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 21st 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

DOI 10.1007/978-3-642-13299-5_1,

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 customers gripe at you?

Table 1: GM Fastlane payoff matrix

	Customer: *Give positive feedback*	*Customer:* *Give negative feedback*
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 conversations. 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.