

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.