

Leading the Herd Astray: An Experimental Study
of Self-fulfilling Prophecies in an Artificial Cultural MarketMATTHEW J. SALGANIK
Princeton UniversityDUNCAN J. WATTS
Yahoo! Research and Columbia University

Individuals influence each others' decisions about cultural products such as songs, books, and movies; but to what extent can the perception of success become a "self-fulfilling prophecy"? We have explored this question experimentally by artificially inverting the true popularity of songs in an online "music market," in which 12,207 participants listened to and downloaded songs by unknown bands. We found that most songs experienced self-fulfilling prophecies, in which perceived—but initially false—popularity became real over time. We also found, however, that the inversion was not self-fulfilling for the market as a whole, in part because the very best songs recovered their popularity in the long run. Moreover, the distortion of market information reduced the correlation between appeal and popularity, and led to fewer overall downloads. These results, although partial and speculative, suggest a new approach to the study of cultural markets, and indicate the potential of web-based experiments to explore the social psychological origin of other macrosociological phenomena.

In markets for books, music, movies, and other cultural products, individuals often have information about the popularity of the products from a variety of sources: friends, bestseller lists, box office receipts, and, increasingly, online forums. Consistent with the psychological literature on social influence (Sherif 1937; Asch 1952; Cialdini and Goldstein 2004), recent work has found that consumers' decisions about cultural products can be influenced by this information (Hanson and Putler 1996; Senecal and Nantel 2004; Huang and Chen 2006; Salganik, Dodds, and

Watts 2006), in part because people use the popularity of products as a signal of quality—a phenomenon that is sometimes referred to as "social" or "observational learning" (Hedström 1998)—and in part because people may benefit from coordinating their choices such as listening to, reading, and watching the same things as others (Adler 1985).

The influence that individuals have over each other's behavior, moreover, can have important consequences for the behavior of cultural markets. In a recent study of an artificial music market, for example, when participants were aware of the previous decisions of others, the popularity of songs was determined in part by a "cumulative advantage" process where early success lead to future success (Salganik et al. 2006). These observed dynamics naturally raise the question of whether *perceived* success alone is sufficient to generate continued success. Can success in cultural markets, in other words, arise solely as a "self-fulfilling prophecy"—a phrase coined by Robert Merton to mean "a *false* definition of the situation evoking a new behavior which makes the original false conception come *true*" [emphasis in original] (Merton

We thank Peter Hausel and Peter S. Dodds for help in designing and developing the MusicLab website, and Peter S. Dodds and members of the Collective Dynamics Group for many helpful conversations. We also thank Hana Shepherd and Amir Goldberg for helpful comments on this manuscript. This research was supported in part by an NSF graduate research fellowship (to MJS), NSF grants SES-0094162 and SES-0339023, the Institute for Social and Economic Research and Policy at Columbia University, and the James S. McDonnell Foundation. The data analyzed in this paper are publicly available from the data archive of the Office of Population Research at Princeton University (<http://opr.princeton.edu/archive/>). Please direct correspondence to Matthew J. Salganik, Department of Sociology, 145 Wallace Hall, Princeton University, Princeton NJ, 08544; mjs3@princeton.edu.

1948).¹ Merton illustrated his idea with a hypothetical bank run—describing how depositors' initially false belief that a bank was going to fail might lead them to withdraw their money causing an actual failure—but the concept was intended to be extremely broad. The bulk of Merton's seminal paper, for example, was devoted to analyzing race-based differences in academic and professional achievement as a self-fulfilling prophecy.

Since Merton, social scientists have considered the possibility of self-fulfilling prophecies in a wide variety of domains, and these studies can be classified according to the unit of analysis at which the self-fulfilling dynamics are thought to take place: individual, dyadic, and collective (Biggs forthcoming). At the level of the individual, for example, a substantial body of work has explored the placebo effect on health outcomes; that is, whether, and under what conditions, individuals experience improved health outcomes that are attributable to receiving a specific substance, but that are not due to the inherent powers of the substance itself (Stewart-Williams and Podd 2004). Moving up to the level of the dyad, researchers have explored the effects of teacher expectations on student outcomes (Rosenthal and Jacobson 1968; Jussim and Harber 2005) and physician prognosis on patient health (Christakis 1999). Finally, at the collective level (i.e., situations involving more than two individuals), sociologists have explored the “performativity” of economic theory (Callon 1998), including whether the predictions of economic models lead people to change their behavior in such a way as to make the original prediction come true (MacKenzie and Millo 2003). A separate body of work by economists, meanwhile, has explored self-fulfilling prophecies with regard to financial panics (Calomiris and Mason 1997), investment bubbles (Garber 1989), and even business cycles (Farmer 1999).

¹ Merton's work on self-fulfilling prophecies was heavily influenced by Thomas and Thomas (1928) who wrote what Merton later called the Thomas Theorem: “if men define situations as real, they are real in their consequences.” For a complete review of the intellectual history, see Merton (1995).

At face value, these studies suggest that over a wide range of scales and domains, the belief in a particular outcome may indeed cause that outcome to be realized, even if the belief itself was initially unfounded or even false. Skeptics, however, remain unconvinced that expectations or perceptions should be treated as fundamental elements of the social environment, on par with individual preferences (Rogerson 1995), and much of the empirical evidence that might persuade them is mixed or ambiguous (Biggs forthcoming). One reason that skeptics remain is that with observational data alone it is virtually impossible to prove that any given real-world event was caused by the self-fulfillment of false beliefs. This difficulty arises because the strongest support for such a claim would require comparing outcomes in the presence or absence of false beliefs, but in almost all cases only one of these outcomes is observed (Holland 1986; Sobel 1996; Winship and Morgan 1999).

Given these limits of observational data, it is no surprise that our best understanding of self-fulfilling prophecies in cultural markets comes from experimental and quasi-experimental methods. For example, by exploiting errors in the construction of the *New York Times* bestseller list, Sorensen (2007) found that books mistakenly omitted from the list had fewer subsequent sales than a matched set of books that correctly appeared on the list. Hanson and Putler (1996), moreover, performed a field experiment in which they directly intervened in a real market by repeatedly downloading randomly chosen software programs to inflate their perceived popularity. Software that received the artificial downloads went on to earn substantially more real downloads than a matched set of software. While these studies offer insight, they may misrepresent the possibility for self-fulfilling prophecies because the manipulations employed were rather modest and therefore only loosely decoupled perceived success from actual success. Further, these studies lacked a measure of the preexisting preferences of participants, and therefore cannot shed any light on how the “quality” of the products

Table 1. Descriptive Statistics of Participants

| Category | Before Inversion (<i>n</i> = 2,211) (% of participants) | After Inversion (<i>n</i> = 9,996) (% of participants) |
|---------------------------------------|--|---|
| Female | 37.9 | 43.9 |
| Broadband connection | 91.3 | 89.4 |
| Has downloaded music from other sites | 69.4 | 65.3 |
| Country of Residence | | |
| United States | 68.1 | 54.7 |
| Brazil | 1.2 | 12.5 |
| Canada | 6.8 | 4.9 |
| United Kingdom | 6.6 | 6.9 |
| Other | 17.3 | 21.0 |
| Age | | |
| 17 and younger | 6.8 | 7.9 |
| 18 to 24 | 29.6 | 26.6 |
| 25 to 34 | 35.8 | 33.9 |
| 35 and older | 27.8 | 31.7 |

involved either amplifies or dampens the effect of initially false information.

