude etc. It is the brand's DNA.

This is why branding has succeed in defining itself as an estoric yet vital science – what could be more important than knowing who you are? When an agency tell its client, “Before we develop this campaign, we really need to focus on your brand identity,” the client hears this as “This is all about to get a lot more expensive,” and why not? If branding can do what it claims, it is a specialization akin to neurosurgery on the brand, as opposed to advertising, which would be more akin to cosmetic surgery.

But unlike surgery, in which the practitioner does all of the work, branding is highly reciprocal. The identity that the specialists develop must align with the expectations of the consumer, or the branding effort will fail. And here we get to a definition that I believe will suit our purposes in examining social media's role: *branding occurs when the company's projected self-identity aligns positively with the consumer's projection of the company's identity.*

Under this definition, we might reach very different conclusions about Microsoft and Apple than the leading indexes did. Microsoft appears to struggle greatly with aligning its self-definition with consumers – an inherent difficulty for a company of its size and breadth. Apple, on the other hand, appears so well aligned that its customers manifest a strong sense of ownership of the brand – a phenomenon I'll take up in more detail in a few pages. Apple's high degree of alignment has allowed it to branch into other areas – personal music players and phones, for instance – while holding on to loyalty and fulfilling customer expectations that the brand's core identity will translate into these new areas. Microsoft, by contrast, has not been able to define what is essentially and authentically “Microsoft” as it enters new arenas, and that is a material weakness: it means that customers will not automatically follow Microsoft into new ventures on the basis of its name, and that hurts the company's ability to diversify.

Indeed, one very useful way of examining the strength of a particular brand is to observe how readily customers will follow the brand into new arenas, and that includes social media. I noted in the last chapter how brands with high consumer engagement would pay a lower costly signal for participation in social media than brands with low engagement, and the same holds true for all of a brand's efforts to evolve with its customers. The CEO of Virgin, Sir Richard Branson, has led his company to a remarkably broad

diversity of ventures under a very cohesive brand identity: the Virgin name is applied to everything from airlines to cell phones, and however different those industries may be, consumers have a sense that their experience of Virgin across these industries will have a common thread that is somehow quintessentially Virgin. Branson's perspective on branding is that it is implicitly contractual – a two-way relationship with each customer, based on a set of agreed-upon characteristics.

I suspect that most brand practitioners would agree with Branson's definition, though the implications of it may be more radical than many brands are ready to take on. If branding is a contract, it is reciprocal to an absolute degree; it is not something that the marketer transmits and the consumer agrees, but rather something that they actively agree on. In game theory terms, this would constitute a 4-4 equilibrium; each party gets exactly what they need from the other, with nothing left over. In this idealized contractual state, there would be no risk of defection at all, because all other outcomes are sub-optimal for both parties.

But this idealized state doesn't exist even for brands with high loyalty; as *Ad Age* and *Infegy* will tell you, Apple scores tens of thousands of negative mentions in the social media space, alongside many more positive ones. While that scoring out not to be treated as a complete measure of brand loyalty, any degree of consumer dissent has to be treated as something less than perfect synergy between brand offerings and consumer needs. So Apple still has to worry about defection, and Apple's customers still have to worry about whether the company is doing right by them in all cases.

The primary reason for the gap between the ideal and the real is that a 4-4 equilibrium requires significant coordination, as the chapter on coordinates games showed, and this in turn requires constant communication. And while branding by its purest definition is supposed to be a two-way street, it has by necessity been a mostly one-way conversation.

### **8.3 The Traditional Limits on Brand Engagement**

What I mean is that in traditional marketing, a brand could not, for all practical purposes, maintain a continuous dialogue with its customers, and its ability to communicate its hoped-for identity was limited to a) its direct interactions with customers and b) mass-media channels, where advertising served as its vehicle. If branding is a contract, then branding's main delivery mechanisms are a very poor way of negotiating a contract. Entire campaigns may be built, rejected, modified, reassessed, etc., while the brand tries to keep pace with ever-changing consumer sentiment.

If traditional branding is a non-stop struggle between consumers and marketers to make themselves understood, but the potential payoff is a 4-4 equilibrium and the sort of loyalty that only a handful of brands inspire, then it is no wonder that well-regarded brand practitioners have achieved such vaunted status. But in practice, the toolsets have been limited. In traditional branding, brands could take the pulse of consumers through continuous primary research, such as focus groups and brand tracking surveys, they could focus on imbuing their brand values throughout their corporate culture, to ensure that every customer touchpoint remained true to the brand identity, and they could ensure that their advertising powerfully expresses the brand values that the research indicates it should. This is, to all practical purposes, what a full-scale branding effort does. And these are worthy activities that are in aid of equilibrium, but the gap between the ideal and the real remains.

Advertising has some insurmountable weaknesses as a branding vehicle, if one accepts the mutual, reciprocal, and contractual definition of the term. As I established in Chapter 3, advertising is *inherently* sub-optimal; in general, consumers would prefer not to be advertised to, and marketers would prefer not to spend money on advertising. It is merely a compromise between consumers' need for free content and advertisers' need to reach consumers. The minimax/maximin arrangement limits advertising's ability to tell a brand story; a minimum requirement for good storytelling is an audience's willingness to listen.

While some brands excel in the use of advertising in a way that at least partly overcomes these weaknesses, many brands that excel at building brand relationships do so with very limited advertising. In fact, I will argue that in such cases, *the lack of advertising is implicitly part of the brand contract*. In other words, not advertising becomes a way of cooperating. I am thinking of brands like Harley-Davidson, which has become legendary in marketing circles for achieving rabid brand loyalty and significant market growth with virtually no advertising spending. It has focused instead on cultivating its owners' groups. And as previously noted, Starbucks has eschewed significant advertising in favor of developing its in-store experience. Starbucks' locations are ubiquitous enough that further exposure through increased advertising is unlikely to make a significant impact on awareness and even less on loyalty. In fact, if Starbucks' biggest danger in its customer relationships is appearing too corporate, then increased advertising would run the risk of exacerbating that problem.

Advertising has limitations as a branding vehicle, yet at the same time, branding does need its vehicles. It is not a passive process, this negotiation of marketer-customer contracts. Both Harley-Davidson and Starbucks

substituted other customer-centric activities that helped build their brands. So what other options does a brand have?

By now you have surmised that I am stacking the deck against what traditional branding can accomplish vs. its stated claims in order to set the stage for what social media can potentially accomplish in the branding arena. But I don't believe social media is any kind of panacea for the fractures between brand goals and marketing vehicles. Rather, it represents both an opportunity and a threat for brands, and the outcome depends on how consumers exercise their new-found authority in the branding arena, and how marketers cultivate those brand relationships.

#### **8.4 Consumers Assert Brand Ownership: Nike**

I'll begin with two examples that illustrate the change that has occurred, and how a single instance can be both an opportunity and a threat, depending on the brand's response. The first is a minor but telling one: in 2007, two brothers launched an online petition to convince Nike to design and release a pair of basketball sneakers based on the ones worn by the "Marty McFly" character in the film *Back to the Future II*. The *Back to the Future* film series is iconic among Gen-Xers nostalgic for 80's style, and the petition quickly grew to 50,000 signatures. 50,000 shoe fans are a force to be reckoned with, but Nike controls 85% of the domestic basketball shoe market, so by any purely quantitative measure, the shoe giant could afford to ignore the McFly petition (McCall 2007).

But they chose not to ignore it. Instead they developed and released the shoe in 2008 as a limited edition, to the delight of the petitioners. The story of the brothers' triumphant effort got picked up in major media outlets, and Nike scored a PR coup.

This type of consumer empowerment story has become so commonplace in social media that it's easy to overlook the deeper implications of the trend. We are now in an era when a consumer seeking to dictate terms on what a company should design and sell not only has a reasonable chance of being heard, but an outside but still reasonable chance of having their wishes fulfilled. This the notion of "brand" is entirely caught up in what a company sells, we are witnessing in the Air McFly incident just one of many ways in which consumers end up working side-by-side with marketers in the brand laboratory.

Nike deserves full faith and credit for letting the consumer into the marketing laboratory to collaborate on the project. But I must also point out that Nike risked very little in the collaboration. The petition had more than

50,000 signatures; Nike released 1,000 pairs of the shoes (Sosa 2008). One would not need a complex formula to calculate that the shoes would quickly sell out and become collectors' items, as they did. It was an instance of 4-4 equilibrium in the brand collaboration game.

The Air McFly incident clearly exemplifies the *opportunity* aspect of increased consumer empowerment in branding, but what about the *threat*? I believe the threat comes from expectations thwarted or fulfilled. This threat is lower for Nike because they're practiced at this form of consumer engagement; they can easily produce limited-edition shoes, and they do so primarily as a loyalty-building tool anyway, rather than as a key revenue driver.

But I find this threat implicit in the surprisingly jaded response that the new shoe elicited in some circles. The influential gadget blog Gizmodo announced the shoe's arrival under the headline, "Nike *Finally* Releasing Back to the Future Part II McFly Sneakers, *Sort Of*." (emphasis mine). The article is a case study in consumer righteousness. It complains that fans had been clamoring for the shoes "for years" (the period between the start of the petition and the shoe's *release* was actually about one year), and that the shoes were not direct replicas but were merely inspired by the movie. The essayist huffs that "It's a start," but declares that he will hold out "for the real deal" (Frucci 2008). So the expectation remains unfulfilled.

Such a sense of entitlement is easy to make fun of, but as part of a larger trend, it also demands to be taken seriously. Nike's sneaker rival Adidas is the target of an online petition – with nearly 5,000 signatures and counting – demanding that the custom Adidas sneaker featured in the Wes Anderson film *The Life Aquatic with Steve Zissou* be produced for purchase (Wloszczyna 2009). Adidas has not said whether they'll produce the shoe. If Adidas does not develop shoes demanded by online petitions, and Nike does so, does Adidas' lack of compliance alter their brand stance? If an otherwise highly collaborative brand doesn't comply with demands for new products or product changes, or even in the case of Nike and the Gizmodo blogger, if the compliance isn't total, does this constitute an act of defection, with negative consequences? These are the difficult question raised by this new branding arena.

## 8.5 Consumers Assert Brand Ownership: Apple

It could be argued that while these shoe petitions do demonstrate an unprecedented degree of consumer empowerment, they don't penetrate very deeply into brand territory; after all, it's not as though consumers tried to *design* the shoe themselves. Enter our next example, the inescapably popular Apple

iPhone. The phone is known for its rabid fan base, but the striking thing from a brand control perspective is that it acquired a portion of that rabid fan base before the phone had even been released. The phone was announced in January 2007 but not released until June 2007, and in that time period, Apple fans asserted an astonishing degree of brand ownership. Dozens of blogs popped up containing nothing but speculation about the new phone, but most striking of all were the user-generated prototypes. Fans developed their own prototype designs, rendered in painstaking detail, based mainly on their aspirations for what the new phone should be. The design concepts range from the sophisticated to the sublime, and many are archived on the blog [appleiphone.blogspot.com](http://appleiphone.blogspot.com).

Once again, the opportunity is obvious; the threat, less so. The speculative prototyping by Apple's fans was all in good fun, and it contributed to the excitement that led to the sale of as many as 700,000 of the phones on the first weekend after its release. If one of the primary goals of good branding is brand evangelism, i.e., the willingness of loyal customers to advocate for the brand of their own accord, then Apple achieved that here as well.

But as with Nike, Apple is being handed a higher standard for consumer engagement. While the Apple fans who developed the prototypes wouldn't actually expect their ideas to become part of the phone's design, their efforts function as a signal to the brand that the customers want a seat at the table – not merely at the product feedback table, but at the brand definition one. For an engaged brand like Apple especially, the newly empowered consumer expects to become part of a continuous feedback loop about the brand direction.

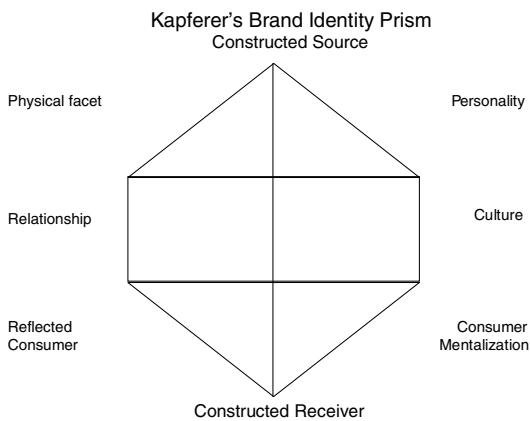
Again, this seemingly revolutionary concept is consistent with a decades-old, widely accepted notion of branding as a deep collaboration between marketers and consumers. The difference is that this notion was an idealization of a relationship that was mostly a one-way conversation: lots of output from marketers, and little bit of input from consumers. That inequity has now been removed, and marketers must now ask themselves whether their ideal can withstand the shock of the real. If branding is truly now a two-way conversation, is it one that marketers should want to have?

