

Web-Based Experiments for the Study of Collective Social Dynamics in Cultural Markets

Matthew J. Salganik,^a Duncan J. Watts^b

^a*Department of Sociology and Office of Population Research, Princeton University*

^b*Yahoo! Research, New York*

Received 7 July 2008; received in revised form 11 April 2009; accepted 1 May 2009

Abstract

Social scientists are often interested in understanding how the dynamics of social systems are driven by the behavior of individuals that make up those systems. However, this process is hindered by the difficulty of experimentally studying how individual behavioral tendencies lead to collective social dynamics in large groups of people interacting over time. In this study, we investigate the role of social influence, a process well studied at the individual level, on the puzzling nature of success for cultural products such as books, movies, and music. Using a “multiple-worlds” experimental design, we are able to isolate the causal effect of an individual-level mechanism on collective social outcomes. We employ this design in a Web-based experiment in which 2,930 participants listened to, rated, and downloaded 48 songs by up-and-coming bands. Surprisingly, despite relatively large differences in the demographics, behavior, and preferences of participants, the experimental results at both the individual and collective levels were similar to those found in Salganik, Dodds, and Watts (2006). Further, by comparing results from two distinct pools of participants, we are able to gain new insights into the role of individual behavior on collective outcomes. We conclude with a discussion of the strengths and weaknesses of Web-based experiments to address questions of collective social dynamics.

Keywords: Cultural markets; Web-based experiments; Superstars; Winner-take-all; Cumulative advantage; Social influence

1. Introduction

Can you study an army in a laboratory? In 1969, the sociologist Morris Zeldich posed this provocative question about the study of group behavior. At the time Zeldich concluded that

Correspondence should be sent to Matthew J. Salganik, Department of Sociology and Office of Population Research, Princeton University, Wallace Hall, Princeton, NJ 08544. E-mail: mjs3@princeton.edu

because logistical difficulties prevented researchers from studying large groups literally “in a lab,” the more realistic way to learn about these groups was by studying the smaller-scale process at work within them (Zelditch, 1969). Almost 40 years later, however, it may be the time to revisit Zelditch’s question with a fresh perspective. In particular, the vast increase in computing power over the past few decades, and the almost limitless pool of participants now available via the Internet have made it possible—although by no means trivial—to design and conduct laboratory-style experiments involving thousands, or even millions, of participants. Moving beyond the classic study of “small groups” (Harrington & Fine, 2000), we can now begin the study of “large groups” (Hedström, 2006).

But first it might be asked, why the ability to study large groups experimentally is needed when we already have the ability to study individuals experimentally? An example that combines the individual-level process of anchoring and adjustment (Tversky & Kahneman, 1974) with the collective dynamics of auctions proves instructive. Individuals have been found to anchor their valuation of objects based on irrelevant (Ariely, Lowenstein, & Prelec, 2003; Simonson & Drolet, 2004) and even false information (Northcraft & Neale, 1987). For example, real estate agents’ valuations of homes have been shown to be affected by false listing prices (Northcraft & Neale, 1987). Therefore, because of this psychological mechanism, one might predict that auctions with higher starting prices would have higher final sale prices because auction participants, like the real-estate agents, would anchor on the higher starting price. Yet in real auctions exactly the opposite can occur: Lower starting prices can lead to higher final prices (Ku, Galinsky, & Murnighan, 2006). In other words, even though the individual psychology of anchoring has been studied extensively, it offers us the wrong prediction about the collective social dynamics of auctions.

More generally than auctions, it seems to be difficult to predict the behavior of systems with many interacting components, even with a reasonable understanding of these lower-level components. That is, just as it is hard to reduce biology to the behavior of molecules (Hartwell, Hopfield, Leibler, & Murray, 1999), it is hard to reduce group behavior to individual psychology (Barton, 1968); in the words of Anderson (1972), “more is different.” This observation does not imply that individual behavioral tendencies tell us nothing about group behavior, as has been argued by proponents of the “collective mind” and the “madness of crowds.” However, it does suggest that models of collective behavior that are based solely on the behavior of a “representative agent” are seriously deficient (Kirman, 1992). The challenge for studying collective behavior is therefore in understanding how the actions and interactions of individuals aggregate to produce collective patterns (Allport, 1924; Coleman, 1990; Hedström, 2005; Schelling, 1978).

Much of the formal work connecting individual behavior and group dynamics has relied on mathematical and agent-based models (Arthur, 1994; Axelrod, 1997; Banerjee, 1992; Bikhchandani, Hirshleifer, & Welch, 1992; Goldstone, Ashpole, & Roberts, 2005; Macy & Willer, 2002; Watts, 2002), but these approaches are invariably subject to the critique that they are based on overly simplistic representations of human behavior and so are of questionable relevance to real-world social phenomena. In this study, therefore, we will attempt to address these same issues with experiments involving human subjects making decisions that—although extremely simple—are not so different from decisions they make in their

everyday lives. Specifically, we will investigate the effect of social influence, a process well studied at the individual level, on the success of products in cultural markets.

The study will be organized as follows. First, we will describe the puzzling nature of success and failure in cultural markets, a collective outcome, and then show how the “multiple-worlds” design of Salganik et al. (2006) allows us to address the role of individual-level social influence on collective outcomes. Unlike most experiments in cognitive science, psychology, and economics that have individuals as the unit of analysis, these experiments have groups of several hundred people as the unit of analysis, and therefore require thousands of participants, thereby running directly into the constraints of the physical laboratory that troubled Zeldich. We were able to circumvent these problems by running our experiment over the Web, allowing us to accommodate many more participants than have been possible in a physical laboratory. After reviewing the results from two previous experiments whose participants were largely American teenagers (Salganik et al., 2006), we will present new results with an entirely distinct pool of participants with different demographics, preferences, and behavior ($n = 2,930$). Comparison of these results reveals additional features of the role of individual behavior on collective outcomes that would not be discernible by studying only one participant population. The study concludes with a discussion of the possibilities and limitations of the Web-based experimental study of large group behavior.

2. Success and failure in cultural markets

An obvious “stylized fact” about markets for cultural products such as books, movies, TV shows, and music is that success in such markets is tremendously unequal (Caves, 2000; Chung & Cox, 1994; De Vany, 2004; Sorensen, 2007; Vogel, 2004). That is, blockbuster movies like *Star Wars* and bestselling books like *Harry Potter* are not merely more successful than average—they are orders of magnitude more successful; so much so that cultural markets are often characterized as “superstar” (Rosen, 1981) or “winner-take-all” (Frank & Cook, 1995) markets. On its own, there is nothing particularly puzzling about this observation: Because the cost of consuming cultural products is generally unrelated to their quality, then as long as some products are distinctly better than others, everyone can enjoy “the best” for essentially the same price as an average product. Thus, one ought to expect that success in cultural markets will be highly skewed at the top end (Rosen, 1981).