In this paper we address the question of self-fulfilling prophecies in cultural markets by means of a web-based experiment where 12,207 participants were given the chance to listen to, rate, and download 48 previously unknown songs from unknown bands. Using a “multiple-worlds” experimental design (Salganik et al. 2006), described more fully below, we were able to simultaneously measure the appeal of the songs and measure the effect of initially false information on subsequent success. When deciding what initially false information to provide participants, we opted for an extreme approach, namely complete inversion of perceived success. This extreme approach is not meant to model an actual marketing campaign (which would likely focus on fewer songs), but rather to completely decouple perceived and actual success so as to explore self-fulfilling prophecies in a natural limiting case. While proponents of self-fulfilling prophecies might suspect that perceived success would overwhelm preexisting preferences and lead the market to lock-in to the inverted state, skeptics might suspect that preexisting preferences would overwhelm the false information and return the songs to their original ordering. Our results were more complex than either of these extreme predictions and therefore suggest the need for additional theoretical and empirical work.

EXPERIMENTAL SET-UP

Participants for our experiment were recruited by sending e-mails to participants from a previous, unrelated experiment (Dodds, Muhamad, and Watts 2003). These e-mails, and the additional web postings they generated, yielded 12,207 participants, the majority of whom lived in the United States and were between the ages of 18 and 34 (Table 1). Our experiment ran from March 14, 2005 to August 10, 2005 (21 weeks), and during that time we recorded a slight increase in the fraction of female participants and an increase in the proportion of participants from Brazil; however, it does not appear that either of these demographic shifts affected our results. Because the experiment was web-based, we had less control over participant recruitment and behavior than in lab-based experiments (Skitka and Sargis 2006). As such, we took a number of specific steps to account for potential data quality problems and these steps are described more fully in an online appendix (available at the SPQ website, www.asanet.org/spq).

Upon arriving at our website, participants were presented with a welcome screen informing them that they were about to participate in a study of musical tastes and that in exchange for participating they would be given a chance to download some free songs by up-and-coming artists. Next, subjects provided informed

consent,² filled out a brief survey, and were shown a page of instructions. Finally, subjects were presented with a menu of 48 songs presented in a vertical column, similar to the layout of popular music sites (Figure 1a).³ Having chosen to listen to a song, they were asked to rate it on a scale of 1 star (“I hate it”) to 5 stars (“I love it”) (Figure 1b), after which they were offered the opportunity to download the song (Figure 1c). Because of the design of our site, participants could only download a song after listening to and rating it. However, they could listen to, rate, and download as many or as few songs as they wished.

Upon arrival to the website, subjects were randomly assigned into one of a number of experimental groups. During an initial set-up phase, 2,211 participants were assigned to one of two “worlds”⁴—independent and social influence—which differed in the available information about the behavior of other participants that was presented with the songs. In the social influence world, the songs were sorted from most to least popular and accompanied by the number of previous downloads for each song. In the independent world, however, the songs were randomly reordered for each participant and were not accompanied by any measure of popularity. Although the presence or absence of download counts was not emphasized, the choices of participants in the social influence condition could clearly be influenced by the choices of previous participants, whereas no such influence was possible in the independent world.

² The research protocols used were approved by the Columbia University Institutional Review Board (protocol numbers: IRB-AAAA5286 and IRB-AAAB1483).

³ The songs were collected by sampling bands from the music website www.purevolume.com. These bands and songs were then screened to insure that they would be essentially unknown to experimental participants. A list of band names and song names is presented later in this paper (Table 2), and additional details on the sampling and screening of the bands are available in Salganik et al. (2006) and Salganik (2007).

⁴ In this paper we will use the term “world” rather than the more standard “condition” to emphasize the fact that although there were two different conditions—independent and social influence—there were a number of different groups to which participants could be assigned (see Figure 2).

After this set-up period, during which the popularity ordering of the songs, as measured by download counts, reached an approximately steady state, we continued to assign subjects to the social influence and independent world, but also created two new social influence worlds (the reason we created two will become clear shortly). In these new worlds, we explored the possibility of self-fulfilling prophecies by inverting the popularity order from the original social influence world. For example, at the time of the inversion “She Said” by Parker Theory had the most downloads, 128, and “Florence” by Post Break Tragedy had the fewest, 9. To construct the initial conditions for the inverted worlds, we swapped these counts, giving subsequent participants the false impression that “She Said” had 9 downloads and “Florence” had 128. As detailed in Table 2, we also swapped download counts for the 47th and 2nd songs, the 46th and 3rd songs, and so on. The task of inverting the songs was slightly complicated by the fact that there were a large number of songs with the same number of downloads; for example three different songs were tied with 13 downloads (Table 2). During the inversion these ties were broken randomly. After this one-time intervention, we updated all download counts honestly as 9,996 new participants listened to and downloaded songs in the four worlds—one unchanged social influence world, two inverted social influence worlds, and one independent world.

Although extremely simple in comparison with real cultural markets, our experimental design (illustrated schematically in Figure 2) offered us greater control than would be possible in real markets (Willer and Walker 2007) and exhibits several advantages over previous studies in this area (Hanson and Putler 1996; Sorensen 2007). First, the “multiple-worlds” (Salganik et al. 2006) feature of the design allows us to isolate the causal effect of an extreme manipulation—in this case complete inversion—by allowing us to compare participant-level, product-level, and collective-level outcomes in unchanged and inverted worlds; previous studies explored only modest manipulations and product-level outcomes. Second,