## 8.6 Kapferer's Prism

To answer that question, we need a theoretical model to help explain how the two-way conversation should work. For this, I need to reach outside of game theory (but not to worry; I'll reach right back in) to some of the original conceptual work on branding.

In 1992, the French marketing theorist Jean-Noel Kapferer introduced an explanatory model for branding that involved a multi-faceted prism, now popularly known as “Kapferer’s prism.” The model is a direct expression of branding as a collaboration between marketers and consumers (Kapferer 2008).

In Kapferer’s model, brand identity takes place in the territory mutually established between a source (the marketer) and the receiver (the consumer). A brand has both a physical dimension (actual products and people, as shown on the left side), and an emotional dimension (the brand’s idea of itself, and the consumer’s idea of the brand, as shown on the right.) The top and bottom portions of the prism are merely what each side brings to the table: Nike brings its Air McFly and the retro-hipster personality that accompanies it; the sneaker fan brings themselves and their idea about themselves as a retro hipster. The brand’s identity – as it pertains to this particular customer and product, not necessarily to Nike as a whole – takes shape in the place where Nike and the sneaker fan meet, agree upon the shoe’s relationship to the customer’s idea about themselves, and agree on a set of shared values. If the collaboration is successful, a relationship is formed around these shared values, e.g., we both love great footwear, we both love retro style, we both think sneakers should be an outlet for self-expression and individuality, etc. Kapferer’s model is therefore highly consistent with Richard Branson’s aforementioned definition; it suggests a contract in which both parties are equally represented.



**Figure 4:** Brand identity prism

Kapferer felt that marketers’ main difficulty in working within the terms of the prism was their ability to separate the audience as a target vs. the audience as a reflected consumer. As a target, the consumer is merely a



representative of the people who want to buy the product, e.g., 15-34 year-old males; this is useful for media-buying and not much more. A reflected consumer, by contrast, is an expression not of what the audience *is* but of what they *want to be*. It is aspirational. In Kapferer's model, the marketer who focuses on the reflected consumer rather than the target is able to tap into the deeper emotional needs that the brand can fulfill, and consequently that marketer will succeed in forming a brand relationship and successfully signaling a cohesive, relatable identity.

The traditional weakness that Kapferer uncovered stems not from marketers' lack of understanding of how branding is meant to work, but rather from marketer's lack of insight into consumer mentalization. Uncovering latent desires is difficult and specialized work; thus we see the emergence of the specialized brand agency on the heels of the Kapferer model. Going after a consumer's stated preferences isn't enough, because in psychoanalytic terms, the consumer's actual preferences may be hidden to themselves.

## 8.7 Rapaille and the Brand Imprint

The market researcher Clotaire Rapaille is famous for his unique methods of uncovering latent desires and translating them into brands. His "imprinting sessions" with potential customers borrow from psychoanalytic practices in stripping away the layers of participants' conscious desires, in order to get at the "imprints" that were formed in childhood. For instance, he persuaded Chrysler to return to round headlights on their new line of Jeeps, because he discovered that people imprinted on the idea of a Jeep as an animal, like a horse, and subconsciously felt that the Jeep's grill needed to be more face-like.

In practice, most brands don't have a Clotaire Rapaille mapping out the territory of consumer mental models. They rely on field research, focus groups and other feedback loops to create a cycle of trial-and-error that, ideally, moves the brand toward the territory of shared cultural values with their consumer. Getting there not only requires the brand to do the tough work of accurately uncovering the consumer's mental model, it also requires that they do so without the kind of interference that so often muddies the waters in complex organizations. If Rapaille had discovered something vitally and verifiably important in the notion that Jeeps should have round headlights, that insight could still have crashed into the prejudices of a Chrysler executive who insists that he or she could never accept the outdated styling of round headlights.

I am belaboring the notion of branding as a delicate dance because, of course, we have a science devoted to delicate dances in game theory. While the tools used may be unique to the field of branding, the process of iteratively uncovering and agreeing to a set of brand qualities – that is, the processing of mutually executing the brand contract – is nothing more or less than a coordination game like the stag hunt. As I've stated, the idealized outcome of proper branding is indeed a 4-4 equilibrium, where consumers relate to the products and brands that map to their aspirations about themselves. Both Chrysler and the car buyer are hunting stag if the headlights – alongside myriad other brand signifiers – fulfill both Chrysler's aspiration to sell more cars and the consumer's aspiration to own a vehicle that makes them feel wild and free.

## 8.8 The Limitations of Traditional Branding

In traditional branding, then, the marketer faces two main challenges in getting to that 4-4 equilibrium: 1) a complex coordination game that involves uncovering latent desires and aligning a large organization behind the fulfillment of those desires, which is an *information* challenge, and 2) a partial reliance on sub-optimal delivery mechanisms, i.e., advertising that consumers prefer to avoid, which is a *delivery* challenge. Of course, if the marketer resolves the first challenge, it makes the second challenge somewhat easier; consumers respond positively to advertising that aligns with their idea of themselves, though they would still prefer to avoid the advertising in an optimal scenario.

The traditional model is based on a certain amount of passivity on the consumer's part: one can envision the consumer parked in front of a television set, absorbing branding signals with each ad, and counter-signaling only with their wallet when they go to make purchases. As I've begun to outline with the Nike and Apple examples, we are entering an era in which this traditional vision has been replaced by one of active and engaged consumers whose input on the brand goes far beyond purchase data and the occasional focus group into sometimes strident demands for how the brand should behave.

Thus the new, social media-based branding territory that I've begun to outline may have the potential to resolve branding's traditional challenges. The *information* challenge is answered by a vast sea of consumer data in the millions of brand conversations occurring in social media each day. The *delivery* challenge is answered by the opportunity for brands to engage with consumers directly on their own terms in social media forums. But

these answers are not perfect; they bring significant coordination challenges regarding how brand data is processed and how online brand communities are engaged. I'll take up each of these two topics – processing the new brand data, and engaging new brand communities – in order.

## 8.9 Introducing Brand Monitoring

The brand-related conversations in the vast sea of social content is indeed uncharted territory, but serious efforts are underway to chart it, or at least to provide navigational tools so that marketers can chart it on their own behalf. The last three years have witnessed the rise of the “social media monitoring tool,” often referred to as the “brand monitoring tool.” As of this writing, approximately 8-10 major providers are vying for dominance in this space. It's a daunting task, because it will require that the firm or firms that emerge as the top players establish a common currency of measurement that its audience of marketers will accept.

The current climate is reminiscent of the emergence of Web analytics tools a decade ago. Today there are only a handful of dominant players in Web analytics – principally Omniture, WebTrends, and Google Analytics – and while there are variations in data collection methodologies across these dominant players, the primary metrics that marketers care about are shared across all of them. It's not quite the level of standardization of, say, VHS or DVD, but it's as close as marketers generally get to agreeing about anything.

Brand monitoring tools are probably several years away from this level of standardization, but some standards have already emerged. The tools are essentially specialized search engines, allowing for keyword-based searches across a broad spectrum of social venues. A brand like Coca-Cola would, for instance, get a broad measure of its level of social media brand conversations by searching for “Coke” or “Coca-Cola” on one of these specialized engines.

With some competitive variation in terms of which specific social venues get indexed, all of the major players include the expected social networking venues: blogs, forums, networks, video channels, etc. To varying degrees, the tools allow marketers to slice and dice the data to their purposes, by uncovering key topics (again, based on keyword density) and allowing marketers to trace the thread of specific conversations over time.

The two very broad metrics that have already emerged as near-standards across providers – and more importantly, among marketers – are 1) number of brand hits, and 2) the ratio of positive to negative sentiment in brand hits, known as “sentiment analysis.” You'll recall that these two metrics were the

basis for *Advertising Age's* indexing of the top brands in social media, so their credence is gaining wide acceptance.

Both metrics are appealing – what brand wouldn't like to know at a glance how popular they are? – but both are problematic. The first metric runs up against my perpetual gripe against quantification, because it treats the number of brand hits as an absolute good. The perverse incentives in this metric are obvious, as I noted earlier with the Egg McMuffin hypothetical scenario. One can readily envision a marketer spamming his/her social network venues in the week before a quarterly report is due in order to boost the raw numbers in this all-important new medium.

The second metric, sentiment analysis, is meant to be a hedge against raw quantification: it suggests that what matters is not simply raw hits, but the level of positive brand affirmation within those raw hits. Good branding efforts would presumably cause the ratio of positive and negative hits to move toward the positive. There are several problems here. The first is that most insightful brand sentiments don't fit neatly into black-and-white categories. A consumer might say, "I like the new Jeep, but the round headlights feel outdated to me." It's a very useful piece of data, unless it's being forced into a bifurcated system, where it appears ambiguous at best.

And that raises the second problem with sentiment analysis: because of the sheer quantity of brand content in social media, most of the providers use "machine analysis" to assign sentiment. In other words, the search engine looks for the recurrence and density of certain terms (such as "like" and "hate") that are deemed to be indicative of sentiment, and they categorize the results in that way. The majority of results are not conclusively positive or negative, and they end up in an "other" category.

As one would expect from any machine-based effort to understand the subtleties of human emotions, sentiment analysis can produce odd results. In a presentation of the limits of this type of analysis, I demonstrate how one monitoring tool showed similar ratios of positive to negative sentiment for both "Chicken McNuggets" and "bird flu." Since many people like Chicken McNuggets, but no one seems to like bird flu, this outcome seems a bit flawed.

To be fair, monitoring providers acknowledge that sentiment analysis requires human validation. The most expensive solutions include a degree of human analysis; they take a sample of the data set, screen it with their own linguists, then apply those results to the entire data set.

I am for the moment dwelling on the toolsets for measuring conversations rather than the conversations themselves because I believe that marketers' use of these tools will greatly influence the future course of brand engagement in social media. To regard brand monitoring as a mere tactic within

the more substantive work of brand engagement is like treating Google as a mere search tool, rather than as the lens through which so much of the Web comes into focus.

As brand conversations grow exponentially, the need to make sense of them and their impact on branding efforts will grow proportionally. In this environment, brand perception can shift more rapidly than occurred in the more controlled environments of traditional advertising and direct customer experiences. Even a highly engaged brand like Target can impair a hard-won reputation with a handful of minor missteps.

In these rarified circumstances, brand monitoring is a critical component of the coordination game. Recall that a coordination game starts from a basis of imperfect knowledge about the other player's moves and tries to overcome that information gap by anticipating the other's actions. Any information on the other player's preferences and predilections – where they prefer to do their stag hunting, for instance – increases the game's chances of success.

Brand monitoring helps to fill that information gap in two important ways: it provides fertile ground for *ethnographic analysis* that gets the marketer closer to understanding the hidden needs and motives of the other player, and it helps the marketer to anticipate preemptive acts of defection that threaten to derail the coordination game and trigger the dreaded death spiral. I'll consider each of these roles in turn.

## 8.10 Brand Monitoring and Ethnographic Analysis

Ethnographic analysis is one of the tools that brand marketers have borrowed from cultural anthropology; it involves *in situ* field analysis of the subject in order to observe their real-world behaviors around the brand outside of the artificial constraints of the focus group. Its appeal to brand specialists lies in its ability to uncover first-hand the spontaneous, unprompted, unrehearsed behaviors of real people interacting with products. One can easily imagine, for instance, that spending a month riding with a local chapter of Harley riders is infinitely more useful to the coordination game than sitting down with a group of bikers around a conference table in a test lab. But the former may also be prohibitively expensive and time-intensive for the brand to undertake, assuming it is even possible for the brand marketer to participate in a way that doesn't taint the results.

Online ethnographic research has some advantages. It cannot, admittedly, provide the visceral experience of real-world, unfiltered brand interactions, but it is unfiltered nonetheless. A brand marketer who spends an

hour a day for a month observing the online interactions in a Harley's owner group stands to gain a wealth of unfiltered insight on what the brand is doing right or wrong; multiply that by the number of viable owner group social networks, and the brand has quickly exceeded what it could achieve by a far more anecdotal field study conducted through a "real life" social media group. Any marketer who has spent time absorbing the wincingly frank and honest brand input that occurs in forums that *aren't* run by the marketer can appreciate the qualitative difference in the feedback in comparison to a focus group.

Take the example of Cuesta Verde and their pricing problem. In a focus group setting, a participant might raise cost sensitivity as an issue, but there could be good reasons not to do so: the participant may wish to avoid the appearance, in front of their peers, of being willing to skimp on the care they provide to their aging parents, or they may simply latch on to other issues that are dominating the group discussion, in the classic problem of cognitive dissonance. Cuesta Verde may then falsely conclude that they don't have a pricing problem. Their failure to recognize the issue and to address it in their brand positioning – by emphasizing value for the investment, for instance – may entirely prevent them from achieving the 4-4 equilibrium of a successful coordination game.