What is more puzzling, however, is that if these “superstar” products—The Beatles, Harry Potter, *Star Wars*—are somehow different from all the other products in ways that are obvious to audiences, one might expect that these same differences would be evident to the producers, publishers, and executives responsible for deciding which projects to support and which to reject. Yet along with inequality of success, a second stylized feature of cultural markets is that they appear to be highly unpredictable, in the sense that attempts to predict hits routinely fail. For example, eight publishers rejected the first Harry Potter book before it went on to sell millions of copies around the world and launch a billion dollar franchise (Lawless, 2005). Executives at the TV network FOX must have bitterly regretted passing on the sitcom *Friends*, but they reveled in their decision to broadcast the mega-hit *American*

Idol after it was rejected by executives at NBC, CBS, and ABC (Carter, 2006). Even some of the most successful cultural products of the 20th century had these troubles; The Beatles struggled to get a record deal (Coleman, 1989), and *Star Wars* almost never got made (Seabrook, 2000).

Just as there are products that are initially rejected, only to end up wildly successful, there are also products that are expected to have wild success that flop (although these are harder to document because they are rarely written about). One such example is Carly Hennessy, an aspiring pop star who impressed MCA Records so much that they invested \$2.2 million in the production and marketing of her first album, *Ultimate High*. Yet, once the album was released, it sold just 378 copies in its first 3 months (Ordonez, 2002). Stories similar to Hennessy's are not uncommon; failures greatly outnumber successes. It is estimated that major record labels lose money on 90% of the artists they sign, and similar, but slightly lower, failure rates are estimated for other cultural products (Vogel, 2004).¹ Ethnographic studies of decision making in cultural industries find that, not surprisingly, industry insiders are very aware of this situation (Bielby & Bielby, 1994; Denisoff, 1975; Faulkner & Anderson, 1987; Gitlin, 1983; Hirsch, 1972), and their feelings are nicely summarized by the well-known quip of the screenwriter William Goldman (1983): "nobody knows anything."

The puzzle we address in this study therefore is the following: If hits are different in some way, why do experts have such difficulty in identifying these products ahead of time?² Rather than looking for an explanation of this puzzle at the level of the individual—either industry executives or consumers—we will argue that the inequality and unpredictability of success, both group-level properties, arise from a process of social influence at the individual level. That is, rather than assuming exogenous and stable preferences, as is common in rational choice models of human behavior and models cultural markets (e.g., Rosen, 1981), our approach is more consistent with a great deal of recent work in experimental psychology suggesting that preferences are neither exogenous nor stable, but are in fact "constructed" by a variety of features of the decision context itself (Bowles, 1998; Payne, Bettman, & Johnson, 1992; Payne, Bettman, & Schkade, 1999; Slovic, 1995). Our explanation, therefore, takes the constructed preferences view of human psychology seriously and adds to it a social component; that is, we propose that expressed preferences are influenced by the observed actions of others, as well as by psychological features of the decision context, such as framing, anchoring, and availability (Ariely, 2008).

3. Social influence and cumulative advantage

The interdependence of decision making and, more specifically, the processes of social influence and conformity have been studied extensively by psychologists, economists, and sociologists for at least 70 years (Asch, 1956; Coleman, Katz, & Menzel, 1957; Crutchfield, 1955; Deutsch & Gerard, 1955; Katz & Lazarsfeld, 1955; Kelley & Thibaut, 1978; Leibenstein, 1950; Luce & Raiffia, 1957; Sherif, 1936; Thibaut & Kelley, 1959). Research in this area has led to two broad insights that are relevant to social influence in cultural

markets. First, in cultural markets, there are simply too many products for individuals to consider—for example, in 1 year alone, about 6,500 new major-label albums are released (Vogel, 2004). As no one can possibly consider all of these products, a natural heuristic for dealing with this choice overload is to assume that the popularity of products is somehow a signal of their quality, a phenomena sometimes called observation learning (Banerjee, 1992; Bikhchandani et al., 1992; Goldstein & Gigerenzer, 2002; Hedström, 1998). Second, in cultural markets, people may benefit from coordinating their choices with others (Adler, 1985); that is, by listening to, reading, or watching the same things, friends, and even strangers, construct common points of interest around which they can interact, thereby fostering notions of commonality and community. Although differentiating between these two kinds of effects can be difficult in practice (Cialdini & Goldstein, 2004; Deutsch & Gerard, 1955; Moretti, 2008), for the purpose of understanding collective dynamics, it may be sufficient to note that both correspond to an individual-level tendency to imitate the behavior of others. From this perspective, popular products will tend to become more popular, leading to “cumulative advantage” or “rich-gets-richer” dynamics, as recent experiments have demonstrated (Hanson & Putler, 1996; Salganik et al., 2006).

Based on simple assumptions about these cumulative advantage processes, a range of mathematical and simulation models have been proposed that produce (a) highly skewed distributions of outcomes (Barabási & Albert, 1999; DiPrete & Eirich, 2006; Simon, 1955; de Solla Price, 1976; Yule, 1925); and (b) unpredictability due to path dependence effects in which initially indistinguishable states can “lock in” to very different equilibria (Arthur, 1994). Results from these models suggest that the puzzling nature of success in cultural markets could follow naturally from the process of social influence at the individual level. Testing this claim empirically, however, has proven more difficult. One difficulty is the sheer proliferation of such models, each of which makes its own assumptions, and yields somewhat different predictions about collective behavior (Adler, 1985; Arthur, 1994; Banerjee, 1992; Bentley, Lipo, Herzog, & Hahn, 2007; Bikhchandani et al., 1992; Chung & Cox, 1994; De Vany & Lee, 2001; De Vany & Walls, 1996; Dodds & Watts, 2004; Granovetter, 1978; Watts, 2002). Further, and more fundamentally, even with the tremendously accurate and fine-grained sales data that are collected by market research firms, it is hard to convincingly measure the quality of products and therefore hard to convincingly argue that some other outcomes could have happened.

Imagine, for example, trying to convince a Harry Potter–obsessed friend that the book’s success was the results of a cumulative advantage process and that the book could have just as easily been a flop (as was predicted by the eight publishers who passed on it). That fan could easily counter that the success of Harry Potter had nothing to do with luck, but stemmed directly from its attributes, which although not what experts in children’s book publishing had anticipated, must have been “what people wanted” (Tucker, 1999).³ One hypothetical way to resolve this disagreement might be to “rewind” the world to the day that Harry Potter was first released and “rerun history” a number of times: If Harry Potter is popular in all of these versions of the world, that would be a strong support for the claim that its success stems from its intrinsic attributes; but if it became popular in one of these worlds and not the others, then we could suspect that its success was determined in part by

chance. In other words, in order to understand the role of chance, we need to observe multiple realizations of the same process. Clearly no such experiment could ever be implemented in existing markets, such as the one that produced Harry Potter. However, we shall argue that it is possible to implement this thought experiment, albeit in a highly simplified form, in an artificial cultural market that we have constructed.

4. Experimental design

The basic framework for our experiment was an artificial cultural “market” created in the form of a website where participants could listen to, rate, and download 48 songs by unknown bands.⁴ Unbeknownst to the participants, we used the “multiple-worlds” experimental design⁵ presented in Fig. 1. Upon arrival to the website, participants were randomly assigned into either the independent condition, where they had no information about the behavior of others, or the social influence condition, within which they were also assigned to a specific “world” where they had information about the behavior of those in their world, but not those in the other worlds. Because all the worlds started with the same initial conditions, had indistinguishable participants, and the same 48 songs, they represent multiple parallel realization of the same process.