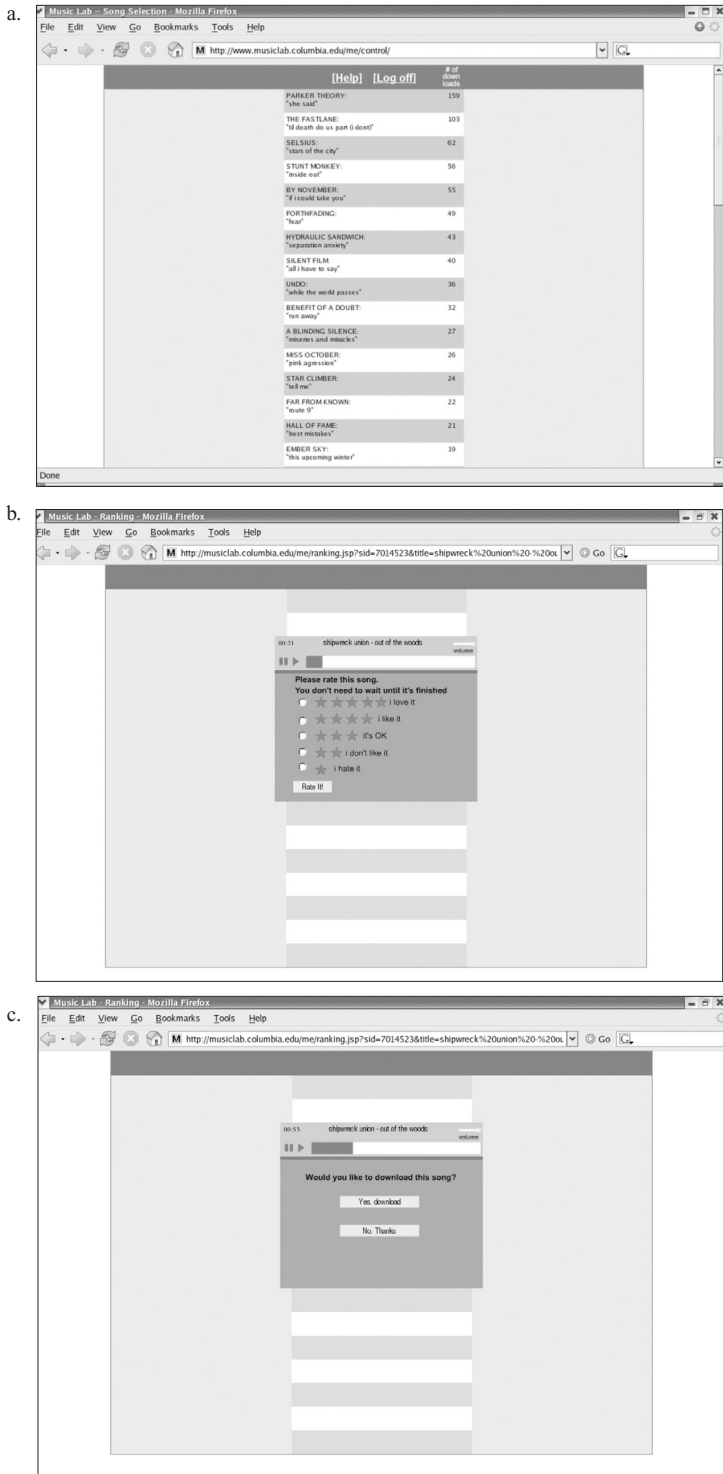


Figure 1. Screenshots from the Experiment
 a. A participant in the social information condition was presented the songs in a single column format sorted by previous number of downloads; screenshot from independent condition not shown.
 b. Having chosen to listen to a song, the participant was asked to rate it on a scale of 1 star (“I hate it”) to 5 stars (“I love it”).
 c. After rating the song, the participant was offered a chance to download the song.

Table 2. The 48 Songs Used in the Experiment Along with their Download Count Before and After Inversion

| Band Name | Song Name | Apparent downloads immediately after inversion | Apparent downloads immediately after inversion |
|----------------------|------------------------------------|--|--|
| Parker Theory | She Said | 128 | 9 |
| Simply Waiting | Went with the Count | 83 | 9 |
| Not for Scholars | As Seasons Change | 81 | 9 |
| Shipwreck Union | Out of the Woods | 66 | 9 |
| Sum Rana | The Bolshevik Boogie | 52 | 9 |
| Dante | Life's Mystery | 48 | 9 |
| Ryan Essmaker | Detour_(Be Still) | 45 | 9 |
| Hartsfield | Enough is Enough | 39 | 9 |
| By November | If I Could Take You | 36 | 9 |
| Star Climber | Tell Me | 35 | 9 |
| 52metro | Lockdown | 31 | 10 |
| Stranger | One Drop | 30 | 10 |
| Forthfading | Fear | 27 | 10 |
| Silverfox | Gnaw | 26 | 10 |
| Selsius | Stars of the City | 23 | 10 |
| Hydraulic Sandwich | Separation Anxiety | 22 | 10 |
| Undo | While the World Passes | 18 | 10 |
| Hall of Fame | Best Mistakes | 17 | 10 |
| The Thrift Syndicate | 2003 a Tragedy | 17 | 10 |
| Unknown Citizens | Falling Over | 16 | 11 |
| Beer bong | Father to Son | 14 | 11 |
| The Fastlane | Til Death do us Part (I don't) | 14 | 12 |
| Evan Gold | Robert Downey Jr. | 13 | 12 |
| Ember Sky | This Upcoming Winter | 13 | 13 |
| Miss October | Pink Aggression | 13 | 13 |
| Silent Film | All I have to Say | 12 | 13 |
| Stunt Monkey | Inside Out | 12 | 14 |
| Far from Known | Route 9 | 11 | 14 |
| Moral Hazard | Waste of my Life | 11 | 16 |
| Nooner at Nine | Walk Away | 10 | 17 |
| Sibrian | Eye Patch | 10 | 17 |
| Drawn in the Sky | Tap the Ride | 10 | 18 |
| Art of Kanly | Seductive Intro, Melodic Breakdown | 10 | 22 |
| Fading Through | Wish me Luck | 10 | 23 |
| Benefit of a Doubt | Run Away | 10 | 26 |
| Salute the Dawn | I am Error | 10 | 27 |
| Cape Renewal | Baseball Warlock v1 | 10 | 30 |
| Go Mordecai | It Does What Its Told | 10 | 31 |
| The Broken Promise | The End in Friend | 9 | 35 |
| Summerswasted | A Plan Behind Destruction | 9 | 36 |
| Secretary | Keep Your Eyes on the Ballistics | 9 | 39 |
| The Calefaction | Trapped in an Orange Peel | 9 | 45 |
| A Blinding Silence | Miseries and Miracles | 9 | 48 |
| Up Falls Down | A Brighter Burning Star | 9 | 52 |
| This New Dawn | The Belief Above the Answer | 9 | 66 |
| Up for Nothing | In Sight Of | 9 | 81 |
| Deep Enough to Die | For the Sky | 9 | 83 |
| Post Break Tragedy | Florence | 9 | 128 |

the popularity of songs in the independent world provides a natural measure of the preexisting preferences of the participant population that can then be compared to outcomes in the social influence worlds, allowing us to

determine the extent to which market information interacts with preexisting preferences; previous studies lack such a measure of participant preferences and thus could not address this important issue. Third, the two

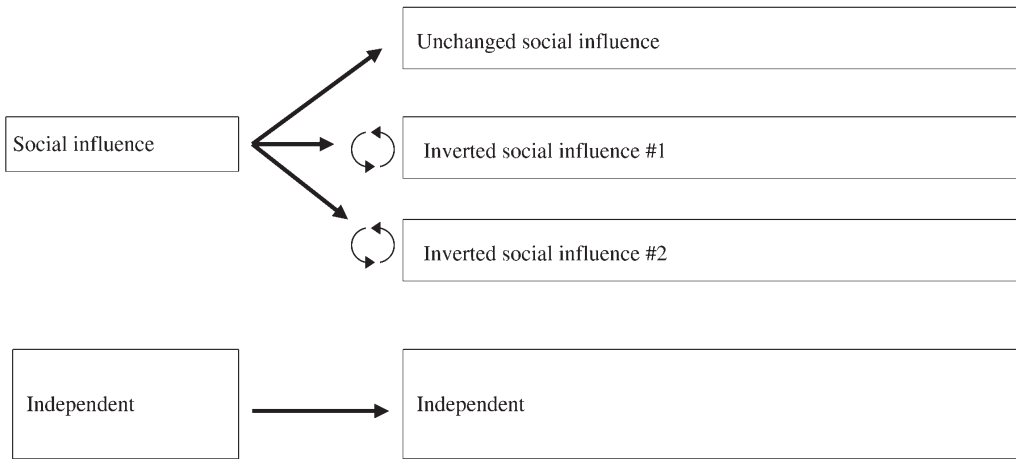
Before inversion**After inversion**

Figure 2. Schematic of the Experimental Design

inverted worlds capture purely random variations in outcomes of the same intervention;⁵ previous studies examined only one such outcome. A final benefit of our framework is that its dynamic nature allows us to track not just the response of individuals to the inversion, but the response of the entire system. This response can be observed over time, avoiding the need to choose some arbitrary point at which to measure an effect.

RESULTS

Descriptive statistics of participant behavior for the pre- and post-inversion periods are presented in Tables 3 and 4. Overall, the experiment involved 87,175 song-listens and 15,168 song-downloads meaning that, on average, participants listened to about seven songs each, of which they downloaded one.