If the ethnographic component of brand monitoring is so crucial to the coordination game, why have so few brands adopted the practice? I believe it is mainly because organizational structures have not evolved sufficiently to take advantage of this changing brand environment. Brand monitoring falls to digital marketers simply because the medium itself is digital, and brand marketers remain mired in traditional methods. Digital marketers are more apt to try to quantify rather than qualify what's occurring in brand conversations online, because that is the predilection of their discipline, and the means by which they evaluate success. When digital and brand marketers finally collaborate on uncovering brand conversations and mining their content for insights, the true value of brand monitoring can be realized.

## 8.11 Diffusing Defection through Brand Monitoring

In the meantime, though, the second contribution of brand monitoring to the coordination game remains well within the grasp of every marketer: the ability to anticipate and diffuse preemptive acts of defection by angry consumers. When the Motrin backlash unfolded over the course of a weekend, many social media marketing commentators savored the opportunity to offer analysis of the barn door after the horses had escaped, i.e., they noted

how the spiral might have been prevented with a good brand monitoring process.

It could be argued that incidents like the Motrin Moms are PR crises, not brand crises, because they usually occur in reaction to a specific provocation, which is the sort of thing that PR firms trained in crisis management are especially adept at handling. And indeed, there is ample cause for involving PR teams in the "first alert" chain for erupting brand crises, especially as such incidents reach the ear of major media, as the Motrin incident did. But I will argue that the brand impact is just as important, because backlash lives on as indexed content, a kind of permanent blemish on a brand's online reputation.

The endurance of negativity is one of the great challenges of the contemporary brand coordination game, and it underscores the need for brands to be fully engaged in the medium. Consider a point of contrast: the Jack-in-the-Box fast food chain nearly went bankrupt in 1993 after an *E. coli* outbreak in its restaurants killed 4 children and sickened hundreds of customers (Martin 1998). But the incident occurred in the pre-Web era; today the chain is thriving, and a Google search on the brand returns no specific hits on the incident in the first page of results. On the vastly more trivial subject of Motrin's lightly mocking video, by contrast, a Google search returns a link to the offending video in the top 10 results, under the headline, "Controversial Motrin Moms Commercial," even though the video only "lived" for a few days. The stakes of the coordination game have changed.

In light of these high stakes, monitoring social media for eruptions of negativity is an essential function for brands to take on, but there is a real danger that the high profile of such incidents will stoke unwarranted fears among brand marketers regarding social media engagement. In practice, most of the activity that has been classified under the "brandjacking" umbrella does not rise to the level of a PR crisis. The brand monitoring service MarkMonitor issues an annual report on brandjacking; they reported nearly half a million incidents in 2008, but more than 400,000 of these were simple acts of cybersquatting, i.e., registering a social media outlet under a brand name for which one does not own the trademark. Less than 2% of the incidents actually involved the use of offensive content (MarkMonitor 2009).

In practice, brands have been very successful in invoking trademark laws to protect themselves when brandjacking occurs. In one notorious incident, a Twitter user known only as "Janet" represented herself as a spokesperson for ExxonMobil and answered questions on the company's behalf for several days before the company invoked their trademark and prevailed on

Twitter to shut the brandjacker down (Diaz 2008). But the more interesting cases are the ones in which brands adopt a “TIT FOR TWO TATS” strategy by tolerating acts of brandjacking. Starbucks has been the target of several brandjacking incidents, including a fake ad in which a young woman enthuses about Starbucks’ Frappuccino beverage before cheerfully pointing out that the cost of a Frappuccino could feed a child in a Sudanese refugee camp for a week. Social media observers predicted swift legal action, but Starbucks has apparently refrained; the video remains on YouTube after 3 years, with over 200,000 views.

Without being able to ascribe specific motives to Starbucks in dealing with such incidents, I will nevertheless argue that forbearance is a smart strategy that furthers the goal of the 4-4 equilibrium. Even if a brand succeeds in having critical content removed – an outcome that is by no means assured where satire is involved – the incident’s fallout and the reams of meta-commentary that it generates live on indefinitely. Forbearance sends the signal that the brand is strong enough to take it, and that their participation in social media is a matter of taking the bad with the good. Anyone who visits YouTube to watch the Frappuccino satire is also presented with links to much more reverent user-generated videos about Starbucks, including one in which Starbucks fan successfully visits all 171 stores in Manhattan in a single day. The video has been viewed nearly twice as often as the satirical one.

The best antidote to brandjacking is not a monitor-and-respond strategy, though it is infinitely better than doing nothing, but rather an engagement strategy. Here I am invoking the well-established prisoner’s dilemma strategy of cooperating on the first move and being tolerant of initial acts of defection. The temptation to brandjack a company – at least in the instances that go beyond mere squatting – stems at least in part from the perception of a large brand as monolithic, unresponsive, and unassailable. It’s a lot more fun to go after a player that deserves to be taken down a few pegs; this has been a basic rule of satire since the ancient Greeks.

Companies that are highly engaged in brand conversations through social media don’t necessarily shield themselves from brandjacking; they merely defuse its explosive potential by shifting the weight of the conversation toward the positive. For instance, the Venezuelan-owned oil and gas company CITGO has been the target of repeated and vehement denunciations in social media because of Venezuelan president Hugo Chavez’ anti-American rhetoric. The company’s response has been, in part, to create an online user-generated content contest that rewards participants for acts of charity within their communities. Notably such a strategy does not at all attempt to engage the Chavez issue directly; rather it simply seeks to start a



different, more positive conversation, and to put company resources behind driving that conversation forward. Thus the effort has been effective not in making the Chavez comments disappear – no social strategy can accomplish that – but in diluting them within a larger conversation about the company's efforts to reward good deeds.

## 8.12 Crowdsourcing Brand Identity

But as I noted earlier, brands can go much further than merely defusing negativity in this new brand environment. If a well-coordinated brand, by Kapferer's definition, involves direct collaboration on both the brand's physical artifacts and its emotional content, then social media may be the *first* medium to make this level of collaboration authentic, practical, and perhaps even necessary to the brand's success. Brands have never before faced the *threat* posed by volumes of permanent and direct brand feedback, and they have never before faced the *opportunity* posed by volumes of permanent and direct brand feedback. It is primarily a matter of how the coordination game is played.

So how should the game be played? In presentations on the subject I have likened the traditional branding process to the ancient practice of building a fortress in the jungle: the jungle must be cleared and the fortress built at great cost and labor, in defiance of nature's encroachment. The fortress must be solid and imposing, and the jungle must be continuously beaten back. The brand, once created, must be promoted so that it rises above the surrounding din, and it must be defended so that it remains inviolate.

The problem is that fortresses are prone to decay. They are an imposition on the landscape, not a natural part of it, and as such, their upkeep requires greater effort for less reward as the jungle asserts itself.

The new branding process that's best suited to the evolving landscape? *Be the jungle*. Learn its ways. Adopt an organic approach that first asks the question, what thrives here? What germinates and grows? The resultant brand will be less singularly imposing because it will be diffuse; it will adapt to local conditions while maintaining its core DNA. It will spread everywhere, without fear of decay. The brand will grow out of the environment but remain conversant with it.

If I haven't hopelessly entangled my argument in the jungle metaphor, let me attempt to literalize it with real-world examples. I believe the best evidence available for the evolution of this new way of branding is the collaboration between marketers and consumers in user-generated content, specifically the crowdsourcing of brand strategy.

Of the many forms of user-generated content now flourishing on the Web – a category that includes product reviews, videos and photos, personal blogging, fan fiction, etc., just to name a few – consumer participation in brand contests is among the most prominent and prolific forms. (I’ll refrain from labeling all of this activity “crowdsourcing,” because I believe crowdsourcing of brand identity occupies a narrower space, as I’ll explain). In any case, consumers show a strong predilection for online activities that serve as an outlet for their own personal expression of a brand’s identity.

This eagerness will not be surprising in the broader context of consumer empowerment that I’ve described in this book; consumers increasingly expect to be granted a high degree of input on branding, especially in the case of collaborative brands that have paid that costly signal. But brand interpretation also has a more deeply rooted appeal. If you accept Clotaire Rapaille’s argument that a successful brand taps into latent desires, then asking consumers to exercise their creativity about brand identity is a double-win for the consumer: they get the pleasure of creative exploration, coupled with the pleasure of wish-fulfillment in making the brand conform to their desires. In other words, it’s fun to come up with a catchy beer slogan; it’s even more fun if the slogan creates a connection between your lived experience and your favorite beer.

The pleasure of this activity is the hidden incentive in cooperation. Without that incentive, consumer participation might seem illogical: why should the same consumers that take pains to minimize exposure to advertising also take pains to participate voluntarily in ad development? It’s useful to remember that consumers unquestionably *want* relationships with brands; they just want them on better terms. This kind of brand collaboration provides those terms. The reciprocal benefit to the marketer is also a double-win: increased consumer loyalty, and valuable branding content that the marketer can actually use.

That last part – using the content – turns out to be the biggest hurdle for brands to leap, but it’s also the most important factor in a successful coordination game. Brand-focused UGC may in fact be the most revolutionary tool that social media offers the marketer, but most marketers will not tap its full potential. The more typical scenario is as follows: suppose you are a well-known fruit juice manufacturer with a tried-and-true brand formula. In order to “activate the base,” you run an online user-generated-content contest and ask users to submit new ad designs and taglines. You get a wealth of thoughtful responses, including some truly break-out creative. You post some winners and give out some prizes.

This is a perfectly legitimate use of UGC, but it is limited in its utility. The consumer recognizes that their participation is conditional, limited to

the contest itself; they're not really involved in the creation of brand artifacts, since the material they create won't live past the contest. The incentive for cooperation doesn't extend much beyond the value of the potential prize. What if instead you ran the ads that consumers had created? What if you started a conversation with your newly recruited brand stewards, soliciting their feedback on new products and campaign ideas? Doing so greatly raises the incentive for cooperation; the consumer then has a true stake in the brand's success.

Taking this final step makes marketers nervous. It's one thing to let consumers into the brand laboratory; it's quite another to let their newly created brand monster break out of the lab and trundle off toward the village. Who's going to explain the crowd of angry torch-bearing villagers to the CEO?

This is why the first brands to engage in this level of brand collaboration are the ones that can afford to pay the costly signal for participation; again, these are not the brands with the largest branding budgets, but the brands with the best track records of collaboration. My own experiences as a marketer include a set of contrasting cases that illustrates this point: I helped to develop UGC promotions for both Columbia Sportswear, a large manufacturer of outdoor clothing with a wide consumer base, and its much smaller subsidiary, Mountain Hardwear, which caters to a loyal core of outdoor enthusiasts. Columbia's UGC content did not live beyond the promotion, but Mountain Hardwear's did, because their tighter brand collaboration lowered the cost of participation.

Mountain Hardwear's contest allowed participants to create their own Mountain Hardwear print ad, using their own words and photographs within a loose template provided to them. Because many of the brand's constituents are possessed of both great adventure stories and the photos to prove it, the contest produced ads that were of comparable quality to professionally produced ones. This is fortunate, since the winning ad ran as a paid advertisement in *Rock & Ice* magazine, as a testament to the brand's commitment to collaboration.

Mountain Hardwear's culture of collaboration is reminiscent of the Vermont-based ice cream maker Ben & Jerry's, whose reliance on consumer input to shape the brand is *itself* a crucial part of the brand's identity. The company's ice-cream line-up contains several consumer-developed flavors and product names; a recent "Do the World a Flavor" contest invited participants to participate in the company's social responsibility efforts by concocting a new ice cream from fair trade ingredients. Far from compromising the company's brand identity, collaboration has helped the company to preserve its culture and its loyal following after the company's acquisition by the consumer packaged goods giant Unilever in 2000; loyal customers

would be hard-pressed to identify any diminishment in their relationship with the brand following the change in ownership.

Now that the payoffs for the brand coordination game have been established in the marketplace, we see larger and more traditional brands following suit and collaborating directly with consumers to develop brand artifacts. One of the more interesting phenomena within this movement is that of traditional brands capitalizing on consumer nostalgia for a bygone era of advertising. In 2008, McDonald's observed the 40<sup>th</sup> anniversary of its Big Mac sandwich by inviting customers to submit new versions of the brand's iconic Big Mac jingle (Johannes 2008). In doing so, McDonald's was tapping into a large vein of nostalgia – the jingle was introduced in 1974 and was ubiquitous on network television for many years – for a period of retrospective innocence in consumer advertising, when a smaller number of advertisers dominated a much smaller media market.

This may be somewhat blind nostalgia – consumers preferred to avoid advertising then as they do now – but it is also an honest reflection of consumers' desire to reconnect with iconic brands, if they can do so on their own terms. For a new generation of consumers not previously exposed to “classic” advertising, the appeal of a jingle competition is simply the fulfillment of expectation; from their perspective, it is perfectly natural to be asked to collaborate.

