

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

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

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

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

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

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

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

7.1 The Theory of Costly Signaling

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

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

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

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

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

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

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

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

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

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

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

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

7.2 Traditional Advertising and Costly Signaling

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

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

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

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

7.3 The Erosion of Costly Signaling in Super Bowl Advertising

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

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

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

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

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

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

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

7.4 How Web Transparency Disrupts Costly Signaling

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

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

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

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

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

7.5 The Evolution of Costly Signaling on the Web

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

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

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

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

7.6 Google Changes the Costly Signaling Currency

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7.7 Paid Search and Popularity-Based Signaling

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

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

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

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

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

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

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

7.8 Noise in the Signaling System

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

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

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

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

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

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

7.9 Popularity-Based Signaling in Social Media

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

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

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

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

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

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

7.10 The Disruptive Effects of Popularity-Based Signaling

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7.11 The Perils of Negativity in Popularity-Based Signaling

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

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

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

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

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

7.12 Popularity-based Signaling in Presidential Politics

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

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

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

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

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

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

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

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

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

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

7.13 Noka and the Disruption of Costly Signaling

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7.14 Sock Puppetry as Noise in the System

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

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

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

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

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

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

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

7.15 Target Learns the Rules of Popularity-Based Signaling

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7.16 Target and the Problem of Capacity

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

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

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

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

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

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

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

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

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

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

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

7.17 Content Aggregators and the Evolution of Costly Signaling

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

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

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

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

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

7.18 The Three Rules of Popularity-Based Costly Signaling

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

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

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

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

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

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

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

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

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

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

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

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

7.19 Paying the Costly Signal to Go Viral

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

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

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

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

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

7.20 A Popularity-Based Success Story: Blendtec

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

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

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

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

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

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

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

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