⁵ Previous work has found that artificial markets, such as the one employed here, can have multiple, stable outcomes (Salganik et al. 2006). Therefore, we wanted to observe as many inverted worlds as possible to get a sense of the range of possible steady states that could result from a specific manipulation. At the same time, we wanted to have sufficient numbers of participants in each world to allow them to actually reach steady state after our exogenous shock (Strogatz 1994). The ultimate choice of two inverted worlds was an attempt to balance the conflicting desires for many worlds and many participants per world.

This ratio of listens to downloads suggests that the download decision was indeed viewed by participants as a consequential one. Moreover, the download decision was strongly related to participants' rating decision; the more stars a participant gave a song, the more likely they were to download that song.

At the time each subject participated, every song in his or her "world" had a specific download count and market rank (i.e., the song with the most downloads had a market rank of 1). Figures 4a–c plot the probability that participants listened to the song of a given market rank. These plots show that participants who were aware of the behavior of others were more likely to listen to songs that they believed were more popular. There was also a slight reversal of this pattern for the least popular songs—for example, a subject in a social influence world was about six times more likely to listen to the most popular song and three times more likely to listen to the least popular song, than to listen to a song of middle popularity rank. The tendency for the subjects to favor the least popular songs, which at first may appear surprising, is consistent with previous work (Salganik et al. 2006) and could simply be an artifact of our experimental set-up—just as the top spot on a list is salient, so too is the bottom.

Table 3. Descriptive Statistics of Subject Behavior Before the Inversion

| | Social Influence ($n = 752$) | Independent ($n = 1,459$) | Total ($n = 2,211$) |
|--------------------------|-----------------------------------|--------------------------------|--------------------------|
| Listens | 5,628 | 11,844 | 17,472 |
| Mean per subject | 7.5 | 8.1 | 7.9 |
| Downloads | 1,133 | 1,691 | 2,824 |
| Mean per subject | 1.5 | 1.2 | 1.3 |
| Pr[download listen] | .201 | .143 | .162 |
| Mean rating (# of stars) | 2.70 | 2.55 | 2.62 |

Table 4. Descriptive Statistics of Subject Behavior After the Inversion

| | Unchanged ($n = 2,015$) | Inverted #1 ($n = 2,014$) | Inverted #2 ($n = 1,970$) | Independent ($n = 3,997$) | Total ($n = 9,996$) |
|--------------------------|------------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------|
| Listens | 14,430 | 12,498 | 12,633 | 30,142 | 69,703 |
| Mean per subject | 7.2 | 6.2 | 6.4 | 7.5 | 7 |
| Downloads | 2,898 | 2,197 | 2,160 | 5,089 | 12,344 |
| Mean per subject | 1.4 | 1.1 | 1.1 | 1.3 | 1.2 |
| Pr[download listen] | .201 | .176 | .171 | .169 | .178 |
| Mean rating (# of stars) | 2.71 | 2.64 | 2.60 | 2.63 | 2.64 |

Alternatively, this behavior could represent a desire for some subjects to compare their preferences with their peers or could be a form of anti-conformist behavior (Simmel 1957; Heath, Ho, and Berger 2006); further experimental work would be required to adjudicate between these possibilities.

A related question is to what extent the observed effect of popularity on individual listening decisions is a consequence of informational versus normative influence (Deutsch and Gerard 1955). In this experiment we suspect that the bulk of the effect was informa-

tional; that is, we suspect that most participants interpreted the popularity of songs as a quality signal regardless of their preference for listening to the same songs as others. We do not know for sure, however, because our experiment was not designed to differentiate between these two kinds of influence since this differentiation was not necessary to address the main questions raised in this paper. Regardless of the precise mechanism involved, we can observe that the behavior of participants in the social influence worlds—both the unchanged (Figure 4a) and inverted

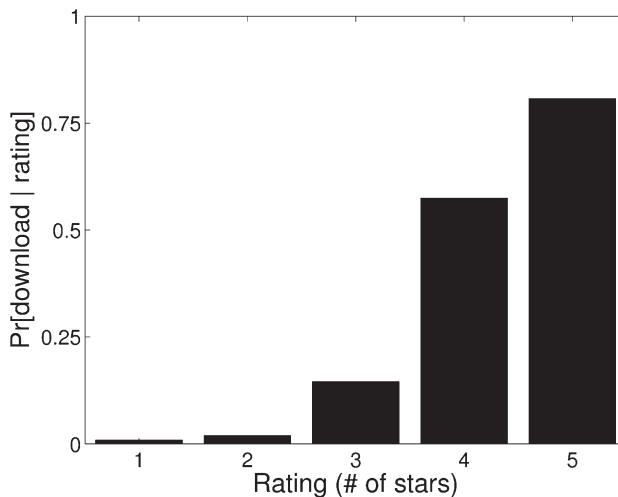


Figure 3. Probability that a Participant Downloaded a Song as a Function of How that Participant Rated the Song

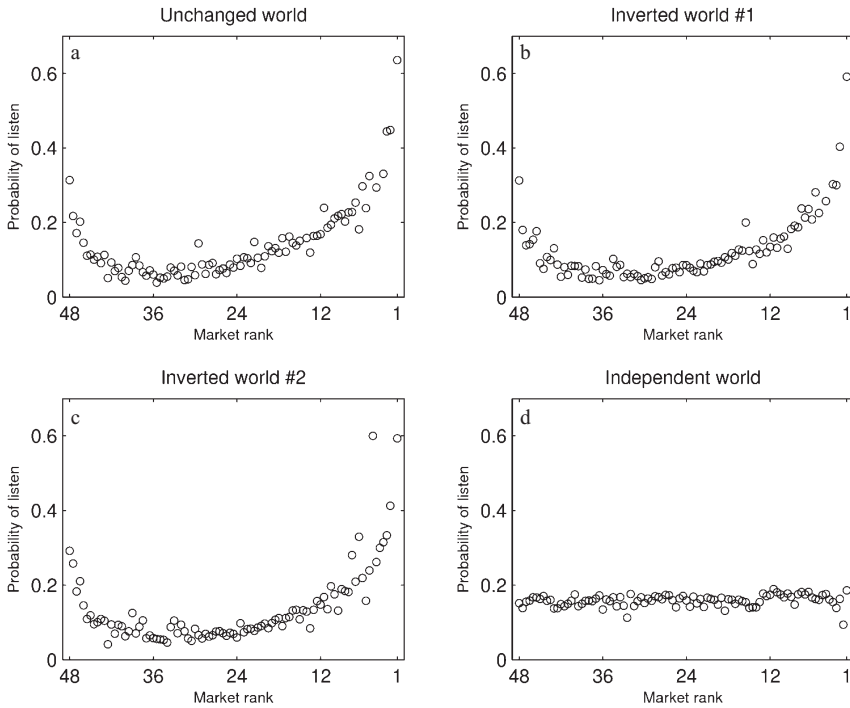


Figure 4. Probability of Listening to a Song as a Function of its Popularity in the Four Worlds

worlds (Figures 4b and 4c)—differed from behavior in the independent world (Figure 4d) where, given our design, song popularity

could have no effect on participants’ listening decisions.