In a similar spirit, Dunkin' Donuts announced a campaign to remind consumers of the company's heritage in, well, donuts. (Ironically, donuts have not been a focal point in the company's advertising for more than a decade). The campaign included an opportunity for users to create and brand their own donut (using a provided list of ingredients), which would then be sold through the stores (Odell 2009). Heinz launched a television ad competition that tapped consumer nostalgia for the classic Heinz ketchup ads of the 1970's; the winning ad and four runner-ups ran as paid advertising on cable networks. An astonishing 2,000 video entries were received, many of comparable quality to professionally produced spots.

Most significantly, the contest gave participants *carte blanche* in exploring the brand's dimensions, and the winning ads showcase a range of expressions of the brand's importance in consumers' daily lives. The winner received a \$57,000 prize, and the runners-up \$5,700 each – a fraction of the cost of a professionally produced spot, over and above the brand loyalty generated and reinforced through the process.

In my judgment, these moments of collaboration, when the coordination game achieves its 4-4 equilibrium, represent the furthest evolutionary stage of the marketer-consumer relationship. Where might it evolve or devolve from here? A further evolution might involve big changes in what

we now regard as traditional advertising: advertisers might find consumers less likely to reach for the “skip” button on their Tivo remotes if they have a chance to see and vote on innovative ads created by consumers like them. And brands might be more inclined to take chances on innovative UGC ads if they can produce greater response at lower cost. This would not only transform branding, but advertising itself.

A devolution, on the other hand, would involve brand engagement wearing thin. As more brands pursue social media marketing more aggressively, the risks of consumer burn-out and marketer bad behavior (the “line-cutter” syndrome) increase exponentially. As with banner advertising, diminished consumer response to brand activity in social activity would have the perverse effect of making marketers more pernicious and less honest in their use of the medium. Consumers would stop participating in brand collaboration activities and would retreat to paid, private networks, free of advertiser intrusion.

Evolution or devolution? This is the critical crossroad at which we find ourselves. The mutual benefits of cooperation, laid out over the past 9 chapters, are evident, and so the evolutionary path is promising. But I believe there is an external factor, seldom discussed, that will greatly determine whether further evolution occurs. I am referring to the question of *scale*. Can the growth of social media itself and the rapid incursion of marketers into the social media arena sustain any meaningful level of cooperation? In the next and final chapter, I will take up this question.

## Chapter 9: Maxwell's Demon and the Dwindling Supply of Consumer Attention

*The physics concept of Maxwell's Demon provides an apt metaphor for the increasing demands on consumer attention levied by social media participation; consumers must continually sort relevant and irrelevant content and connections in order to make their participation worthwhile. As more marketers participate at a greater volume in social media, they face the threat of consumer exhaustion; how much of their dwindling supply of attention will consumers devote to brands? The Volunteer's Dilemma, in which players must set aside their short-term interests for the long-term good, illuminates this question. Marketers' increasing demand for quantifiable results can create a perverse incentive to maximize short-term gains, at the risk of alienating consumers in a cooperative arena. The use of "counterreinforcers" that hold marketers accountable to acceptable rules of engagement may prevent mutual defection.*

Why is theoretical physics such a rich source of metaphors for describing the postmodern condition? In seminal postmodern works like Thomas Pynchon's *The Crying of Lot 49* or Terry Gilliam's *Brazil*, the protagonists deal with the problem of information overload: when confronted with more information than we can ever hope to sort through, how do we tell the difference between the useful and the useless, and how do we use information to tell the difference between our allies and those who plot against us? This dilemma is illuminated (though not resolved) by physics concepts like the relative truth of a thing based on the observer's position (Einstein's Theory of Relativity), the tendency to alter events just by observing them (Heisenberg's Uncertainty Principle), and the idea that multiple truth-claims can exist side-by-side in a multiple, infinite universe (Schrodinger's Cat). In other words, we're told that we live in a universe of infinite possibilities and no single, governing hierarchy, so we shouldn't be at all surprised when we can't find a decent steak house using Google Maps.

In my own hopelessly postmodern fashion, I also find it helpful to draw on theoretical physics to explain the problem of information overload in social media marketing. Therefore I offer a metaphor that (for me) reflects the razor's edge we currently walk between a utopian and dystopian future for social media: Maxwell's Demon.

Maxwell's Demon is a theoretical concept invented by the Scottish physicist James Maxwell in 1871. Maxwell was offering a challenge to Newton's Second Law of Thermodynamics, which describes entropy, i.e. the tendency of things to fall apart. He envisioned a box of hot and cold molecules bouncing around; under the Second Law, these highs and lows will eventually even out to an unremarkably lukewarm state, much like network television.

But Maxwell envisioned a creature, a "demon" in the box, whose sole job is to sort the hot molecules from the cold. If the energy the demon uses to sort molecules is less than the amount of energy retained by keeping hot molecules together, the demon can defy entropy and even create perpetual motion (Baeyer 1998). The big question that physicists continue to debate is this: how much energy is used in the act of sorting?

## 9.1 The Problem of Overtaxed Attention

For consumers and marketers alike, success or failure in their online experience entirely hinges on their ability to sort and isolate relevant information: finding your target audience among billions, joining with like-minded communities, locating relevant search results, finding your ex-classmates, and on and on. In this sense, the dominant company in the arena – Google – is nothing more or less than a giant Maxwell's Demon. It doesn't create content; it sorts it, and it enlists millions of lesser demons – you and me – to help it do that.

The utopian/dystopian dilemma for social media is this: if we can create ways to sort information that keep pace with the growth of information, we create a utopia of relevance and connectedness. If we fail to do so, then social media will eat itself: the demands of keeping up with the social stream will outweigh its relevance to users, and they will retreat to fragmented communities of deeper relevance but less connectedness.

Each moment that a user spends on the Web is spent as a Maxwell's Demon. When searching for content, the user must sort the relevant from the irrelevant and click on the best results. On Facebook, they must decide whose posts to read and whose to ignore. On YouTube, they review the popularity of video based on views and scores, then decide what's worth watching and what's not. If the demands of all of this sorting activity become

overwhelming, the user's limited and overtaxed attention begins to flag. They "friend" someone on Facebook that they'd rather not connect with. They ignore ads that are relevant to them while being inadvertently drawn in by ones that aren't, only to regret it afterwards. They become annoyed by all their friends' requests for them to view videos or enter contests, and they begin to ignore them. They later decide that their social media participation is more trouble than it's worth, and they begin to drop out.

What we have here are two related problems: one practical, the other somewhat existential. The practical problem is whether advertisers can sustain the attention of consumers long enough or well enough to conduct the kinds of coordination games described in previous chapters. Doing so will also require advertisers to maintain a cooperative stance, i.e., to play by the still-evolving rules of these new media rather than to defect for short-term gain. The existential problem implicates the first: can consumers sustain sufficient attention and discernment for meaningful social media participation *at all*, let alone in interactions with marketers? I'll begin with the practical matter of how marketers sustain attention.

## 9.2 The Volunteer's Dilemma and the Tragedy of the Commons

I have outlined in previous chapters how social media marketing represents a shift from a cost-based signaling system to an attention-based signaling system, i.e., the brands that are most engaged with their consumers will, proportionally, gain more attention in social media than those that are not. This system is a boon to companies like Blendtec and Mountain Hardwear, who can now better interact with customers and prospects outside of the costly signaling system of paid advertising. But as the payoff tables make clear, *all* rational marketers want to get better results for less money, and this attention-based social media marketplace looks alluringly like an opportunity to get something for nothing. It is therefore replete with perverse incentives for bad behavior, as examples like Target's Rounders program clearly demonstrate.

This problem of how to get participants to behave themselves in a free and open system with no central governing authority is one that I have taken up in previous chapters in describing the *free rider* scenario, e.g., the person cutting in line in the bakery. You'll recall that the free rider problem crops up in the iterated prisoner's dilemma. To review that scenario: in a one-off prisoner's dilemma, the rational course of action is to defect, since the other player's course of action is unknown. But in the iterated prisoner's dilemma,



if the other player's repeated actions demonstrate a willingness to cooperate, then cooperation is the rational course of action, as the payoffs will be higher.

We can apply this same logic to collective scenarios, in which an individual can gain more in the short term by defecting from the group's collective interest, even though doing so ruins the group's long-term interests and hurts the individual as well. In game theory this problem has been called the Volunteer's Dilemma, which, like its close cousin the Iterated Prisoner's Dilemma, is mainly concerned with how to enforce cooperation wherever it serves participants' mutual long-term interests.

In the Volunteer's Dilemma, the participant must decide whether to make a short-term sacrifice in order to preserve a collective good that they themselves participate in. That sacrifice will cost the individual more in the short-term, and therein lies the dilemma: it is tempting to act as a free rider for short-term gain, especially since making a sacrifice is no guarantee that others will make the same sacrifice. Minding your place in line in the bakery is an apt example, but game theory offers more dramatic ones. The most famous is the "Tragedy of the Commons," first articulated in an article by Garret Hardin in *Science* in 1968.

Hardin invoked a 19<sup>th</sup> century philosophical tract contemplating the problem of overgrazing on lands held in common. When a herdsman decides whether to add another animal to his grazing flock, he reasons that the positive consequences – increased revenue – will accrue to him alone, while the negative consequences – overgrazing – will be shared in common with the other herdsman. He therefore rationally decides to add more animals, and could continue to do so into infinity, even though the resource itself is tragically finite. His fellow herdsman could be expected to follow the same logic and the same tragic course, until the common grazing ground is destroyed. As Hardin eloquently describes this outcome, "Ruin is the destination toward which all men rush, each pursuing his own best interest in a society that believes in the freedom of the commons. Freedom in a commons brings ruin to all" (Hardin 1968).

In contemplating marketers' use of free social network resources like Facebook, YouTube, and Twitter, one might object that these resources are altogether different from shared grazing lands, in that they are as infinite as their owners' willingness to continue adding server capacity. In theory, bad behavior by some marketers will not remove or even reduce the ability of other marketers to get their message out in these forums, since access is, to date, unlimited. But as you have probably surmised, server capacity is not the finite resource at risk of overgrazing; the finite resource is the *consumer's attention*.

### 9.3 Social Media's Capacity Problem

Before considering how marketers will deal with the problem of the diminishing supply of consumer attention, I feel compelled to offer evidence of the problem itself. Let us count the ways: In June 2009, Nielsen reported that Twitter usage had grown 2,000% in one year's time. This anniversary coincided with what some predicted would be a meltdown of all Twitter-related applications, as the number of posts surpassed 2.1 billion, the maximum number of entries that most databases are equipped to handle (based on a 32-bit signed integer). Twitter has since been recoded to handle the additional capacity, and as of this writing, nearly 6 billion "tweets" have been posted.

While the growth in the number of new blogs being created has slowed (presumably because everyone has one now), Technorati reports that there are approximately 900,000 new blog posts added to the blogosphere *every 24 hours*.

Of course, users may reduce their Web usage in order to spend more time on their mobile devices, but they won't escape content saturation there; the number of mobile applications is expected to reach 100,000 by the end of 2010, and 10 million by 2020. The number of text messages sent in the U.S. alone in 2008 is estimated at 95 billion.

With so much social content available, perhaps users can rely on expert guidance – a social media guru – to help sort it all out. But one must first sort out the gurus: there are 5,855 self-proclaimed social media experts on Twitter alone (Ochman 2009). One of the consequences of instantly available information is that one does not need to pay the costly signal of earned experience in social media in order to attain guru status; redistributing information is sometimes enough.

I share these numbers to give some sense of the scale of competition for consumers' attention; that competition, in turn, creates the overgrazing problem. It would be possible, in fact, to plot out this overgrazing mathematically, given the right data set. For instance, Nielsen's most recent report of Web usage shows that the average Facebook user spends a bit over 5 hours per month on the network. Suppose that average user is connected to 5 brands through Facebook, and they spend 1 out of 5 of their Facebook hours reading updates from those brands.

The problem for the brand is that their allotted attention from that user is likely to decrease as the user's Facebook usage increases. As the user increases the number of brands that they connect with, the average attention given to each brand declines. The user may increase their total time on Facebook, but they cannot increase it indefinitely, and time spent on

brands will inevitably lose out to time spent on friends and family as the user increases their number of personal connections. The brand continues to push updates, offering links to discounts and other costly enticements, but the user's attention still declines, as it must, which in turn makes the brand more aggressive. We are in the familiar death spiral of mutual defection; death means the user de-friends the brand, and the connection is lost.

