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

E. Anderson, *Social Media Marketing: Game Theory and the Emergence of Collaboration*,
DOI 10.1007/978-3-642-13299-5_2,

"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 Kruschev 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.

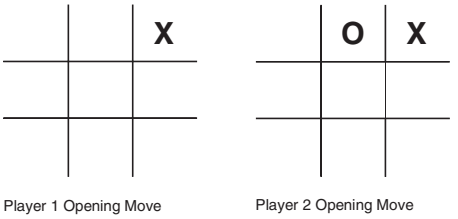

Figure 1: Tic-Tac-Toe

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

	Reagan: Heads	*Reagan: Tails*
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.

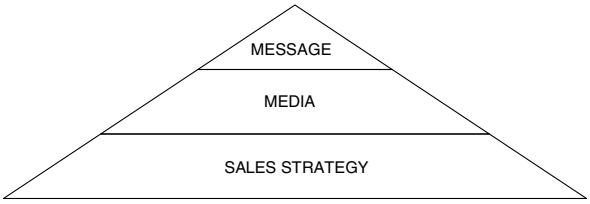

Figure 2: Hierarchy of marketing factors

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

	Me: Shop STP	*Me: Don't shop STP*
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.