The participants, although not aware of these multiple worlds, were told that they were participating in a study of music tastes and that, in exchange for participating, they would be offered a chance to download new music by up-and-coming artists.⁶ After providing their informed consent,⁷ participants completed a brief survey that collected simple demographic information such as age, gender, country of residence, etc. Next, participants were sent to a page of instructions and then, finally, were presented with a song menu (Fig. 2), similar to the layout of other online music downloading sites. For participants in the social influence worlds, the band names and song names were accompanied by the number of previous downloads in their world, whereas for participants in the independent condition no such information was provided. Participants could click on a song at which point it began to play. While the song was playing, participants were asked to rate the song on a scale of 1 star (“I hate it”) to 5 stars (“I love it”). After rating the song, participants were offered the chance to download it, and then were taken back to the song menu and were able to listen to




Fig. 1. Schematic view of the experiment design used in these experiments. This design allows researchers to observe the effects of the individual process of social influence on collective outcomes.

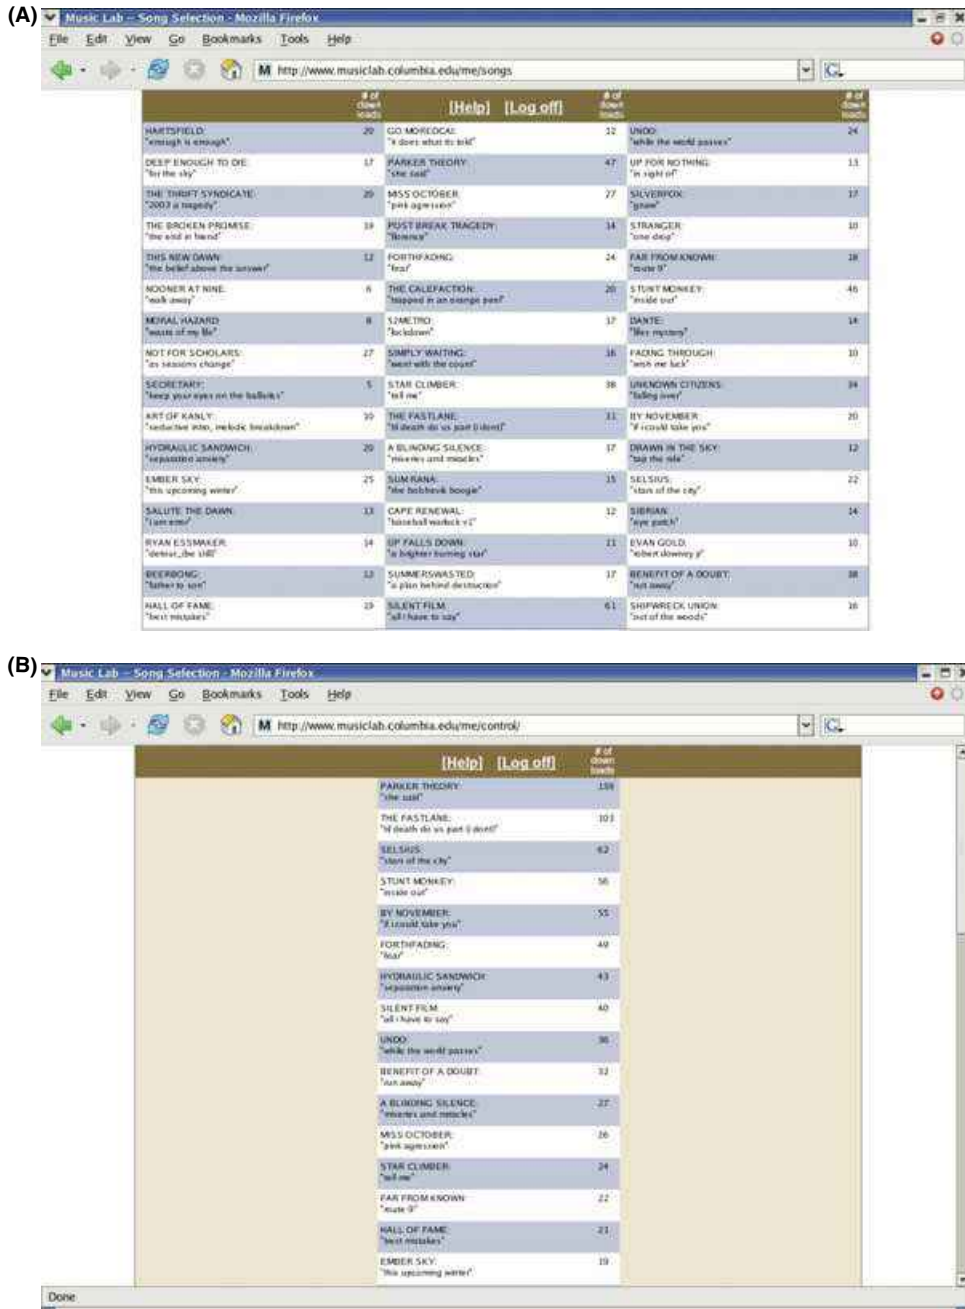


Fig. 2. Screenshots of the song selection screen from Experiment 1 (A) and Experiments 2 and 3 (B). In Experiment 1 the placement of songs was random, but that in Experiments 2 and 3 the songs were sorted by popularity. Screenshots from the independent condition were identical, except the songs were always ordered randomly and the download counts were not displayed.

other songs. The download counts presented to the participants were updated accurately and in real time.⁸ Because of the design of our website participants could not download songs without listening to them, but they could listen to, rate, and download as many songs as they wished.

5. Previous results

Using the same experimental set-up, Salganik et al. (2006) conducted two successive experiments that varied in the amount of social influence to which individual decision makers were exposed. These experiments took place from October 7, 2004 to December 15, 2004 (69 days) and December 15, 2004 to March 8, 2005 (83 days), and collectively involved a total of 14,341 participants. The participants were mainly recruited from the now defunct, teen-interest website, <http://www.bolt.com>, using a combination of text ads and banner ads. Consistent with Bolt's demographics, most participants were teenagers living in the United States (Table 1).⁹ In each experiment, participants arriving to our website were randomly assigned to one of eight social influence worlds or the one independent world, such that the independent world received twice as many participants as the social influence worlds (the reason for this difference will become clear when we describe our measure of unpredictability). Given a total of 14,341 participants over the two experiments, therefore, each social influence world had about 700 participants, while each independent world had about 1,400.

In Experiment 1, songs were presented in a 16×3 grid not sorted by popularity, while in Experiment 2, the songs were presented in a single-column format sorted by popularity (Fig. 2). In the independent condition of both experiments, the songs were presented in the same format, but without any information about popularity and in a random order. Thus, by

Table 1
Descriptive statistics of participants in Experiments 1, 2, and 3

Category	Experiment 1 (<i>n</i> = 7,149) (% of Participants)	Experiment 2 (<i>n</i> = 7,192) (% of Participants)	Experiment 3 (<i>n</i> = 2,930) (% of Participants)
Female	36.4	73.9	38.0
Broadband connection	74.1	69.0	90.6
Has downloaded music from other sites	60.4	62.4	69.3
Country of residence			
United States	79.8	81.8	68.4
Canada	4.5	4.4	6.3
UK	4.4	4.7	6.6
Other	11.0	9.1	18.7
Age (years)			
17 and younger	11.5	16.0	1.5
18–24	27.8	34.9	5.7
25–34	38.5	39.2	29.8
35 and older	22.3	9.9	63.1

comparing the results within each experiment, we could observe the effects of individual social influence on collective outcomes and, by comparing across experiments, we could observe the effect of increasing the “strength” of this individual-level process (i.e., by making the social information more salient to the users).