## 9.4 The Risks of Quantification

This problem is greatly exacerbated by the current land rush to *quantify* the success of social media, which inevitably harms the cause of promoting quality. Consider another hypothetical scenario: a marketing manager convinces her VP that the brand could deepen their customer engagement using Twitter. The VP agrees, but demands that the manager establish a goal of 1,000 Twitter followers in the first quarter. Since the process of acquiring Twitter followers is largely organic, i.e., consumers encounter the brand's Twitter feed and decide it is worth following, then the only way for the marketer to achieve this rapid-results goal is to offer an incentive for sign-ups through an online promotion.

The promotion succeeds in garnering the requisite Twitter follower count, but because these users signed up mainly to get the incentive, their engagement level is low. The marketer, in turn, is under pressure to monetize the channel, so she tries to counter low responsiveness with more frequent offers, which causes annoyed users to "un-follow" the brand. The VP concludes that social media has failed the brand, when, in fact, the brand has failed social media.

In this scenario, the Twitter strategy failed because it contained *perverse incentives*, i.e., it induced consumers to sign up for reasons that were at variance with the brand's actual goal for being on Twitter. Consumers signed up for a prize or a discount irrespective of their long-term interest in connecting with the brand, and so the brand traded the *quality* of its followers for a *quantity* of followers that would look attractive on paper. The ultimate cost to the brand was far greater than if it had never been on Twitter at all.

Lest it appear that I am throwing the quantitative baby out with the perverse incentive bathwater, let me acknowledge that it is entirely reasonable for marketers to find ways to quantify the effects of social media, and it is self-evidently reasonable to set a goal of increasing one's quantity of social media followers. In an economic downturn, the ability to quantify results is often what gets a marketing initiative funded, and funding social media initiatives can bring marketers closer to the 4-4 equilibrium.

The danger lies in treating quantification in absolute rather than relative terms, e.g., an arbitrary goal of  $x$  followers rather than gradually building a following over time. The absolutist approach is what leads to the problem of overgrazing; one can privilege quantity or quality, but never both at the same time. Hardin makes this point emphatically, citing game theory founder John von Neumann: “*It is not mathematically possible to maximize for two (or more) variables at the same time*” (emphasis mine).

I will take the risk of belaboring this point because many marketers in the current climate seem determined to do what is mathematically impossible. The rush toward monetization of social media marketing creates perverse incentives for overgrazing; a case in point is the recent emergence of paid tweeting, i.e., the practice of paying an influential social networker to tout a product on Twitter in the context of their everyday tweeting. A *New York Times* article on the practice describes one influential Twitter user with 50,000 followers being paid to endorse personalized M&Ms candy (Stone 2009). This is a textbook example of trying to maximize both variables: the influencer is prized by the marketer for his/her large following, which is the result of their authenticity and credibility. But paid tweeting diminishes their authenticity and credibility, and once the payola is revealed or even suspected, the influencer will lose followers and credibility – both the quantity and the quality of the engagement are diminished.

But as with the overgrazers in Hardin’s metaphor, a single act of defection is easy to justify, because the single act alone will not ruin Twitter’s credibility, and the cost of its diminished credibility is born by the whole community, not by the solo defector. That defector cannot be blamed when the problem has become chronic across the network, and users abandon in droves.

It has been argued that as long as sponsors provide “transparency,” i.e., that they acknowledge when a message is sponsored, that users will not abandon the venue; they can simply choose to ignore or discount the message as they would any other advertising. But then what, actually, is the point? If social connections between marketers and consumers bear the promise of a 4-4 equilibrium, but they are instead used like any other ad medium, what has either player gained? We would, in fact, see the same continuous performance decline in click-based Twitter or Facebook sponsorships as we did in banner advertising, in a cycle of mutual defection. In the *New York Times* piece, an owner of one of the social media sponsorship companies denied that sponsorships would diminish trust in social network connections, saying, ““All we are trying to do is get consumers to become marketers for us.””

Therein lies the problem. Consumers are not marketers. As this study has tried to show, consumers and marketers are mutually dependent adversaries in the marketing game; each cannot do what the other does. When consumers naturally enthuse about brands in social media, their credibility comes from the fact that they are *not marketers*. Turning them into paid marketers destroys their credibility and ruins the system.

A marketer who read the *New York Times* article on Twitter with the slightest degree of hindsight would probably recognize the potential for ruination through overgrazing, because the problem is strongly reminiscent of the over-saturation of banner advertising in the name of quantification in the late 90's, and the resultant consumer defection. But then what might the marketer conclude on the basis of this insight? I suggest there are two plausible responses:

- 1) "If this trend continues, Twitter sponsorships won't be viable for very long. I had better take advantage now, while consumers are still responsive."
- 2) "If this trend continues, Twitter sponsorships won't be viable for very long. I'm going to refrain from contributing to the problem."

What I wish to point out is that *neither of these responses solves the problem*. The first response is obviously an outright defection that will accelerate the decline of the medium as surely as pop-under ads did in banner advertising. The second response is noble, cooperative, and forward-looking in its view that restraint is necessary to preserve the long-term good, but self-restraint will not prevent others from abusing the system. In order to prevent overgrazing of social media marketing, there must be behavioral controls in place that go beyond individual restraint.

## 9.5 Social Traps and Counterreinforcers in Social Media

The question of what kinds of behavioral controls would prevent overgrazing is the subject of "Social Traps," a breakthrough study by John Platt that analyzes the problem of reconciling individual self-interest to the collective good from the perspective of behavioral psychology. Platt builds on the groundwork laid by Garrett Hardin and Thomas Schelling in analyzing the volunteer's dilemma, and he brings in Skinnerian behavioral psychology's emphasis on how positive or negative behaviors are reinforced. In Platt's simple and compelling formulation, social traps occur when a given behavior produces positive results for the individual and negative results for the

group. As long as the individual is only accountable to themselves, the negative behavior is self-reinforcing, resulting in “locked-in behavior,” even though the individual’s long-term interests are imperiled by the behavior. This occurs because social traps typically involve a delay between the short-term gain and the long-term loss; a farmer may get several years of good grazing from the commons before it is destroyed. This problem of “individual goods and collective bads” can’t be solved by the sacrifice of one or two heroes; positive group behavior must somehow be enforced (Platt 1973).

Platt offers several ways out of the social trap, some of which are applicable to the current social media marketing dilemma. The most important of these is the notion of “counterreinforcers.” Since destructive behavior is self-reinforcing in the social trap, counterreinforcers discourage this behavior by offer some negative consequence that the player must evaluate before taking the action. For instance, if the herdsmen on the commons instituted a fee for every grazing animal added beyond a certain quota, then any herdsman acting in his short-term interest would have to weigh this cost against the profitability of adding another animal.

By social media’s very nature, the formation of such formal rules of collective engagement is rare, but counter-reinforcement is not. In any online community, implicit rules of engagement spring up very quickly, and they tend to be rigorously reinforced by its membership. You’ll recall the analysis of the psychological rewards of punishing bad behavior covered in Chapter 4; social media allows participants to go the extra mile in doling out punishment – particularly in the form of verbal castigation – at very little cost. Those who participated in the piling-on of negative reactions to the Motrin video paid very little: an investment of less than 5 minutes in viewing the offending video and responding on Twitter.

In this respect, the risks of a death spiral in social media are very different than they are in, say, banner advertising, where the marketer’s level of control allows them to heap on more and more ad impressions, chasing the elusive click. In social media marketing, consumers exert great control over their level of exposure to brands and can easily dole out punishment in a variety of ways: negative feedback, de-friending, etc. It has often been noted that consumers in social media “vote with their feet,” i.e., they quickly and easily drop social connections with brands that don’t pay off for them.

We could imagine, for instance, celebrities and/or influencers who engage in sponsored tweets being “un-followed” if the implicit rules of the community decreed that sponsored tweets are obnoxious and unwanted. Fearing for the integrity of their own personal brand in the face of this counter-reinforcement, the celebrity would have a strong incentive to drop the sponsorship. The consumer facing an unwanted marketing intrusion into,

say, their movie-watching experience largely stands alone; their negative reaction never surfaces as a counter-reinforcer, and so the bad behavior continues. But an individual negative reaction on Twitter can be mustered into collective outrage in a matter of minutes. Thus unlike a common grazing ground, social media is a common ground with a built-in set of constraints against bad behavior; it's very easy to get kicked of the collective. Marketers have a greater incentive to play by the rules.

## 9.6 Voting with their Feet: Why Quality Matters in Social Media

But marketers' implicit agreement not to be obnoxious is a rather low bar to set for a medium that offers opportunities for deep engagement, and indeed, winning the attention game will require more than avoiding bad behavior. Competing for a consumer's declining attention on a social network demands an emphasis on *quality*. If a consumer's increasingly divided attention span on Twitter means that they will only follow a handful of brands that provide them with valuable content, then there is a built-in incentive for brands to solve the problem with higher-quality Twitter content. If a consumer will only watch 1 out of every 100 brand-sponsored videos on YouTube, then *quality*, as measured in votes and popularity, will be the deciding factor. If the consumer's attention further subdivides to the point at which they will only watch 1 out of every 1000 brand videos, then quality must increase accordingly. The loss of attention raises the cost of the signal that brands rely on to connect through social media, so that only brands able to pay the social cost of deep engagement will succeed. Brands that persist in the "something for nothing" view of social media will simply be squeezed out as available attention declines.

Given the social media acceleration I have outlined, it is axiomatic that consumers will become much more selective about their brand engagements in social media, even as most brands are still getting their sea legs. Does this mean that some brands will simply fail at social media? It does. In fact, given the potential dangers of saturation, it is *necessary* that some brands fail at social media, so that others can succeed. Defining success on the basis of quality rather than quantity is social media marketing's best chance at becoming a mature marketing medium.

Because analysis of social media marketing is still in its Unbridled Enthusiasm stage, in which every new venue is treated as the next big thing, very little attention has been paid to the qualitative factors that determine a brand's success in the space. As I have already noted, the popularity-based

system of costly signaling will allow some brands to succeed at lower cost than others. But how does social media separate the wheat from the chaff?

The best work recognizing the importance of attention as a limited commodity in the social media game is an overlooked study from HPLaboratories' Social Computing Lab, "Crowdsourcing, Attention and Productivity." The study's subject is the relationship between popularity and productivity in YouTube videos, but its findings are broadly applicable to social media marketing. The study's authors raise the question of whether a "tragedy of the commons" is unfolding on YouTube, where over-competition for user attention discourages users from producing new content. In marketing terms, this would impact both brands' willingness to provide content and users' willingness to produce their own brand-related content (such as the Starbucks fan's chronicle of his efforts to visit every Manhattan Starbucks).

The study found that attention was indeed the valued commodity that YouTube uploaders pursued, independent of financial gain. The attention paid by other users, measured in views and comments, very strongly correlated to the likelihood that users would produce more content, and lack of attention had the inverse effect, to the point where users that lacked attention would stop producing videos (Huberman 2009). In other words, the factor that hedges against oversaturation and the tragedy of the commons is the phenomenon of users "voting with their feet"; contributors and brands that don't achieve good quality scores – in the forms of views, comments, and votes – will decide that the costly signal of popularity is too dear. The resultant equilibrium will indeed exclude some brands and force all participating brands to work harder to gain popularity, but the tragedy of the commons can be averted.

## **9.7 Pancake People and the Problem of Information Overload**

There are promising signs, then, that sufficient counter-reinforcement will deter marketer defection in social media marketing, and thus prevent a tragedy of the commons. This leaves us with the more existential problem of information overload in social media. The exponential demands on a user's attention created by the explosion of content raise the question of how users will maintain a sufficient level of discernment to find useful information. This is a critical issue for social media marketing: as consumers become more reliant on peer content to make decisions about brands, their ability to discern differences between high and low quality, truth and fiction, and relevance and irrelevance will greatly shape their purchase decisions. As a Maxwell's Demon tasked with sorting valuable and non-valuable content in



every single Web interaction, will the consumer ultimately be empowered or overwhelmed?

This complex problem boils down to a simple question: are our powers of discernment waxing or waning? As my final analogy in this study, I wish to borrow the playwright Richard Foreman's concept of the "pancake people." In an essay in the cultural studies journal *Edge*, Foreman laments the loss of depth we suffer when knowledge becomes instantly available. We become "spread wide and thin as we connect with that vast network of information accessed by the mere touch of a button." Foreman contrasts this condition with the traditional "cathedral" structure of knowledge, in which individuals acquired information in layers, as they acquired the skills to make sense of that information. In a traditional course of study, analysis of political systems, for instance, would proceed from a basic understanding of forms of government to their ideological permutations. Today, countless political opinions can be accessed at a keystroke, with no prerequisite to understand their ideological basis, and no built-in method for discerning informed opinions from the dangerously uninformed.

Still, respondents to Foreman's lament argued, we may be better off as pancake people, having replaced one form of ignorance – limited access to knowledge – with a less debilitating one: too much knowledge. If we can have our cake and eat it too – that is, if we can develop powers of discernment that allow us to sort information rationally while having instant access to this vast array of information – then we'll be vastly better off. Previous epochs in which the availability of information suddenly surged, e.g., the advent of the printing press, created similar anxieties, but ultimately the greater supply of knowledge had a positive impact on human culture.