Given that the subjects were influenced by the behavior of others (when they were made

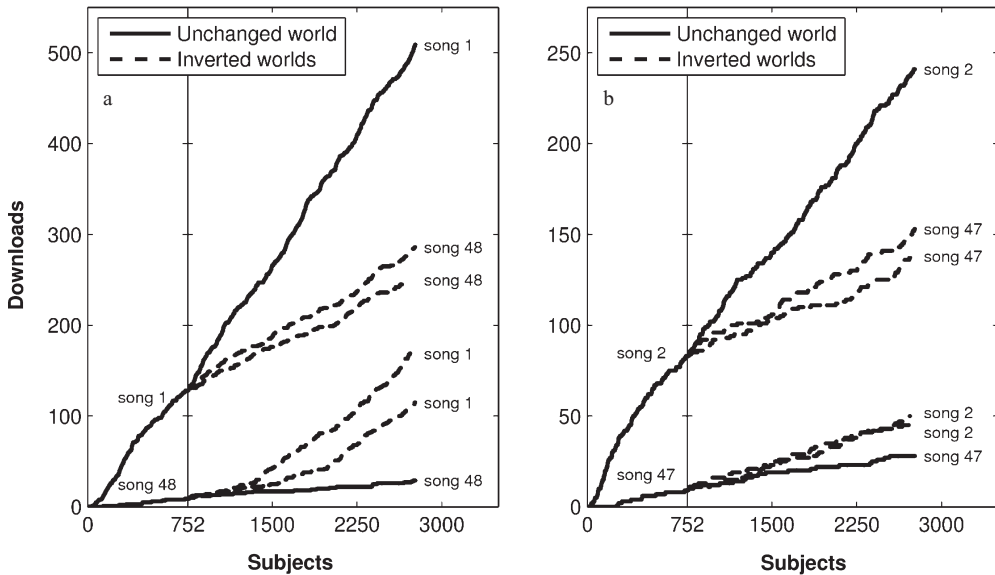


Figure 5. Popularity Dynamics Before and After the Inversion for the Most and Least Popular Song




Figure 6. Projected final rank K_i in the Unchanged World (a), and the Inverted Worlds (b and c), and as a Function of Market Rank at the End of the Set-up Period

aware of it), it is natural to explore the implications of the inversion on the subsequent success of specific songs. Figure 5 displays the number of downloads over time both in the unchanged (solid lines) and inverted (dashed lines) worlds for two pairs of songs: Figure 5a compares song 1, defined as the most popular song at the end of the set-up period (“She Said” by Parker Theory), with song 48, defined as the least popular song (“Florence” by Post Break Tragedy); and Figure 5b compares song 2 (“Went with the Count” by Simply Waiting) with song 47 (“For the Sky” by Deep Enough to Die). In both cases, download trajectories in the unchanged world were similar before and after the inversion.

The trajectories in the inverted worlds, however, were quite different. In Figure 5a, song 48 earned downloads at a faster rate as a consequence of the inversion (i.e., the slopes of the two dashed lines for song 48 are steeper than the corresponding solid line), while the pattern is reversed for song 1, which suffered from the inversion. A similar pattern is displayed in Figure 5b for song 47, which also benefited from the inversion, and song 2, which also suffered. For all four songs, the initially false perception of their popularity, arising from the inversion, caused their real popularity to change in the direction of the false belief.

It is impossible to say with certainty whether these dynamics would have led to permanent effects on the popularity of the songs, or whether the observed effects were merely transitory. The problem is that

although large, our pool of participants proved insufficient for either inverted world to reach a new steady state; note, for example, that in Figure 5a (dashed lines) song 1 still appears to be gaining on song 48 in both inverted worlds, where, interestingly, this does not seem to be the case for song 2 (Figure 5b, dashed lines). In spite of this limitation, we were nevertheless able to estimate the “final” rank of each song (i.e., the rank that would have been achieved if the experiment had run for many more participants), by linearly extrapolating the download trajectories for all 48 songs in each world.⁶ Figure 6A shows that ranks before the inversion and projected final ranks were highly correlated in the unchanged world ($r = 0.84$); whereas by contrast Figures 6b and 6c show they had a very weak relationship in the inverted worlds ($r = 0.16$ in both worlds).⁷ Figures 6b and 6c also confirm our intuition about Figure 5; that is, it appears that in both inverted worlds song 1 will eventually return to the top spot, but that song 2 will not return to the second spot.

⁶ The extrapolation of the download trajectories is based on a linear-least squares fit over the last 1,000 subjects in each world. Conclusions based on these projections are robust to the particular number of subjects used in the fitting (Salganik 2007). It is also the case that the projected final rank was highly correlated with the ranking at the end of the experiment ($r = 0.92$ in the unchanged world and $r = 0.83$ and 0.84 in the inverted worlds).

⁷ The vertical banding in Figure 6 is caused by the large number of songs that had the same market rank at the end of the set-up period (see Table 2).




Figure 7. Relationship Between Estimated Impact of Inversion, Δ , and Market Rank at the End of the Set-up Period (Both Estimates Presented on the Same Axis)

Further, given the projected outcomes of each song, we can compare the projected final rank in the unchanged and inverted worlds to estimate the long-term impact of the inversion on the popularity of individual songs. More specifically, the impact of the inversion on a specific song, Δ_p , is defined to be the difference between the projected final rank in the unchanged world and the projected final rank in an inverted world ($K_{i,unc} - K_{i,inv}$). Since we have two inverted worlds, moreover, we have two estimates of our impact measure Δ_p . Figure 7 shows the impact of inversion as a function of the rank during the set-up period and reveals, for example, that the most popular song during the set-up period was projected to finish at the same rank in the inverted and unchanged worlds; thus the impact of inversion for this song was 0.⁸ The second most popular song during the set-up period, however, was projected to finish higher in the unchanged world than in the inverted worlds (5th vs 20th and 23rd), so for this song the impact of the inversion was negative (i.e., it was hurt by the inversion). Overall, the final rankings of almost all songs seem to be permanently affected by the inversion, where songs that were promoted by the inversion

tended to do better in the long run, and songs that were initially demoted tended to do worse. Thus, in our experiment, the manipulation of market information, combined with a process of social influence, seemed to lead to long-term changes in the popularity of songs.

Whereas Figures 6 and 7 show the outcomes experienced by individual songs, Figure 8 shows the outcome experienced by the entire “market”—specifically, it shows the Spearman rank correlation $\rho(t)$ between popularity at the end of the set-up period and popularity in the three social influence worlds as a function of time.⁹ During the set-up period (to the left of the vertical line) the initial social influence world quickly converged to an approximate steady state, as evidenced by the continuing high value of $\rho(t) \approx 1$ for the unchanged world (solid line) after the inversion (to the right of the vertical line). By contrast, the two inverted worlds (dashed lines) started, by definition, with $\rho(t) = -1$, after which both increased monotonically towards what appears to be an asymptotic limit around zero. In other words, popularity in the inverted worlds moved to a state that had, in effect, no relationship with the popularity before the

⁸ The vertical banding in Figure 7 is caused by the large number of songs that had the same market rank at the end of the set-up period (see Table 2).

⁹ We used the Spearman rank correlation instead of the more familiar Pearson product-moment correlation (r) because the former is non-parametric measure of monotonic association, rather than just linear association (Kendall and Gibbons 1990).