Before considering our main question about inequality and unpredictability of success, we first measure whether the information about the popularity of the songs affected participants’ behavior at all, and whether this effect was larger in Experiment 2, where the information was made more salient. To construct such a measure, we note that at the time that each subject participated, every song in his/her world had a specific market share and market rank (e.g., the song with the highest market share had a market rank of 1). Fig. 3 plots the probability that a participant listened to a song of a given rank (independent of which song occupied that rank at that time) in Experiments 1 and 2. Fig. 3A shows that, in Experiment 1, when the songs were not sorted by popularity, participants in the social influence worlds had about a 20% chance of listening to the most popular song, but only about a 7% chance of listening to a song of average popularity (solid circles). In the independent world (open circles), the probability of listening was independent of market rank, which is as expected given our design. Fig. 3B shows the same results from Experiment 2 and reveals that, as expected, sorting the songs by popularity caused participants to be more affected by market rank. For example, in Experiment 2, participants had almost a 50% chance of listening to the most popular song, but only about a 5% chance of listening to a song of middle market rank.¹⁰

These results therefore confirm that the popularity of the songs affected participants’ choices and generally led them to listen to the more popular songs—a result that is consistent with the large literature on social influence and conformity. Now, we turn our attention to measuring the consequences of this individual-level social influence on collective outcomes, particularly the inequality and unpredictability of success. To begin with, we define the success of a song based on its market share of downloads,

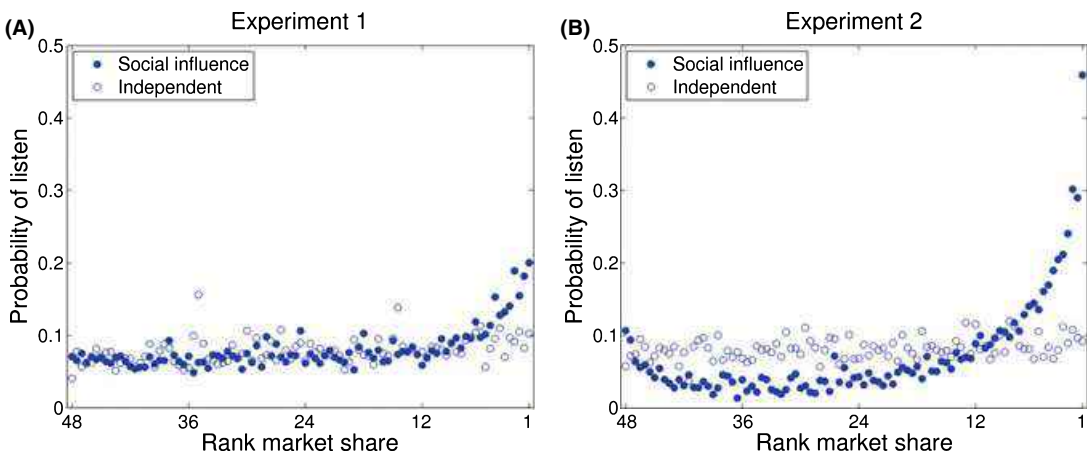


Fig. 3. Probability of listening to a song of a given market rank in Experiment 1(A) and Experiment 2(B). Participants in Experiment 2 were more likely to listen to more popular songs.

$$m_i = \frac{d_i}{\sum d}$$

We then measure the inequality of success in terms of the Gini coefficient, a common measure of inequality that is defined as

$$G = \frac{\frac{1}{S^2} \sum_{i=1}^S \sum_{j=1}^S |m_i - m_j|}{2 \cdot \sum \frac{m}{S}}$$

where m_i is the market share of song i and S is the number of songs (Allison, 1978; Coulter, 1989).¹¹ The Gini coefficient can be thought of as the expected difference in market share between two randomly chosen songs that has been scaled to fall between 0 (complete equality) and 1 (complete inequality).

Fig. 4A shows that, in Experiment 1, success was more unequal in the social influence worlds than the independent world. To assess whether differences this large could have occurred by chance (i.e., to assess statistical significance), we used a modified bootstrap procedure (Efron & Tibshirani, 1993); that is, we randomly resampled the participants (with replacement) from the independent condition to produce 1,000 replicate independent conditions that were similar in size to the social influence worlds and calculated the Gini coefficient for each of these replicates. The value presented in Fig. 4A is the mean Gini coefficient of these replicates, and the interval is the range within which 95% of these replicate values fell. The probability that a randomly chosen independent world replicate was more unequal than a randomly chosen social influence world was small ($p < .01$), indicating that the observed differences were statistically significant. Fig. 4B shows that a similar pattern was observed in Experiment 2 with success in the social influence worlds being more unequal than in the independent world ($p < .01$).

Further, the between-experiment comparison shows that in Experiment 1, when the songs were randomly ordered, the inequality in the social influence worlds ranged from 0.28 to 0.42, and that, in Experiment 2, when the songs were sorted by popularity, it ranged from 0.45 to 0.56. In other words, every social influence world in Experiment 2 was more unequal than the most unequal in Experiment 1.¹² This difference was not only statistically significant ($p < .01$) but also substantively large—by comparison, the magnitude of this difference is similar to the difference in income inequality between Western European countries such as France, Spain, and the Netherlands ($G \sim 0.34$) and developing countries such as Nigeria, Peru, and Venezuela ($G \sim 0.50$) (UNDP, 2004). Finally, Figs. 4C and 4D plot the dynamics of the Gini coefficient as the experiments progressed and show that the approach to steady state was qualitatively similar in both experiments. Further, they show that we would not anticipate very different results had the experiment continued.¹³

Next we turn our attention to the unpredictability of success. Recall that the cumulative advantage models predicted that we will observe different outcomes across realizations of exactly the same process (Arthur, 1994). To quantify this, we define the unpredictability of