While we may indeed evolve to this best of all possible information epochs, we are clearly not there yet. A 2007 study by the British Library on the "information behavior" of Generation Y college students convincingly showed that we have not yet developed the discernment skills necessary to make good use of the glut of available information. The study showed, for instance, an alarming lack of in-depth reading: about 60 percent of e-journal readers consume no more than three pages; the average time spent on e-book and e-journal sites are "four and eight minutes respectively." Users also spent as long searching for information as they did consuming it – a sure sign that Maxwell's Demon is failing. Most alarmingly, the study found that the so-called "Google Generation" struggles with constructing Google searches that accurately reflect the information they're seeking, and then struggles again with discerning relevant search results from irrelevant ones. If the generation raised on Google can't use it properly, what hope is there?

## 9.8 Social Media's Answer to Google's Flattening Effects

The hope may, in fact, lie with social media. Even as user-generated content like blogs, wikis, and forums add significantly to the information glut, they also provide alternative means of accessing information, adding new dimension and perspective to an epistemological scenario that was entirely dominated by Google just half a decade ago. Instead of merely sorting through endless search results as a beset-upon Maxwell's demon, I can locate content through experts on message boards, through primary sources pre-sorted on Wikipedia, through Amazon lists created by reviewers I trust, through Facebook peers, experts on Twitter, and on and on. These resources replace the lonely, dimensionless, pancake-like search with multiple, competing perspectives that demand the use of my critical thinking skills, even as they simplify the information-gathering process.

The media theorist Douglas Rushkoff makes a similar argument in response to Foreman; he argues that the great leap forward in information-gathering is our ability to tap into collective intelligence in a way that shows us multiple perspectives all at once. While the changes in information-gathering have indeed undermined traditional informational authority, we gain the ability to sort through multiple authorities without privileging one over another; thus "our capacity to contend with multiple dimensions is increased." One can see this multi-perspective balance in play in coverage of major news events, such as the 2009 post-election protests in Iran. In similar past events, such as the Tiananmen Square protests in China in 1989, the public's access to information was constrained to major media outlets, whose access was easily constrained by the Chinese government. In the Iran protests, those same constraints on major outlets were in place, but the public had an astonishing degree of real-time access through Twitter feeds, blog posts, and Web videos transmitted from cell phones. The media consumer had the opportunity to weigh these perspectives against official accounts, and the net gain in comprehension of the event is beyond dispute.

The means by which the contemporary media consumer apprehends world events is analogous to the means by which they will apprehend marketing content. Traditional advertising will continue to provide the "official" account, while consumer perspectives on the brand will provide additional dimensions. Advertisers wishing to influence those perspectives will participate in these channels as well. The savvy consumer will not uncritically accept any single one of these perspectives, but will take ownership of a multi-dimensional perspective that represents a net gain in their ability to make smart purchase decisions.

Social media's potential as an antidote to information overload may ultimately lie in its capacity for list-making. The linguist Umberto Eco's affectionate history of list-making, *The Vertigo of Lists*, argues that Western Civilization's penchant for lists has been a critical means of organizing knowledge and seizing control of one's environment. In an interview with *Der Spiegel*, Eco laments the flattening effects of the Google epistemology in a critique reminiscent of Foreman's "pancake people." Eco argues that for young consumers not raised on traditional epistemologies, "Google is a tragedy. Schools ought to teach the high art of how to be discriminating" (Beyer 2009).

Eco is not alone in his view that education has a responsibility to deal with the pancake problem. The concept of "information literacy" has steadily been gaining ground among educators since the advent of the Web; its purpose is to promote methodologies for organizing, synthesizing, and evaluating information, most particularly the unorganized, non-synthesized, and non-evaluated content indexed by the Web. While such skills are obviously important, the movement is controversial, because it attempts to impose standards of competence on one's ability to use an ever-changing medium, and because it is spearheaded by the American Library Association, which, it could be argued, has a vested interest in maintaining libraries' traditional control over information retrieval. The question is whether such a movement is necessary, if in fact the Web can evolve its own capacity for list-making that enhances users' ability to sort and evaluate knowledge.

There is growing evidence that the Web can evolve such a capacity. The online encyclopedia Wikipedia, the content of which is produced and edited by volunteers, was found by the journal *Nature* to have a degree of accuracy comparable to the *Encyclopedia Britannica* (Giles 2005). Wikipedia effectively overcomes much of the indeterminacy of multiple and competing Web-based perspectives by funneling those perspectives into a rigorous peer-based editorial process that demand reliable citations but allows conflicting points of view to co-exist, provided they meet the citation criteria.

The growth of content aggregators like Digg and StumbleUpon are also indicative of the evolution of a list-making capacity. As noted previously, such services add a significant dimension to popularity-based signaling by allowing users to apply peer judgment – even narrowing down their selections to trusted peers only – to their content sorting process. As social media participation increases, i.e., users who were once content consumers only evolve to become content producers, the quality of the sifting will improve, and users will be able to "unflatten" their perspective. In my search for material on game theory, for instance, I could expand my point of view to everything Google produces, or I could confine it to the recommendations of

a handful of recognized, pedigreed experts. In combining the two, I would be accomplishing exactly what Rushkoff advocates: the ability to conditionally hold multiple perspectives at once, with a resulting enrichment of my knowledge of the subject.

Where does all of this leave the age-old, mutually dependent conflict between marketers and consumers? I will end with a set of predictions:

Consumer empowerment through social media is inevitable and permanent, irrespective of the specific channel or technology; the freedom of information that comes with increasingly reliance on peer perspectives cannot be reversed. It is axiomatic that marketing itself will continue to evolve to accommodate these changes, as a simple matter of increased payoffs.

This does not mean, however, that every marketer will participate, or that all who participate will reach the 4-4 equilibrium of mutual collaboration. Traditional advertising will persist for as long as free, sponsored content remains desirable, which is to say, indefinitely. The criteria for successful social media participation by marketers will become more stringent, not less, as demand on consumer attention increases. Persistent consumer backlash against marketers' overstepping in social media will gradually evolve a set of norms for that participation, and many brands will choose not to pay for that costly signal. The evolution of these standards for quality, collaboration, and transparency will allow social media marketing to survive and thrive.

# Bibliography

- 100 Most Valuable Global Brands 2009. (n.d.). Retrieved from <http://www.millwardbrown.com/sites/optimor/Media/Pdfs/en/BrandZ/BrandZ-2009-Report.pdf>
- 171 Starbucks. (2007, July 16). Retrieved December 23, 2009, from <http://www.youtube.com/watch?v=CwYxuV2dVzw>
- 2008 Super Bowl Consumer Intentions and Actions Survey (Rep.). (2008, January 28). Retrieved December 21, 2009, from Retail Advertising Marketing and Advertising website.
- A person like me now most credible spokesperson for companies; trust in employees significantly higher than in CEOs, Edelman Trust Barometer Finds. (2006, January 23). Retrieved December 15, 2009, from <http://www.edelman.com/news/ShowOne.asp?ID=102>
- Angier, N. (2002, July 23). Why We're So Nice: We're Wired to Cooperate – NYTimes.com. Retrieved December 15, 2009, from <http://www.nytimes.com/2002/07/23/health/psychology/23COOP.html>
- Ankeny, J. (2009, July 20). GetJar forecasts 100,000 mobile apps by end of 2009 – FierceMobileContent. Retrieved December 15, 2009, from <http://www.fiercemobilecontent.com/story/getjar-forecasts-100-000-mobile-apps-end-2009/2009-07-20>
- Apple iPhone. (n.d.). Retrieved December 15, 2009, from <http://appleiphone.blogspot.com/>
- Axelrod, R. (2006). *The Evolution of Cooperation*. New York: Perseus Books Group.
- Ayres, C. (2009, July 22). Video: Revenge is best served cold. Retrieved December 20, 2009, from [http://www.timesonline.co.uk/tol/comment/columnists/chris\\_ayres/article6722407.ece](http://www.timesonline.co.uk/tol/comment/columnists/chris_ayres/article6722407.ece)
- Badham, J. (Director). (1983). *War Games* [Motion picture on DVD]. USA: MGM.
- Baeyer, H. C. (1998). *Maxwell's demon why warmth disperses and time passes*. New York: Random House.
- Barash, D. P. (2004). *The Survival Game How Game Theory Explains the Biology of Cooperation and Competition*. New York: Holt Paperbacks.
- Barbaro, M. (2008, January 28). Target Tells a Blogger to Go Away. Retrieved December 14, 2009, from [http://www.nytimes.com/2008/01/28/business/media/28target.html?\\_r=3](http://www.nytimes.com/2008/01/28/business/media/28target.html?_r=3)
- Basulto, D. (2005, January 28). A Little Bit Broken, A Little Bit Perfect. Retrieved December 15, 2009, from <http://www.tcsdaily.com/printArticle.aspx?ID=012805B>
- Beito, L. R. (1999). *Leadership effectiveness in community policing*. Bristol, IN: Wyndham Hall P.
- Belkin, L. (2008, November 17). Moms and Motrin. Retrieved December 20, 2009, from <http://parenting.blogs.nytimes.com/2008/11/17/moms-and-motrin/>
- Best Global Brands List | 2009. (n.d.). Retrieved December 14, 2009, from [http://www.interbrand.com/best\\_global\\_brands.aspx](http://www.interbrand.com/best_global_brands.aspx)

- Beyer, S., & Gorris, L. (2009, November 11). SPIEGEL Interview with Umberto Eco. *Der Spiegel*. Retrieved December 22, 2009, from <http://www.spiegel.de/international/zeitgeist/0,1518,659577,00.html>
- Blendtec's Channel. (n.d.). Retrieved December 22, 2009, from <http://www.youtube.com/user/blendtec?blend=1&ob=4>
- Blog the Debt Away [Editorial]. (2007, March 05). *New York Times*. Retrieved December 20, 2009, from <http://query.nytimes.com/gst/fullpage.html?res=9B07E1DC1431F936A35750C0A9619C8B63>
- Briggs, C. (2009, January). BlendTec Will It Blend? Viral Video Case Study. Retrieved from [http://www.socialens.com/wp-content/uploads/2009/04/20090127\\_case\\_blendtec1.pdf](http://www.socialens.com/wp-content/uploads/2009/04/20090127_case_blendtec1.pdf)
- Carlton, D. (1975). *Dynamics of the arms race*. London: C. Helm.
- Chaney, P. (n.d.). Skittles Social Media Experiment, Fad or Trend | Marketing Profs Daily Fix Blog. Retrieved September 05, 2009, from [http://www.mpdailyfix.com/2009/03/skittles\\_social\\_media\\_experime.html](http://www.mpdailyfix.com/2009/03/skittles_social_media_experime.html)
- Chen, J. (2008, September 24). Blog growth slows; more bloggers are bringing home the bacon. Retrieved December 15, 2009, from <http://arstechnica.com/old/content/2008/09/blog-growth-slows-more-bloggers-are-bringing-home-the-bacon.ars>
- ComScore and Starcom USA Release Updated Natural Born Clickers Study Showing 50 Percent Drop in Number of U.S. Internet Users Who Click on Display Ads – comScore, Inc. (2009, October 01). Retrieved December 15, 2009, from [http://www.comscore.com/Press\\_Events/Press\\_Releases/2009/10/comScore\\_and\\_Starcom\\_USA\\_Release\\_Updated\\_Natural\\_Born\\_Clickers\\_Study\\_Showing\\_50\\_Percent\\_Drop\\_in\\_Number\\_of\\_U.S.\\_Internet\\_Users\\_Who\\_Click\\_on\\_Display\\_Ads](http://www.comscore.com/Press_Events/Press_Releases/2009/10/comScore_and_Starcom_USA_Release_Updated_Natural_Born_Clickers_Study_Showing_50_Percent_Drop_in_Number_of_U.S._Internet_Users_Who_Click_on_Display_Ads)
- Corporate Information – Technology Overview. (n.d.). Retrieved December 22, 2009, from <http://www.google.com/corporate/tech.html>
- Corporate Recruiters Survey 2007* (Rep.). (2008, January 09). Retrieved December 21, 2009, from Graduate Management Admissions Council website: [http://www.gmac.com/NR/rdonlyres/A0E40452-11CC-4838-AC8E-EB6A2D18213B/0/2007CorpRecruit\\_FINAL-web.pdf](http://www.gmac.com/NR/rdonlyres/A0E40452-11CC-4838-AC8E-EB6A2D18213B/0/2007CorpRecruit_FINAL-web.pdf)
- Crosby, J. (2007, December 01). Bloggers seeing red over Target's little secret | StarTribune.com. Retrieved December 14, 2009, from <http://www.startribune.com/business/11987331.html>
- D'Angelo, F. (2009, October 26). Happy Birthday, Digital Advertising! – Advertising Age – DigitalNext. Retrieved December 16, 2009, from [http://adage.com/digitalnext/post.php?article\\_id=139964](http://adage.com/digitalnext/post.php?article_id=139964)
- Davis, M. D. (1997). *Game theory a nontechnical introduction*. Mineola, N.Y: Dover Publications.
- Defending Whales: Mr Splashy Pants goes viral! (2007, November 28). Retrieved December 21, 2009, from [http://weblog.greenpeace.org/whales/2007/11/mr\\_splashy\\_pants\\_goes\\_viral.html](http://weblog.greenpeace.org/whales/2007/11/mr_splashy_pants_goes_viral.html)
- Delaney, A. (2009, September 30). Bank Of America: Another Customer Gets Bank's Attention Via YouTube. Retrieved December 20, 2009, from [http://www.huffingtonpost.com/2009/09/30/bank-of-america-another-c\\_n\\_304630.html](http://www.huffingtonpost.com/2009/09/30/bank-of-america-another-c_n_304630.html)
- Diaz, S. (2008, August 06). Internet brand-jacking: What can be learned from Exxon Mobil? | Between the Lines | ZDNet.com. Retrieved December 23, 2009, from <http://blogs.zdnet.com/BTL/?p=9602>