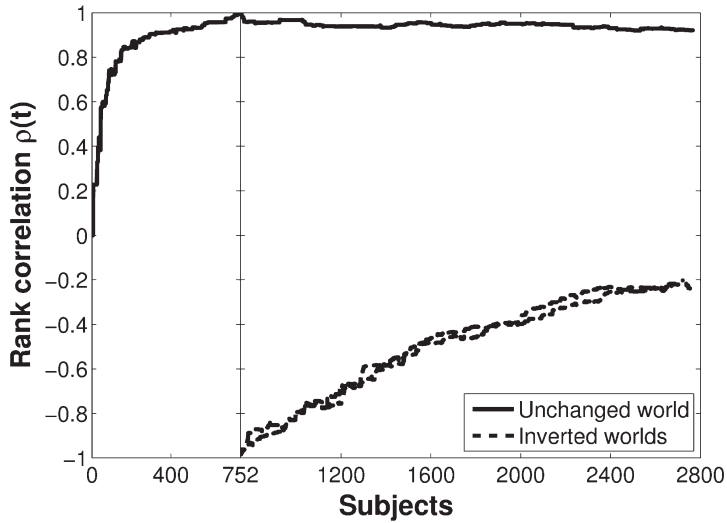


Figure 8. Spearman Rank Correlation, $\rho(t)$, between Popularity at the End of the Set-up Period and Popularity in the Unchanged and Inverted Worlds

inversion. Figures 7 and 8 therefore represent slightly different perspectives on the effect of the inversion: whereas Figure 7 suggests that almost all songs, considered individually, did experience some effect on long-term popularity as a consequence of the inversion, Figure 8 suggests that these effects were not strong enough for the inversion to lock in at the level of the entire market.

The Role of Appeal

An obvious explanation for the failure of the inversion to lock in is that the songs themselves were of different quality, and that these differences were more salient than the perception of popularity. Previous theoretical work by economists has indeed emphasized the importance of intrinsic quality on outcomes (Rosen 1981), and attempts have been made to measure the quality of cultural products (Hamlen Jr. 1991; Krueger 2005). Unfortunately, no generally agreed upon measure of quality exists, in large part because quality is largely, if not completely, a social construction (Gans 1974; Becker 1982; Bourdieu 1984; DiMaggio 1987; Cutting 2003; Frith 2004). Fortunately, our experimental design permits us to proceed without resolving these conceptual difficulties by measuring instead the intrinsic “appeal” of the

songs to our pool of participants. Because the behavior of the participants in the independent condition reflects their true preferences (i.e. they were not subject to social influence), the appeal α_i of each song i can be defined as the market share of downloads that that song earned in the independent condition,

$\alpha_i = \frac{d_i}{\sum d}$, where d_i is the number of downloads for song i in the independent condition (Salganik et al. 2006).¹⁰

We emphasize that our measure of appeal is not truly a measure of quality, in part because it is specific to our subject pool. Were we to rerun the experiments with a new population of participants, such as senior citizens from Japan, our procedure would almost certainly result in a different measure of appeal. Further, because the measure of appeal is based on market share, it is constrained to sum to 1; thus if we added another song to our

¹⁰ An alternate measure of the appeal of a song is the probability of downloading that song given a listen (see, for example, Aizen et al. 2004). We chose not to use this measure, however, because it does not include a measure of the attractiveness of the song and band names, something that was found to vary across songs. Our proposed measure, because it is based on the probability of listen and the probability of download given listen, includes both the attractiveness of the song itself and the attractiveness of the song and band name.

experiment, the measured appeal of the other songs would decrease. While these characteristics limit the general applicability of our measure, fortunately they do not affect the current analysis.

Figure 9 compares our measure of appeal to our estimates of the impact of the inversion, Δ , and shows that “bad” songs benefited from the inversion, while the “good” songs suffered. Moreover, Figure 9 shows that the success of the very “best” songs was essentially unaffected, even though these were typically the most severely penalized by the inversion. This tendency for the highest appeal songs to recover their original ranking coupled with the tendency for lower appeal songs to maintain at least some of the benefits of the inversion prevented the songs from either locking in to their inverted ordering or returning to their preinversion state.

We can also use our measure of appeal to see how closely the projected outcomes in the social influence worlds reflect the “true” preferences of our population, as revealed by their choices in the independent condition. In the unchanged social influence world, the projected outcomes were strongly, but not completely, related to appeal (rank correlation = 0.82). Therefore, even in the absence of manipulation, the presence of social influence produced outcomes that did not perfectly reflect the true preferences of the participant population

(Salganik et al. 2006). The inverted worlds, however, were much less reflective of population preferences (rank correlation = 0.40 and 0.45), suggesting—perhaps not surprisingly—that markets in which perceived popularity has been manipulated will in general be less revealing of true preferences than markets in which popularity is allowed to emerge naturally.

The Consequences for Participants

A final and unexpected consequence of the inversion was a substantial reduction in the overall number of downloads. As shown in Figure 4, subjects in all social influence worlds tended to listen to the songs that they thought were more popular. In the inverted worlds, however, the songs that appeared to be more popular tended to be of lower appeal; thus, subjects in the inverted world were more exposed to lower appeal songs. For example, in the unchanged world, the ten highest appeal songs had about twice as many listens as the ten lowest appeal songs, but in the inverted worlds this pattern was reversed with the ten lowest appeal songs having twice as many listens. As a consequence, subjects in the inverted worlds left the experiment after listening to fewer songs and were less likely to download the songs to which they did listen (Table 4). Together, these effects led to a substantial reduction in downloads: 2,197 and 2,160 in




Figure 9. Relationship between Estimated Impact of Inversion, Δ , and Song Appeal (Both Estimates Presented on the Same Axis) with Dashed Line Showing Loess Fit to the Data Intended as a Visual Aid (Cleveland 1993)

the inverted worlds, compared with 2,898 in the unchanged world.

The combination of increased success for some individual songs (Figure 7) on the one hand, and decreasing overall downloads, on the other hand, suggests that the choice to manipulate market information may resemble a social dilemma, familiar in studies of public goods and common-pool resources (Dawes 1980; Yamagishi 1995; Kollock 1998), but less evident in market-oriented behavior. Specifically, Figure 7 suggests that any individual band could expect to benefit by artificially inflating their perceived popularity, regardless of their true appeal or the strategies of the other bands; thus all bands have a rational incentive to manipulate information. When too many bands employ this strategy, however, the correlation between apparent popularity and appeal is lowered, leading to the unintended consequence of the market as a whole contracting, thereby causing all bands to suffer collectively (Dellarocas 2006).¹¹

DISCUSSION AND CONCLUSION

Although Merton's concept of the self-fulfilling prophecy is appealing both for its elegance and its generality, social scientists have encountered difficulty in demonstrating empirically that self-fulfilling prophecies actually occur, and that observed outcomes do not instead reflect exogenous factors like intrinsic differences in quality or convergence to rational equilibria. In particular, convincingly rejecting alternative explanations in favor of a self-fulfilling prophecy requires one to compare the outcomes of otherwise identical "versions" of history, some of which include the false perception of reality, and others of which do not—clearly an impossibility

for most real-world social processes. Moreover, in many cases of interest, including cultural markets, the identification of self-fulfilling prophecies is complicated by the presence of multiple "scales" in the dynamics; that is, the decisions required to render the false belief true are made by individuals, but the false belief itself concerns some collective property, like aggregate number of downloads, over which no one individual has much control. Causal explanations of collective social phenomena that invoke self-fulfilling prophecies are therefore rendered problematic not only by the absence of counterfactuals in most observational data, but by the analytical complexity of the micro-macro problem (Schelling 1978; Coleman 1990; Hedström 2005).

By conducting an experiment on a large enough scale that we can observe both individual choices and collective dynamics simultaneously, our study sidesteps these difficulties, allowing us not only to identify the presence of self-fulfilling prophecies—when they occur—but also to begin to quantify their effects. We are able to show, for example, that although inversion of market information can lead to substantial differences in the success of individual songs, the effect on the overall market ranking was not as dramatic as we had anticipated—many "good" songs recouped much of their original popularity in spite of our manipulation. Our experiment therefore provides some ammunition both for proponents of self-fulfilling prophecies, and also for skeptics, by suggesting that cultural markets can exhibit self-fulfilling prophecies, but that their effects may be limited by preexisting individual preferences.