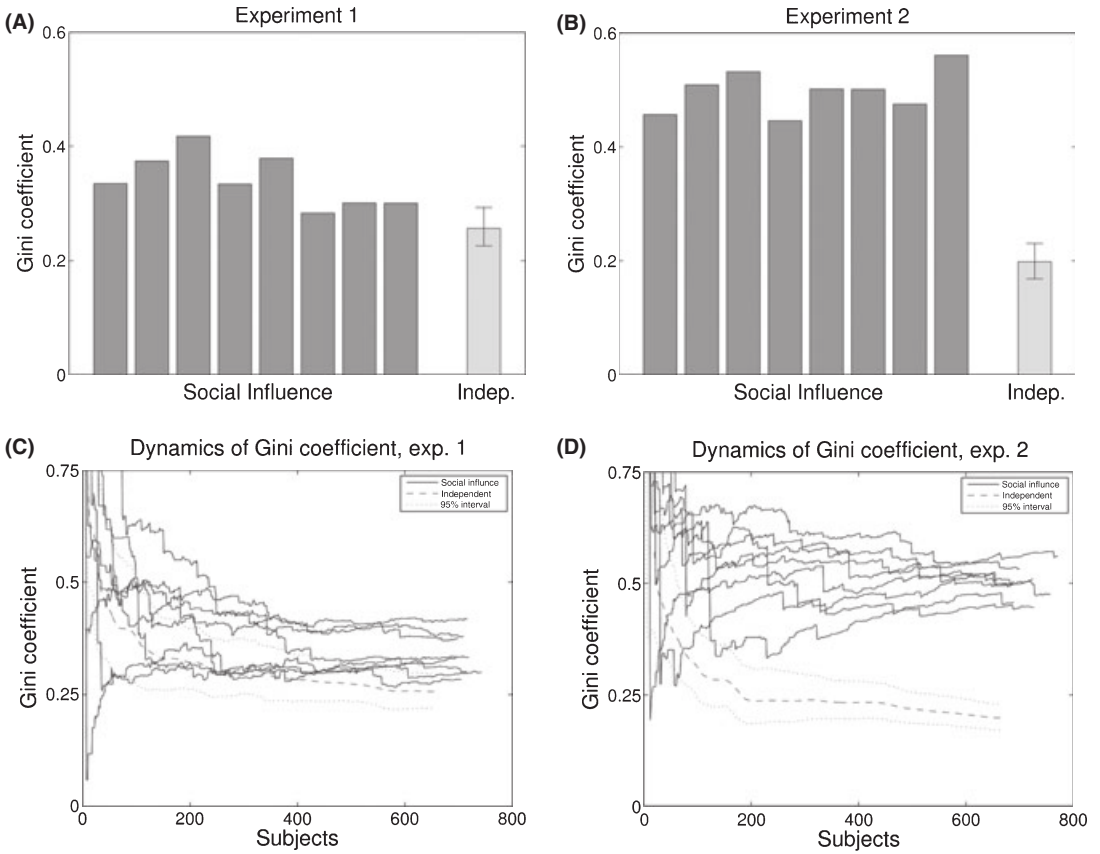


Fig. 4. Inequality of success in Experiment 1 (A and C) and Experiment 2 (B and D) as measured by the Gini coefficient. In both experiments, success in the social influence worlds was more unequal than success in the independent world. Also, success in the social influence worlds in Experiment 2 was more unequal than success in the social influence worlds in Experiment 1.

specific song, u_i , to be the average difference in market share between two randomly chosen success outcomes. That is,

$$u_i = \frac{\sum_{i=1}^R \sum_{j=i+1}^R |m_{ij} - m_{ik}|}{\binom{R}{2}}$$

where m_{ij} is the market share of song i in world j and R is the number of realizations. If a song earns the same market share in all realizations, then its unpredictability is defined to be zero. Note that zero unpredictability by this measure does not mean that in practice it would be easy to predict the song’s success, but since the song has the same outcome in all realizations, precise prediction cannot be ruled out in principle. To the extent that a song experiences different levels of success in different realizations, however, these differences

represent *inherent* unpredictability that cannot be eliminated any more than one can eliminate the uncertainty in the roll of a die. The unpredictability of an experimental condition was defined to be the average unpredictability of the songs in that condition, $U = \sum_S u_i/S$.

Recall, however, that there was only one realization of an independent world and so it seems impossible to calculate unpredictability, which was defined based on differences across multiple realizations of the same process. It was for this reason that the independent condition had twice as many participants as each social influence world. In order to calculate the unpredictability for the independent condition, we randomly split the independent world of about 1,400 participants into two independent worlds of about 700 participants each. The unpredictability in the independent condition was then calculated by comparing outcomes in these split worlds, and this procedure was repeated 1,000 times.

Figs. 5A and 5B show that the unpredictability of success was greater in the social influence condition than the independent condition. To calculate a measure of statistical

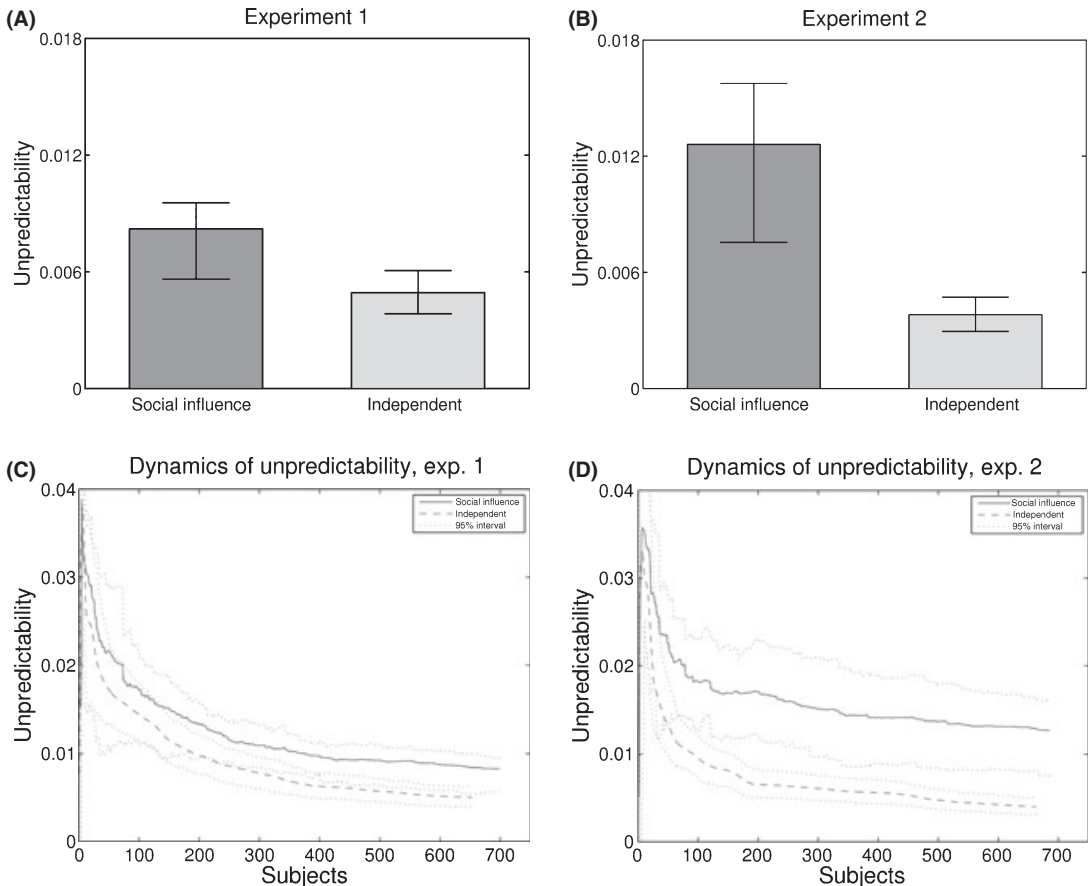


Fig. 5. Unpredictability of success in Experiment 1 (A and C) and Experiment 2 (B and D). In both experiments, success in the social influence worlds was more unpredictable than success in the independent world. Also, success in the social influence worlds in Experiment 2 was more unpredictable than success in the social influence worlds in Experiment 1.

significance, we compared the distribution of replicate values from the independent condition to the distribution of calculated unpredictability values for the 28 ($8 \times 7/2$) possible pairs of social influence worlds (the 95% intervals for these distributions are shown in the figure). The probability that the unpredictability of a randomly chosen pair of social influence worlds was greater than the unpredictability of a random split of the independent condition was small ($p < .01$) in both Experiments 1 and 2. Thus, we conclude the difference in unpredictability across the conditions is statistically significant. Moreover, comparing the values of our unpredictability measure between Experiments 1 and 2, we also find that increasing the amount of social influence increased the unpredictability of success by about 50%. We note that although a difference this large is unlikely, it is not statistically significant by common standards ($p \approx .06$). Finally, Figs. 5C and 5D again show that we would not have likely seen different results if the experiments had continued and that the approach to steady state was qualitatively similar in both experiments.