- Dixit, A. K., & Nalebuff, B. J. (1993). *Thinking Strategically The Competitive Edge in Business, Politics, and Everyday Life*. Boston: W. W. Norton & Company.
- Dixit, A. K., & Nalebuff, B. J. (2008). *The Art of Strategy A Game Theorist's Guide to Success in Business and Life*. New York: W. W. Norton.
- Do the World a Flavor. (n.d.). Retrieved December 23, 2009, from <http://www.benjerry.com/activism/inside-the-pint/do-the-world-a-flavor/>
- Do You Know Who's On Twitter? (2009, August 28). Retrieved December 15, 2009, from <http://www.emarketer.com/Articles/Print.aspx?1007250>
- E Quervain, D. (1994). The Neural Basis of Altruistic Punishment. *Science Magazine*, 27(305), 1254-1258.
- Eco, U. (2009). *The Infinity of Lists*. Rizzoli.
- Ehrenberg, A. (1974). Repetitive Advertising and the Consumer. *Journal of Advertising Research*, April, 25-34.
- Elliott, S. (2000, February 02). THE MEDIA BUSINESS: ADVERTISING – THE MEDIA BUSINESS: ADVERTISING; Expensive Super Bowl spots seem to work for many dot-coms as traffic to their sites grows. – NYTimes.com. Retrieved December 15, 2009, from <http://www.nytimes.com/2000/02/02/business/media-business-advertising-expensive-super-bowl-spots-seem-work-for-many-dot.html>
- E-Marketing Report 2000* (Rep.). (2008, October 06). Retrieved December 16, 2009, from Morgan Stanley Dean Witter website: <http://www.docstoc.com/docs/1744758/emarketing2000ReportMorganStanley>
- Evans, L. (2009, March 02). Skittles & Social Media – Obviously a Company that Doesn't Get It from. Retrieved December 14, 2009, from [http://www.searchmarketinggurus.com/search\\_marketing\\_gurus/2009/03/skittles-social-media-obviously-a-company-that-doesnt-get-it.html](http://www.searchmarketinggurus.com/search_marketing_gurus/2009/03/skittles-social-media-obviously-a-company-that-doesnt-get-it.html)
- Excellent series about Noka chocolate at Dallas Food -. (2006, December 26). Retrieved December 14, 2009, from [http://blog.kitchenmage.com/2006/12/excellent\\_serie.html](http://blog.kitchenmage.com/2006/12/excellent_serie.html)
- Fain, D., & Pedersen, J. (2005). Sponsored Search: A Brief History. *Bulletin of the American Society for Information Science and Technology*.
- Feature: State of the Blogosphere 2009 – Technorati. (n.d.). Retrieved December 14, 2009, from <http://technorati.com/blogging/feature/state-of-the-blogosphere-2009/>
- Feltovich, N., Harbaugh, R., & To, T. (2002). Too cool for school? Signalling and countersignalling. *RAND Journal of Economics*, 33(4), 630-639.
- Frontline: the persuaders: interviews: clotaire rapaille |. (n.d.). Retrieved December 15, 2009, from <http://www.pbs.org/wgbh/pages/frontline/shows/persuaders/interviews/rapaille.html>
- Frucci, A. (2008, June 26). Nike Finally Releasing Back to the Future Part II McFly Sneakers, Sort Of – Nike – Gizmodo. Retrieved December 23, 2009, from <http://gizmodo.com/5020072/nike-finally-releasing-back-to-the-future-part-ii-mcfly-sneakers-sort-of>
- Gaffney, J. (2009, August 17). Bringing Coca-Cola to the Party. Retrieved December 15, 2009, from [http://www.digidaydaily.com/stories/bringing\\_coca\\_cola\\_to\\_the\\_party/](http://www.digidaydaily.com/stories/bringing_coca_cola_to_the_party/)
- Giles, J. (2005). Internet encyclopaedias go head to head. *Nature*, 438, 900-901. Retrieved December 22, 2009, from <http://www.nature.com/nature/journal/v438/n7070/full/438900a.html>
- Graham, J. (1999, July 26). Escaping the Cult of Click-Through – ClickZ. Retrieved December 16, 2009, from <http://www.clickz.com/816321>

- Groundswell: MyStarbucksIdea.com. (n.d.). Retrieved December 14, 2009, from <http://www.forrester.com/Groundswell/embracing/mystarbucksidea.html>
- Hardin, G. (1968). The Tragedy of the Commons. *Science*, 162, 1243-1248.
- Harmon, A. (2003, August 25). Finding Comfort in Strangers With an Online Diet Journal – NYTimes.com. Retrieved December 21, 2009, from <http://www.nytimes.com/2003/08/25/technology/25DIET.html>
- Heath, T. (2006, November 24). In Sales, Nike Is a Shoe-In – washingtonpost.com. Retrieved December 15, 2009, from <http://www.washingtonpost.com/wp-dyn/content/article/2006/11/23/AR2006112301066.html>
- Heinz Names Winner of Consumer Generated Ad Contest. (2008, April 30). Retrieved December 15, 2009, from [http://promomagazine.com/contests/archive/heinz\\_names\\_winner\\_consumer\\_generated\\_ad\\_contest\\_0430/](http://promomagazine.com/contests/archive/heinz_names_winner_consumer_generated_ad_contest_0430/)
- Hespos, T. (2009, August 11). Tracking Non-Click Conversions – ClickZ. Retrieved December 16, 2009, from <http://www.clickz.com/818361>
- Hespos, T. (2009, July 29). Between A Rock And A Hard Place – ClickZ. Retrieved December 16, 2009, from <http://www.clickz.com/816431>
- Huberman, B. A., Romero, D. M., & Wu, F. (2009). Crowdsourcing, Attention, and Productivity. *Journal of Information Science*, 35(6), 758-769.
- Information Behaviour of the Researcher of the Future* (Rep.). (2007, October 18). Retrieved December 22, 2009, from British Library/JISC website: <http://www.jisc.ac.uk/media/documents/programmes/reppres/ggworkpackageii.pdf>
- IProspect Search Engine User Behavior Study* (Rep.). (2006). Retrieved December 21, 2009, from IProspect website: [http://www.iprospect.com/premiumPDFs/WhitePaper\\_2006\\_SearchEngineUserBehavior.pdf](http://www.iprospect.com/premiumPDFs/WhitePaper_2006_SearchEngineUserBehavior.pdf)
- Jarvis, J. (2007, October 17). Dell Learns to Listen. Retrieved December 21, 2009, from [http://www.businessweek.com/bwdaily/dnflash/content/oct2007/db20071017\\_277576.htm?chan=top+news\\_top+news+index\\_top+story](http://www.businessweek.com/bwdaily/dnflash/content/oct2007/db20071017_277576.htm?chan=top+news_top+news+index_top+story)
- Johannes, A. (2008, June 20). McDonald's Celebrates Big Mac Anniversary with Song Contest. Retrieved December 15, 2009, from [http://promomagazine.com/contests/archive/mcdonalds\\_celebrates\\_big\\_mac\\_0620/](http://promomagazine.com/contests/archive/mcdonalds_celebrates_big_mac_0620/)
- Kapferer, J. (2008). *The New Strategic Brand Management Creating and Sustaining Brand Equity Long Term*. London: Kogan Page.
- Keeney, D. (2006, December 30). Ethical Considerations In Posting Comments to Blogs – 30-Dec-06: Public Relations article: contact Dan Keeney. Retrieved December 14, 2009, from <http://www.dpkr.com/en/art/index.asp?222>
- Keizer, G. (2007, January 10). Spam Sets Record, Accounts For 94% Of E-mail – Spam – InformationWeek. Retrieved December 17, 2009, from <http://www.informationweek.com/news/internet/showArticle.jhtml?articleID=196802782>
- Kelley, S., & Turley, A. (2002). *Analyzing the Relationship Between Ad Content and the Effectiveness of Super Bowl Commercials*. Lecture presented at 2002 AMA Faculty Consortium on Sports Marketing, Lexington, Kentucky.
- Kohler, E. (2007, November 30). More on Target Rounders | The Deets. Retrieved December 14, 2009, from <http://www.thedeets.com/2007/11/30/more-on-target-rounders/>
- Li, C., & Bernoff, J. (2008). *Groundswell: Winning in a World Transformed by Social Technologies*. New York: Harvard Business School P.
- Link schemes – Webmasters/Site owners Help. (n.d.). Retrieved December 22, 2009, from <http://www.google.com/support/webmasters/bin/answer.py?hl=en&answer=66356>



- Lizza, R. (2008, May 05). Campaign Journal: Bill Vs. Barack . Retrieved December 14, 2009, from [http://www.newyorker.com/talk/2008/05/05/080505ta\\_talk\\_lizza](http://www.newyorker.com/talk/2008/05/05/080505ta_talk_lizza)
- MarkMonitor. (2009, March 09). *MarkMonitor Year-in-Review Report Finds Online Abuse of Major Brands Was a Growth Industry for Fraudsters in 2008* [Press release]. Retrieved from <http://markmonitor.com/pressreleases/2009/pr090309.php>
- Martin, R. (1998, March 09). Foodmaker settles E. coli contamination lawsuit. *Nation's Restaurant News*. Retrieved December 22, 2009, from [http://findarticles.com/p/articles/mi\\_m3190/is\\_n10\\_v32/ai\\_20379708/](http://findarticles.com/p/articles/mi_m3190/is_n10_v32/ai_20379708/)
- Masel, J. (2007). A Bayesian model of quasi-magical thinking can explain observed co-operation in the public good game. *Journal of Economic Behavior*, 64, 216-231.
- McCall, M. (2007, November 26). The McFly 2015 Project. Retrieved December 14, 2009, from <http://beta.picks.yahoo.com/picks/2578/the-mcfly-2015-project>
- Miller, J. D. (2003). *Game Theory at Work How to Use Game Theory to Outthink and Outmaneuver Your Competition*. New York: McGraw-Hill.
- Morrissey, B. (2009, October 12). Success Factors for Brands on Facebook. Retrieved December 20, 2009, from [http://www.adweek.com/aw/content\\_display/news/digital/e3i-30e7feb16ddb0207759efe2444962050](http://www.adweek.com/aw/content_display/news/digital/e3i-30e7feb16ddb0207759efe2444962050)
- Most Expensive Chocolates. (n.d.). Retrieved December 14, 2009, from [http://www.forbes.com/2006/02/07/cx\\_sy\\_0208featslide\\_2.html?thisSpeed=35000](http://www.forbes.com/2006/02/07/cx_sy_0208featslide_2.html?thisSpeed=35000)
- Motrin Ad Makes Moms Mad. (2008, November 16). Retrieved December 21, 2009, from <http://www.youtube.com/watch?v=LhR-y1N6R8Q>
- Nash, J. (2002). *Essential John Nash* (H. W. Kuhn, Ed.). Princeton, N.J: Princeton UP.
- Needleman, R. (2009, August 15). Starbucks: Stay as long as you want | Rafe's Radar – CNET News. Retrieved December 13, 2009, from [http://news.cnet.com/8301-19882\\_3-10310495-250.html](http://news.cnet.com/8301-19882_3-10310495-250.html)
- Neff, J. (2009, October 06). Apple, Microsoft Are Megabrand Kingpins in Social Radar Index. Retrieved December 14, 2009, from [http://adage.com/abstract.php?article\\_id=139386](http://adage.com/abstract.php?article_id=139386)
- Nelson, E. (2009, June 05). *Dell's Journey in Conversation and Community*. Lecture presented at Social Commerce Summit. *Docstoc*. Retrieved from <http://www.docstoc.com/docs/5817646/Chapter-Two-How-Dell-is-Integrating-Social-Media-into-Our-Business>
- Neumann, J. V., & Morgenstern, O. (2007). *Theory of Games and Economic Behavior (Commemorative Edition) (Princeton Classic Editions)*. New York: Princeton UP.
- Nicole, K. (2007, December 12). Mr. Splashy Pants and the Tale of a Hijacked PR Campaign. Retrieved December 21, 2009, from <http://mashable.com/2007/12/12/mr-splashy-pants-greenpeace/>
- Nielsen, J. (2007, August 20). Banner Blindness: Old and New Findings (Jakob Nielsen's Alertbox). Retrieved December 15, 2009, from <http://www.useit.com/alertbox/banner-blindness.html>
- NOKA Chocolate – How Luxury Chocolate is Made. (2007, April 09). Retrieved December 23, 2009, from <http://www.youtube.com/watch?v=4JeiB8J2Buw>
- Noyes, K. (2008, June 25). E-Commerce News: E-Marketing: Racy JCPenney Ad Stripped From YouTube. Retrieved December 15, 2009, from <http://www.ecommercetimes.com/story/Racy-JCPenney-Ad-Stripped-From-YouTube-63570.html?welcome=1214575663&wlc=1252351087>
- Ochman, B. (2009, June 29). Updated List of Social Media Gurus on Twitter: a Public Service for CMOs, from What's Next Blog. Retrieved December 15, 2009, from [http://www.digidaily.com/stories/updated\\_list\\_of\\_social\\_media\\_gurus\\_on\\_twitter\\_a\\_public\\_service\\_for\\_cmos\\_from\\_what\\_s\\_next\\_blog/](http://www.digidaily.com/stories/updated_list_of_social_media_gurus_on_twitter_a_public_service_for_cmos_from_what_s_next_blog/)