Naturally, our experiment is unlike real cultural markets in a number of respects that render our results more suggestive than conclusive. For example, unlike in our experiment where popularity was manipulated at a single time and in a highly artificial manner (i.e., total inversion), distortion in the real world may occur repeatedly, and may also exhibit considerable subtlety and variety. As an extreme example of manipulation, total inversion seems a natural first case to consider; but there are of course many other possible manipulations that one could explore, even

¹¹ The dilemma faced by the bands appears to be more similar to common-pool resource situations than public goods situations because the benefit that a band receives may be related to their proportion of the total contribution, not just to the total contribution (Apestequia and Maier-Rigaud 2006). However, this statement is hard to make precisely because the payoff functions for the bands are unknown. For more on the difference between common-pool resource and public goods situations see Apestequia and Maier-Rigaud (2006).

within our simplified experimental framework. Moreover, whereas our subjects were exposed to only a single source of influence—download counts—information in real cultural markets exhibits richer content (Chevalier and Mayzlin 2006) and comes from a variety of sources (Katz and Lazarsfeld 1955). Also, because our experiment had only 48 songs, the average participant listened to about one-seventh of the music in our market and some participants listened to almost all the songs—a feat that would be impossible in real cultural markets where the number of products is overwhelming (Caves 2000). Finally, by focusing exclusively on consumer decisions, we did not account for the decisions of institutional actors such as music executives, radio stations, and cultural critics; actors who can have important effects on outcomes as has been demonstrated by numerous scholars working in the “production of culture” school (Hirsch 1972; Peterson 1976; Frith 1978; DiMaggio 2000; Peterson and Anand 2004; Dowd 2004).

Precisely how our results would change under more realistic conditions is difficult to predict. We suspect, for example, that our finding that the highest appeal songs tend to succeed regardless of interference may derive from the relatively small number of songs, which prevented the “best” songs from escaping notice even in the inverted worlds. Thus this finding may not generalize to more realistic scenarios in which the number of songs is much greater. Moreover, because we only performed one type of manipulation on one set of songs, it is unclear how our findings would be affected either by less severe distortions or by using a set of songs that are more (or less) similar in terms of appeal. Nor is it obvious how the results would have differed had our subjects been exposed to a stronger (or weaker) form of social influence.

In spite of these ambiguities, which we hope will be addressed with additional experiments or simulations, we believe that our findings are likely to have applicability beyond the specific scope of the experiment itself, and thereby add to our general understanding of self-fulfilling prophecies in cultural markets. We also believe this experiment may have implications for experimental soci-

ology and social psychology more generally by showing the potential for web-based experiments to operate on a scale that is not possible in a physical lab (Skitka and Sargis 2006). Our experiment involved more than 12,000 participants—a number which, to place in the context of traditional psychology experiments, is larger than the total enrollment of many universities. Even larger experiments are practical today, and likely to become increasingly so as web-related technology continues to develop. Although there are a number of important issues to consider when conducting web-based experiments—some of which are shared with laboratory experiments, and some of which are novel—we suspect that the ability to run experiments involving tens, or even hundreds, of thousands of participants will open exciting new areas of theory development and testing.

For example, both sociologists (DiMaggio 1997) and psychologists (Schaller and Crandall 2003) have recently taken an interest in the psychological foundations of culture, arguing that “Individuals’ thoughts, motives, and other cognitions govern how they interact with and influence one another; these interpersonal consequences in turn govern the emergence, persistence, and change of culture” (Schaller and Crandall:4). Economists, sociologists, and physicists, moreover, have proposed numerous mathematical and simulation models that purport to represent how interpersonal influence—a micro-level phenomenon—aggregates to produce macro-level phenomena like information cascades, winner-take-all markets, and the successful diffusion of innovations.

Although these modeling exercises have led to some intriguing and even counterintuitive insights, they have also been confounded by the difficulty of reconciling models either with micro-level or macro-level empirical data. At the micro-level, empirical difficulties arise because social influence experiments are not generally designed to differentiate between the different “rules” governing individual behavior that are assumed, sometimes implicitly, in various models (Lopez-Pintado and Watts 2008). And at the macro-level, empirical verification is plagued by ambiguities of cause and effect; that is, very different

individual-level rules can sometimes generate very similar collective outcomes, while at other times very different collective outcomes can be generated by indistinguishable rules (Granovetter 1978; Watts 2002). By dramatically increasing the scale at which controlled experiments can be conducted, while still retaining the ability to measure individual-level responses, web-based experiments, of which the one described here is merely illustrative, can therefore help address the micro-macro problem inherent in understanding the psychological foundations of cultural markets and other macrosociological phenomena (Hedström 2006).

REFERENCES

- Adler, Moshe. 1985. "Stardom and Talent." *American Economic Review* 75(1):208–12.
- Aizen, Jonathan, Daniel Huttenlocher, Jon Kleinberg, and Antal Novak. 2004. "Traffic-based Feedback on the Web." *Proceedings of the National Academy of Sciences USA* 101 (Supp. 1):5254–60.
- Apestequia, Jose, and Frank P. Maier-Rigaud. 2006. "The Role of Rivalry: Public Goods Versus Common-Pool Resources" *Journal of Conflict Resolution* 50(5):646–63.
- Asch, Solomon E. 1952. *Social Psychology*. Englewood Cliffs, NJ: Prentice-Hall.
- Becker, Howard S. 1982. *Art Worlds*. Berkeley, CA: University of California Press.
- Biggs, Michael. Forthcoming. "Self-Fulfilling Prophecies." In *The Oxford Handbook of Analytical Sociology*, edited by Peter Bearman and Peter Hedström. Oxford, UK: Oxford University Press.
- Bourdieu, Pierre. 1984. *Distinction*. Translated by R. Nice. Cambridge, MA: Harvard University Press.
- Callon, Michel, ed. 1998. *The Laws of the Markets*. Oxford, UK: Blackwell.
- Calomiris, Charles W. and Joseph R. Mason. 1997. "Contagion and Failures During the Great Depression: The June 1932 Chicago Banking Panic." *American Economic Review* 87(5): 863–83.
- Caves, R. E. 2000. *Creative Industries: Contracts between Art and Commerce*. Cambridge, MA: Harvard University Press.
- Chevalier, Judith A. and Dina Mayzlin. 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research* 43:345–54.
- Christakis, Nicholas A. 1999. *Death Foretold: Prophecy and Prognosis in Medical Care*. Chicago, IL: University of Chicago Press.
- Cialdini, Robert B. and Noah J. Goldstein. 2004. "Social Influence: Compliance and Conformity." *Annual Review of Psychology* 55:591–621.
- Cleveland, William S. 1993. *Visualizing Data*. Summit, NJ: Hobart Press.
- Coleman, James S. 1990. *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.
- Cutting, James E. 2003. "Gustave Caillebotte, French Impressionism, and Mere Exposure." *Psychonomic Bulletin & Review* 10(2):319–43.
- Dawes, Robyn M. 1980. "Social Dilemmas." *Annual Review of Psychology* 31:169–93.
- Dellarocas, Chrysanthos. 2006. "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms." *Management Science* 52(10):1577–93.
- Deutsch, M., and H. B. Gerard. 1955. "A Study of Normative and Informative Social Influences Upon Individual Judgement." *Journal of Abnormal Social Psychology* 51:629–36.
- DiMaggio, Paul. 1987. "Classification in Art." *American Sociological Review* 52(4):440–55.
- . 1997. "Culture and Cognition." *Annual Review of Sociology* 23:263–87.
- . 2000. "The Production of Scientific Change: Richard Peterson and the Institutional Turn in Cultural Sociology." *Poetics* 28:107–36.
- Dodds, P. S., R. Muhamad, and D. J. Watts. 2003. "An Experimental Study of Search in Global Social Networks." *Science* 301 (5634):827–9.
- Dowd, Timothy J. 2004. "Production Perspectives in the Sociology of Music." *Poetics* 32:235–46.
- Farmer, Robert E. A. 1999. *Macroeconomics of Self-fulfilling Prophecies*. Cambridge, MA: MIT Press.
- Frith, Simon. 1978. *The Sociology of Rock*. London, UK: Constable.
- . 2004. "What is bad music?" Pp 15–36 in *Bad Music*, edited by Christopher J. Washburne and Maiken Derno. New York: Routledge.
- Gans, Herbert. 1974. *Popular Culture and High Culture*. New York: Basic Books.
- Garber, Peter M. 1989. "Tulipmania." *Journal of Political Economy* 97(3):535–60.
- Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83(6):1420–43.
- Hamlen Jr., William A. 1991. "Superstardom in Popular Music: Empirical Evidence." *The Review of Economics and Statistics* 73 (4):729–33.
- Hanson, Ward A. and David S. Putler. 1996. "Hits and Misses: Herd Behavior and Online Product Popularity." *Marketing Letters* 7(4):297–305.
- Heath, Chip, Ben Ho, and Jonah Berger. 2006. "Focal points in coordinated divergence." *Journal of Economic Psychology* 27:635–47.