One feature that our measure of unpredictability obscures is that the success of the songs was not completely random, but rather, that some songs consistently did better than others. An advantage of our experimental design is that it allows us to assess whether that might be because some of the songs were more appealing than others. In general, the “quality” of cultural products is notoriously difficult—some might say impossible—to measure objectively (Becker, 1982; Bourdieu, 1984; Cutting, 2003; DiMaggio, 1987; Gans, 1974); thus, we instead refer to the “appeal” of a song, which we define in terms of its popularity in the independent condition. We note that what we are calling “appeal” is not necessarily the same as quality. Nevertheless, because the behavior of users in the independent condition is, by definition, unaffected by the behavior of others, it provides a natural measure of preexisting population preferences—something that has also proven difficult to measure convincingly in previous studies of cultural markets (Hamlen, 1991; Krueger, 2005). More precisely, we measure the appeal of the songs to our pool of participants by the market share of downloads in the independent world:

$$a_i = \frac{d_i}{\sum d}$$

where d_i is the number of downloads for song i in the independent condition (Salganik & Watts, 2008; Salganik et al., 2006). Fig. 6 shows that higher appeal songs tended to do better, but there is substantial scatter in the relationship, particularly for the highest appeal songs. In other words, outcomes were more unpredictable for higher appeal songs than for lower appeal songs, suggesting that it is easier to predict failure than to predict success.

Further, if one were to average the success across the eight social influence worlds, this average success is strongly associated with the underlying appeal; however, it is nonlinear in the sense that small changes in appeal at the top of the distribution are associated with larger changes in average success, especially in Experiment 2.¹⁴ Thus, one way to consider these results is that the outcome of each particular social influence world contains some signal relating to the underlying appeal of the songs and some noise that results from the interdependence of decision making. When the social influence is stronger, as in Experiment 2, that noise appears to be greater. We speculate that if one could generate thousands of

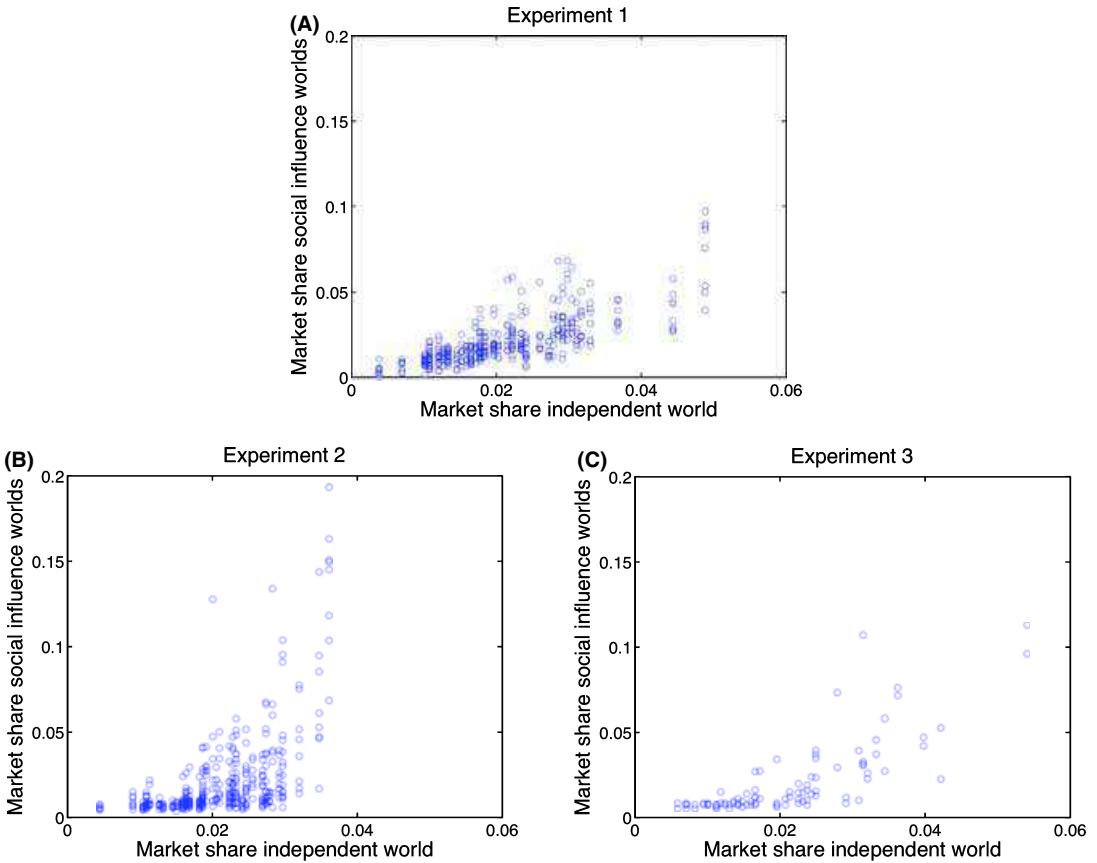


Fig. 6. The relationship between appeal, as measured by market share in the independent condition, and success, as measured by market share in the social influence worlds in Experiments 1(A), 2(B), and 3(C).

social influence worlds and average their outcomes that this average success would be strongly related to the underlying appeal of the products. However, in practice, we cannot observe these parallel worlds; thus, the ratio of signal to noise in the outcomes in real cultural markets is an open question.

Overall, the results from Experiments 1 and 2 provided strong support for the argument that social influence at the individual level is simultaneously responsible for increased inequality and unpredictability in collective outcomes—in this case, the distribution of market share. Although simple to state, this finding nevertheless exhibits a curious paradox: On the one hand, by revealing the existing popularity of songs to individuals, the market provides them with real, and often useful, information; but on the other hand, if they actually use this information, the market inevitably aggregates less useful information. This result, which is analogous to “information cascades” in economics (Banerjee, 1992; Bikhchandani et al., 1992), suggests, in turn, that social institutions that make us aware of the behavior of others—the *New York Times* bestseller list, the *Billboard* album charts, and lists of

top-grossing movies—do provide a useful service to individuals, but only at the cost of increasing the overall inequality and unpredictability of the markets themselves.

6. A new pool of participants

There are risks, however, in drawing conclusions from studies that were based mostly on the behavior of American teenagers, a population that may be unusual in many ways. For example, there is some evidence that younger people may be more susceptible to social influence (Park & Lessig, 1977; Pasupathi, 1999) and that conformity levels vary by country (Bond & Smith, 1996). For this reason, we now investigate how the individual and collective-levels result that we observed in Experiment 2 would change if we used a different group of participants. Further, by observing the same process with two distinct pools of participants, we were able to make additional comparisons that were not possible with just one pool of participants.