- Odell, P. (2009, March 18). Dunkin? Donuts Returns to its Roots? Doughnuts? in \$10 Million Campaign. Retrieved December 15, 2009, from <http://promomagazine.com/contests/dunkindonutscampaign/>
- Ostrom, E., Gardner, R., & Walker, J. (1994). *Rules, Games, and Common-Pool Resources*. Ann Arbor: University of Michigan P.
- Owyang, J. (2009, January 28). Community Managers Must Deliver ROI: Commandments For Surviving a Recession. Retrieved December 21, 2009, from <http://www.web-strategist.com/blog/2009/01/28/community-managers-must-deliver-roi-tips-for-surviving-a-recession/>
- Pasick, A. (2005, July 12). Ninety pct of DVR users skip TV ads. Retrieved December 17, 2009, from [http://www.redorbit.com/news/entertainment/172001/ninety\\_pct\\_of\\_dvr\\_users\\_skip\\_tv\\_ads\\_new/index.html](http://www.redorbit.com/news/entertainment/172001/ninety_pct_of_dvr_users_skip_tv_ads_new/index.html)
- Platt, J. (1973). Social Traps. *American Psychologist*, 28, 641-651.
- Poundstone, W. (1992). *Prisoner's dilemma*. New York: Doubleday.
- RealtyTrac. (2009, September 10). *Foreclosure activity remains near record level in August* [Press release]. Retrieved December 20, 2009, from <http://www.realtytrac.com/content-management/pressrelease.aspx?channelid=9&itemid=7381>
- Reynolds, C. (2009, July 07). Smashed guitar, YouTube song; United is listening now – Travel – LATimes.com. Retrieved December 15, 2009, from <http://travel.latimes.com/daily-deal-blog/index.php/smashed-guitar-youtu-4850/>
- Robinson-Jacobs, K. (2007, January 25). Plano chocolatier feels blogger's bite. Retrieved December 23, 2009, from <http://www.dallasnews.com/sharedcontent/dws/bus/stories/012507dnbusnoka.1caa3f0.html>
- Schelling, T. (1982, March 19). *Ethics, Law, and the Exercise of Self-Command*. Lecture presented at Tanner Lectures on Human Values in University of Michigan. Retrieved December 15, 2009, from <http://www.tannerlectures.utah.edu/lectures/documents/schelling83.pdf>
- Schelling, T. (2006). *Micromotives and Macrobehavior*. New York: W. W. Norton.
- Schelling, T. C. (2006). *The Strategy of Conflict*. New York: Harvard UP.
- Schultz, H. (2008, January 07). Howard Schultz's Memo. Retrieved December 21, 2009, from [http://www.businessweek.com/bwdaily/dnflash/content/jan2008/db2008017\\_860745.htm](http://www.businessweek.com/bwdaily/dnflash/content/jan2008/db2008017_860745.htm)
- Skyrms, B. (2003). *The Stag Hunt and the Evolution of Social Structure*. New York: Cambridge UP.
- Skyrms, B. (2004). *Stag hunt and the evolution of social structure*. Cambridge: Cambridge UP.
- Social Media Marketing and PR Benchmark Survey 2008* (Rep.). (2008). Retrieved [marketingsherpa.com](http://marketingsherpa.com)
- Social Networking Sites Account for More than 20 Percent of All U.S. Online Display Ad Impressions, According to comScore Ad Metrix – comScore, Inc. (2009, September 01). Retrieved December 15, 2009, from [http://comscore.com/Press\\_Events/Press\\_Releases/2009/9/Social\\_Networking\\_Sites\\_Account\\_for\\_More\\_than\\_20\\_Percent\\_of\\_All\\_U.S.\\_Online\\_Display\\_Ad\\_Impressions\\_According\\_to\\_comScore\\_Ad\\_Metrix](http://comscore.com/Press_Events/Press_Releases/2009/9/Social_Networking_Sites_Account_for_More_than_20_Percent_of_All_U.S._Online_Display_Ad_Impressions_According_to_comScore_Ad_Metrix)
- Sosa, J. (2008, July 07). Nike messes with my head | Hutchinson Leader. Retrieved December 15, 2009, from <http://www.hutchinsonleader.com/community/jorge-sosa/nike-messes-my-head>
- Spence, M. (1973). Job Market Signaling. *Quarterly Journal of Economics*, 87(3), 355-374.

- Starbucks Ad. (2006, August 23). Retrieved December 23, 2009, from [http://www.youtube.com/watch?v=VnbT7qt6RF4&feature=player\\_embedded](http://www.youtube.com/watch?v=VnbT7qt6RF4&feature=player_embedded)
- Steel, E. (2009, March 03). What Are You Doing? (#Skittles). Retrieved December 14, 2009, from <http://blogs.wsj.com/digits/2009/03/03/what-are-you-doing-skittles/>
- Stone, B. (2009, November 21). A Friend's Tweet Could Be an Ad. *New York Times*. Retrieved December 22, 2009, from [http://www.nytimes.com/2009/11/22/business/22ping.html?\\_r=2](http://www.nytimes.com/2009/11/22/business/22ping.html?_r=2)
- Surowiecki, J. (2000, September 04). How Kozmo is Getting Killed by its Customers. *The New Yorker*, 34.
- Starantino, Q. (Director). (1992). *Reservoir Dogs* [Motion picture]. USA: Lions Gate.
- Tate, R. (2009, October 27). On Firing Day, Busy Wired Editor Had Other Places To Be – Chris Anderson – Gawker. Retrieved December 15, 2009, from <http://gawker.com/5391319/on-firing-day-busy-wired-editor-had-other-places-to-be>
- Taylor, W. (2007, December 18). Permission Marketing. Retrieved December 17, 2009, from <http://www.fastcompany.com/magazine/14/permission.html>
- The Pancake People, or, The Gods Are Pounding My Head. (2005, March 08). Retrieved December 23, 2009, from [http://www.edge.org/3rd\\_culture/foreman05/foreman05\\_index.html](http://www.edge.org/3rd_culture/foreman05/foreman05_index.html)
- The Top Brands in Social Media Report 2008. (n.d.). Retrieved December 22, 2009, from <http://www.immediatefuture.co.uk/the-top-brands-in-social-media-report-2008/>
- Time Spent on Facebook up 700%, but MySpace Still Tops for Video | Nielsen Wire. (2009, June 02). Retrieved December 15, 2009, from [http://blog.nielsen.com/nielsenwire/online\\_mobile/time-spent-on-facebook-up-700-but-myspace-still-tops-for-video/](http://blog.nielsen.com/nielsenwire/online_mobile/time-spent-on-facebook-up-700-but-myspace-still-tops-for-video/)
- Top U.S. Web Brands and Parent Companies for September 2009 | Nielsen Wire. (2009, October 14). Retrieved December 15, 2009, from [http://blog.nielsen.com/nielsenwire/online\\_mobile/top-u-s-web-brands-and-parent-companies-for-september-2009/](http://blog.nielsen.com/nielsenwire/online_mobile/top-u-s-web-brands-and-parent-companies-for-september-2009/)
- Troaca, F. (2008, August 29). 95.4 Billion Text Messages Sent by US Mobile Users in 2008 – New record for the US – Softpedia. Retrieved December 15, 2009, from <http://news.softpedia.com/news/95-4-Billion-Text-Messages-Sent-by-US-Mobile-Users-in-2008-92684.shtml>
- Tucker, A. W. (1983). The Mathematics of Tucker: A Sampler. *The Two-Year College Mathematics Journal*, 14(3), 228-232.
- United Breaks Guitars. (2009, July 06). Retrieved December 20, 2009, from <http://www.youtube.com/watch?v=5YGc4zOqozo>
- US Online Ad Spend Turns the Corner. (2009, December 11). Retrieved December 20, 2009, from [http://www.emarketer.com/Article.aspx?R=1007415&Ntt=paid+search&No=0&xsrc=article\\_head\\_siteearchx&N=0&Ntk=basic](http://www.emarketer.com/Article.aspx?R=1007415&Ntt=paid+search&No=0&xsrc=article_head_siteearchx&N=0&Ntk=basic)
- USA, Federal Trade Commission, Do Not Call Registry. (2008, July). *Annual Report to Congress for FY 2007*. Retrieved December 17, 2009, from <http://www.ftc.gov/os/2008/07/P034305FY0dncreport.pdf>
- Vargas, J. A. (2008, November 20). Obama Raised Half a Billion Online | 44 | washingtonpost.com. Retrieved December 14, 2009, from [http://voices.washingtonpost.com/44/2008/11/20/obama\\_raised\\_half\\_a\\_billion\\_on.html](http://voices.washingtonpost.com/44/2008/11/20/obama_raised_half_a_billion_on.html)
- Vine, R. (2008, September 01). Mad Men obsessives Twitter away | Technology | The Guardian. Retrieved December 15, 2009, from <http://www.guardian.co.uk/technology/2008/sep/01/internet.blogging>

- Vranica, S. (2009, December 17). Pepsi's Drinks Will Sit Out the Super Bowl. Retrieved December 22, 2009, from <http://online.wsj.com/article/SB10001424052748703581204574600322164130250.html>
- Wahlberg, D. (2002, July 18). Why cooperate? It's a pleasure, says Emory study. *Atlanta Journal-Constitution*.
- Wallace, D. F. (2000). *Up Simba!: 7 Days on the Trail of an Anticandidate*. New York: Warner Books.
- Walsh, J. (2007, April 04). Salon's new letters registration policy – Salon.com. Retrieved December 14, 2009, from [http://www.salon.com/opinion/walsh/salon/2007/04/04/letters\\_registration/index.html](http://www.salon.com/opinion/walsh/salon/2007/04/04/letters_registration/index.html)
- Warren, C. (2009, August 27). Audi Taps its Facebook Fans to Help Design Car of the Future. Retrieved December 15, 2009, from <http://mashable.com/2009/08/27/audi-facebook/WebConnect>
- Site Price Index. (n.d.). Retrieved December 16, 2009, from <http://www.sitepriceindex.com/Avgprice.html>
- Weston, N. (2006, December 19). NOKA chocolate exposed! -. Retrieved December 14, 2009, from <http://www.slashfood.com/2006/12/19/noka-chocolate-exposed/>
- What's Noka Worth? (2006, December 09). Retrieved December 14, 2009, from <http://www.dallasfood.org/modules.php?name=News&file=article&sid=78>
- White, R. (1999). What can advertising really do for brands? *International Journal of Advertising*, 18(1).
- Wloszczyna, S. (2005, January 13). USATODAY.com – Steve Zissou's cool shoes prove tough to fill. Retrieved December 15, 2009, from [http://www.usatoday.com/life/movies/news/2005-01-13-zissou-shoes\\_x.htm](http://www.usatoday.com/life/movies/news/2005-01-13-zissou-shoes_x.htm)
- Wu, J. (2001, November 01). Anatomy of a Dot-Com. *Supply Chain Management Review*.
- Yu, F. (2000, October 04). The Unbearable Lightness of Ad Revenue. *Business World*