- Hedström, Peter. 1998. "Rational Imitation." Pp 306–27 in *Social Mechanisms: Analytical Approach to Social Theory*, edited by Peter Hedström and Richard Swedberg. Cambridge, UK: Cambridge University Press.
- . 2005. *Dissecting the Social*. Cambridge, UK: Cambridge University Press.
- . 2006. "Experimental Macro Sociology: Predicting the Next Best Seller." *Science* 311:786–7.
- Hirsch, Paul M. 1972. "Processing Fads and Fashions: An Organization-Set Analysis of Cultural Industry Systems." *American Journal of Sociology* 77(4):639–59.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81:945–60.
- Huang, Jen-Hung and Yi-Fen Chen. 2006. "Herding in Online Product Choice." *Psychology & Marketing* 23(5):413–28.
- Jussim, Lee and Kent D. Harber. 2005. "Teacher Expectations and Self-Fulfilling Prophecies: Knowns and Unknowns, Resolved and Unresolved Controversies." *Personality and Social Psychology Review* 9(2):131–55.
- Katz, Elihu and Paul F. Lazarsfeld. 1955. *Personal Influence*. Glencoe, IL: Free Press.
- Kendall, Maurice and Jean Dickinson Gibbons. 1990. *Rank Correlation Methods*. London, UK: Edward Arnold.
- Kollock, Peter. 1998. "Social Dilemmas: The Anatomy of Cooperation." *Annual Review of Sociology* 24:183–214.
- Krueger, Alan B. 2005. "The Economics of Real Superstars: The Market for Rock Concerts in the Material World." *Journal of Labor Economics* 23(1):1–30.
- Lopez-Pintado, Dunia and Duncan J. Watts. 2008. "Social Influence, Binary Decisions, and Collective Dynamics." *Rationality and Society* 20(4):399–443
- MacKenzie, Donald and Yuval Millo. 2003. "Constructing a Market, Performing Theory: The Historical Sociology of a Financial Derivatives Exchange." *American Journal of Sociology* 109(1):107–45.
- Merton, Robert K. 1948. "The Self-Fulfilling Prophecy." *Antioch Review* 8:193–210.
- . 1995. "The Thomas Theorem and The Matthew Effect." *Social Forces* 74(2):379–424.
- Peterson, Richard A. 1976. "The Production of Culture." *The American Behavioral Scientist* 19:669–84.
- Peterson, Richard A. and N. Anand. 2004. "The Production of Culture Perspective." *Annual Review of Sociology* 30:311–34.
- Rogerson, Richard. 1995. "Book Review: The Macroeconomics of Self-Fulfilling Prophecies by Roger E. A. Farmer." *Journal of Economic Literature* 33(2):841–2.
- Rosen, Sherwin. 1981. "The Economics of Superstars." *American Economic Review* 71(5):845–58.
- Rosenthal, Robert and Lenore Jacobson. 1968. *Pygmalion in the Classroom: Teacher Expectations and Pupils' Intellectual Development*. New York: Holt, Rinehart and Winston.
- Salganik, Matthew J. 2007. Success and Failure in Cultural Markets, PhD thesis, Department of Sociology, Columbia University, New York.
- Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science* 311:854–6.
- Schaller, Mark and Christian S. Crandall, eds. 2003. *The Psychological Foundations of Culture*. Mahwah, NJ: Lawrence Erlbaum.
- Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: W. W. Norton.
- Senecal, Sylvain and Jacques Nantel. 2004. "The Influence of Online Product Recommendations on Consumers' Online Choices." *Journal of Retailing* 80:159–69.
- Sherif, Muzafer. 1937. "An Experimental Approach to the Study of Attitudes." *Sociometry* 1(1):90–8.
- Simmel, Georg. 1957. "Fashion." *American Journal of Sociology* 62(6):541–58.
- Skitka, Linda J., and Edward G. Sargis. 2006. "The Internet as Psychological Laboratory." *Annual Review of Psychology* 57:529–55.
- Sobel, Michael E. 1996. "An Introduction to Causal Inference." *Sociological Methods and Research* 24:353–79.
- Sorensen, Alan T. 2007. "Bestseller Lists and Product Variety." *Journal of Industrial Economics* LV(4):715–38.
- Stewart-Williams, Steve and John Podd. 2004. "The Placebo Effect: Dissolving the Expectancy Versus Conditioning Debate." *Psychological Bulletin* 130(2):324–40.
- Strogatz, Steven H. 1994. *Nonlinear Dynamics and Chaos*. Reading, MA: Perseus.
- Thomas, W. I., and Dorothy Swaine Thomas. 1928. *The Child in America: Behavior Problems and Programs*. New York: Knopf.
- Watts, Duncan J. 2002. "A Simple Model of Global Cascades in Random Networks." *Proceedings of the National Academy of Sciences, USA* 99(9):5766–71.
- Willer, David and Henry A. Walker. 2007. *Building Experiments: Testing Social Theory*. Stanford, CA: Stanford University Press.
- Winship, Christopher and Stephen L. Morgan. 1999. "The Estimation of Causal Effects from

Observational Data." *Annual Review of Sociology* 25:659–706.

Yamagishi, Toshio. 1995. "Social Dilemmas." Pp 331–35 in *Sociological Perspectives on Social*

Psychology, edited by Karen S. Cook, Gary Alan Fine, and James S. House. Boston, MA: Allyn and Bacon.

Matthew J. Salganik is an assistant professor in the Department of Sociology and Office of Population Research at Princeton University. He received his PhD in sociology from Columbia University in 2007. His research interests include social networks, sampling methods for hidden populations, and web-based social research.

Duncan J. Watts is principal research scientist at Yahoo! Research, where he directs the Human Social Dynamics group, and a professor of sociology at Columbia University. His research interests include the structure and evolution of social networks, the origins and dynamics of social influence, and the nature of distributed "social" search. He holds a BSc in Physics from the University of New South Wales and a PhD in Theoretical and Applied Mechanics from Cornell University.