To recruit this different set of participants, we sent 13,546 e-mails to subjects from a previous, unrelated Web-based study (Dodds, Muhamad, & Watts, 2003). Given the nature of Dodds et al.'s (2003) study, we expected these participants to be quite different than the participants in Experiments 1 and 2. From March 14, 2005 to April 7, 2005 (24 days), these e-mails, plus the Web-postings they generated, yielded a pool of 2,930 participants that were assigned into two social influence worlds and one independent world, corresponding to a similar number of participants per world as in Experiments 1 and 2.¹⁵

Most important for the goals here, these participants differed from those in Experiment 2 in a number of ways, as we had expected. First, the two groups were different demographically with participants in Experiment 3 being older, more male, and more international (Table 1). Second, the participants differed in their behavior—in particular, participants in Experiment 3 listened to twice as many songs, on average, as those in Experiment 2, a difference we had not anticipated (Table 2). Finally, participants in Experiment 3 had different musical preferences than those in Experiment 2. Fig. 7A compares the appeal measures from Experiment 2 and 3 and shows that the different pools of participants had slightly different musical preferences

Table 2
Descriptive statistics of participant behavior in Experiments 2 and 3

	Experiment 2			Experiment 3		
	Social Inf. 8 Worlds (<i>n</i> = 5,746)	Indep. 1 World (<i>n</i> = 1,446)	Total 9 Worlds (<i>n</i> = 7,192)	Social Inf. 2 Worlds (<i>n</i> = 1,471)	Indep. 1 World (<i>n</i> = 1,459)	Total 3 Worlds (<i>n</i> = 2,930)
Listens	20,217	5,643	25,860	10,591	11,844	22,435
Listens per participant	3.5	3.9	3.6	7.2	8.1	7.7
Downloads	8,106	2,192	10,298	2,040	1,691	3,731
Downloads per participant	1.4	1.5	1.4	1.4	1.2	1.3
Pr[download listen]	0.401	0.377	0.398	0.193	0.143	0.166
Mean rating (no. stars)	3.31	3.22	3.29	2.70	2.55	2.62

Note: Social Inf., social influence; Indep., independent.

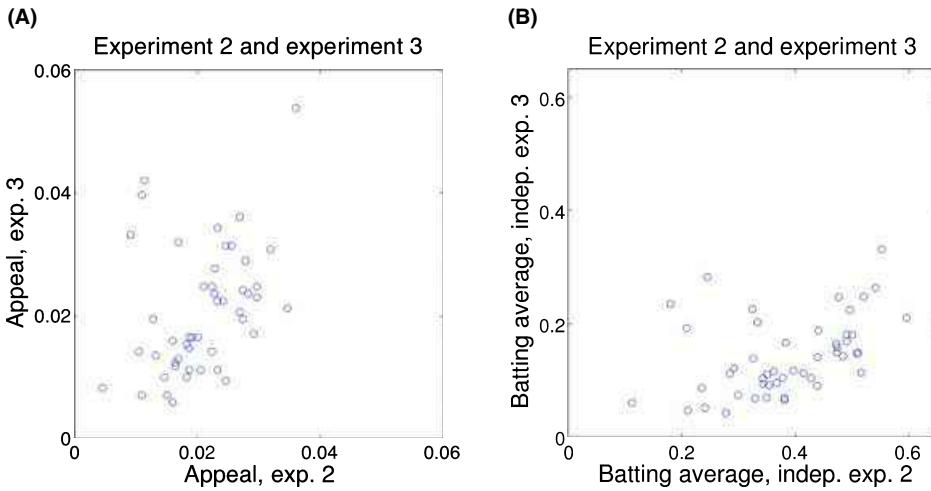


Fig. 7. Comparing the preferences of the participants in Experiments 2 and 3 as measured by appeal (A) and batting average (B).

($r = .37$). Further, participants in Experiment 3 liked the songs less on average than the participants in Experiment 2 as can be seen in Fig. 7B, which compares the probability that a listen for a specific song will lead to a download—what Aizen, Huttenlocher, Kleinberg, and Novak (2004) call the “batting average.” Overall, 46 of the 48 songs had lower batting averages in Experiment 3,¹⁶ a finding that is consistent with the lower song ratings that were reported in Experiment 3 (mean of 2.6 stars) compared with Experiment 2 (mean of 3.3 stars) (Table 2).¹⁷ Thus, we conclude that, as expected, the participants recruited for Experiment 3 differed along at least three important dimensions: demographics, behavior, and musical preferences.

In spite of these differences, participants from Experiment 3 were influenced by the behavior of others in a similar way to participants in Experiment 2. Fig. 8 plots the probability that a participant listened to a song of a given rank (independent of which song occupied that rank at that time). Participants in Experiment 3 had higher listening rates overall, and thus all the effects observed in Experiment 2 were more pronounced in Experiment 3. In qualitative terms, however, the impact of social influence, as measured by the difference between independent (open circles) and social influence (closed circles) conditions, was similar in both experiments.

Fig. 9 shows that the results at the collective level were also qualitatively similar. For example, in both Experiments 2 (Fig. 9A) and 3 (Fig. 9B), the inequality of success in the social influence worlds was greater than the inequality in the independent world ($p < .01$). Further, a comparison across experiments shows the amount of inequality in the social influence worlds was similar to the amount of the inequality in the social influence worlds in Experiment 2. The independent worlds, however, appear to have different levels of inequality: $G = 0.19$ in Experiment 2 and $G = 0.27$ in Experiment 3. We suspect that this increased inequality in the independent condition can be attributed to the greater variation in the probability of download given listen in Experiment 3 than in Experiment 2 (Fig. 6B). Figs. 9C

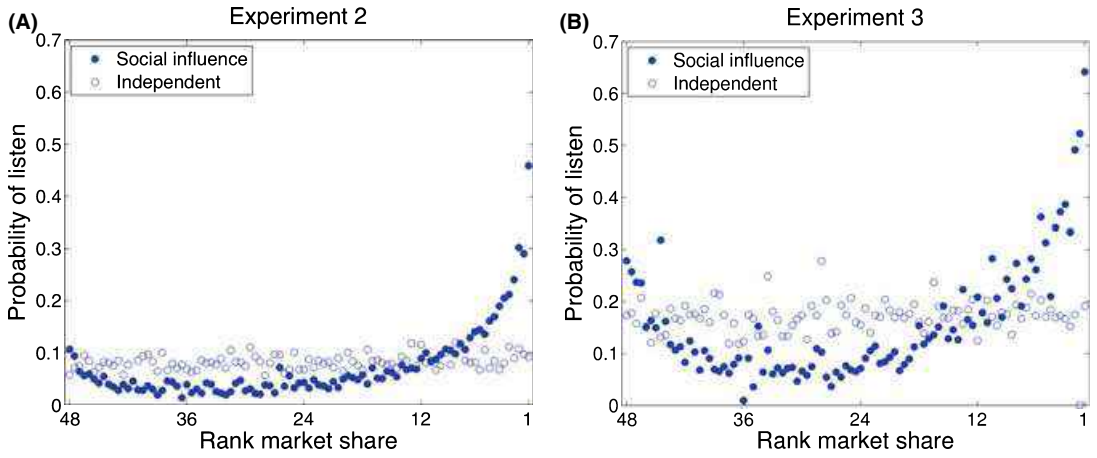


Fig. 8. Probability of listening to a song of a given market rank in Experiments 2 (A) and 3 (B). The patterns in Experiments 2 and 3 were qualitatively similar.

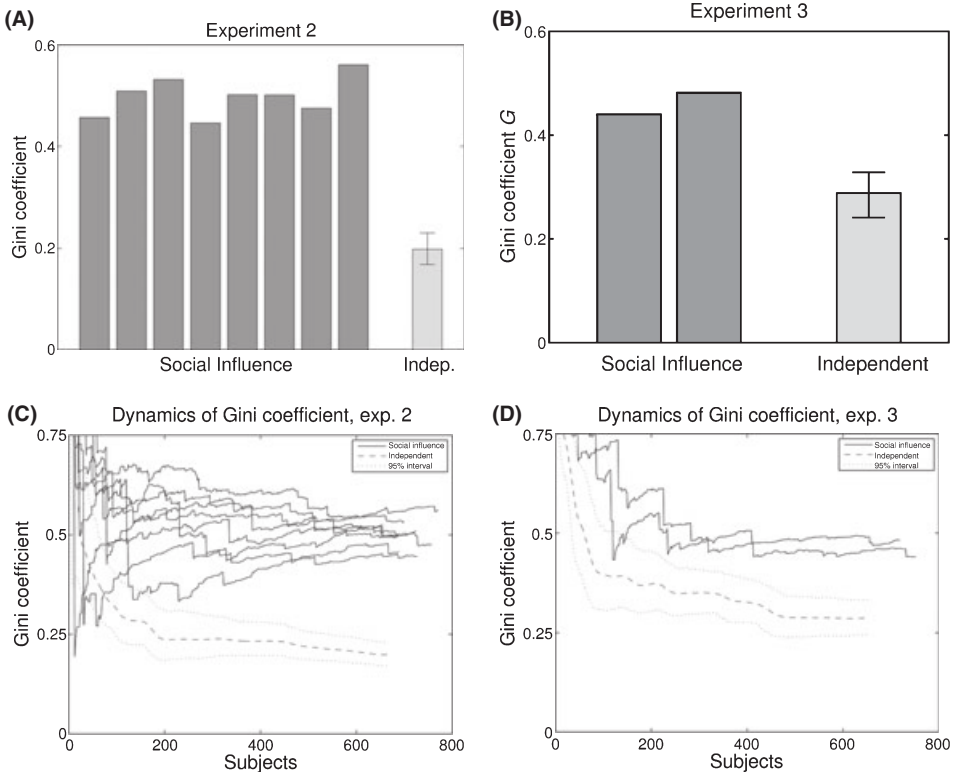


Fig. 9. Inequality of success in Experiments 2 (A and C) and 3 (B and D) as measured by the Gini coefficient. The inequality of success in Experiment 3 was similar to the inequality of success in Experiment 2.

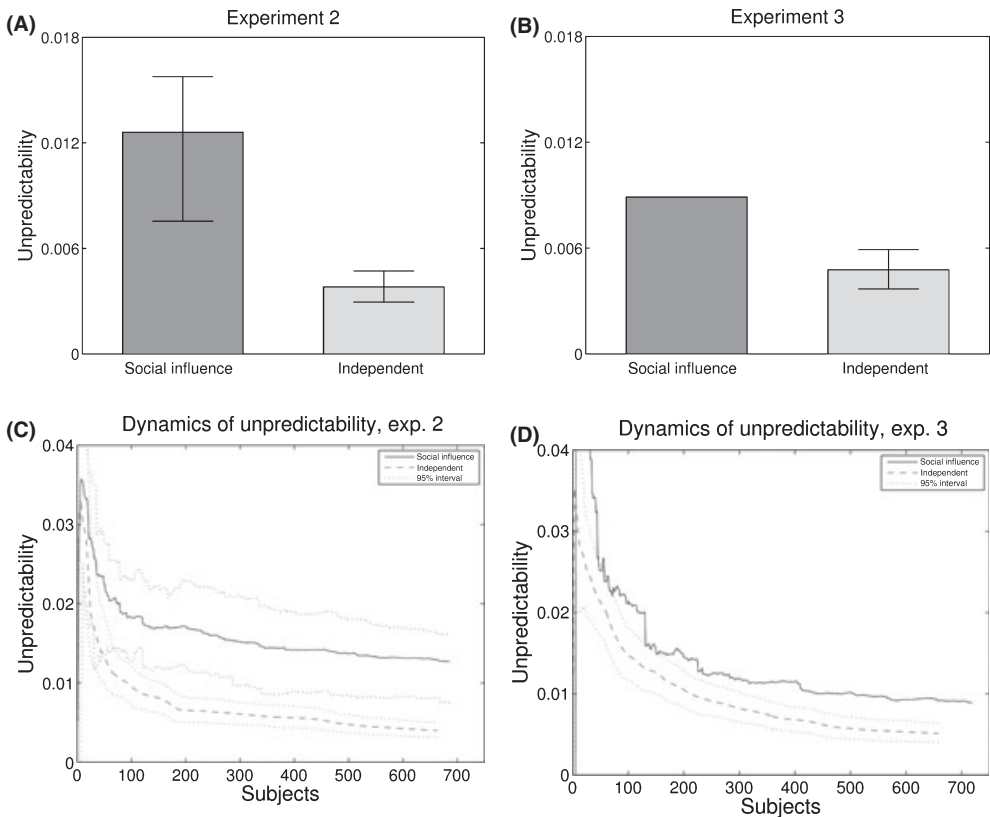


Fig. 10. Unpredictability of success in Experiments 2 (A and C) and 3 (B and D). The unpredictability of success in Experiment 3 was similar to the unpredictability of success in Experiment 2.

and 9D show that the full dynamic trajectories—not just the steady states—were similar across experiments.

Considering now the unpredictability of success, Fig. 10 shows that in both Experiments 2 (Fig. 10A) and 3 (Fig. 10B), the unpredictability of success was greater in social influence condition than the independent condition ($p < .01$). A comparison across experiments, however, shows that the unpredictability in the social influence condition seems to have decreased from Experiments 2 to 3. We suspect that this difference could be because participants in the social influence condition in Experiment 3 listened to so many more songs than those in Experiment 2: 7.2 compared with 3.5, respectively. This extensive listening behavior ensured that all songs had a reasonable number of listens and therefore prevented the high appeal songs from escaping notice. Finally, Figs. 10C and 10D show that—as with the inequality measure above—the dynamics of the unpredictability measure were qualitatively similar in both experiments.

Collectively, these three experiments illustrate an important point made by the sociologist Stanley Lieberman (1997): When considering interventions to a social system, it is important to distinguish between changes that will affect the distribution of outcomes while preserving

relative positions, and changes that, by contrast, will affect relative positions while keeping the distribution of outcomes unchanged. For example, although the aggregate features of Experiments 2 and 3 were similar, the results for specific songs were different (as might be expected given the differences in preferences). Consider “Life’s Mystery,” an acoustic ballad by the singer Dante; this song did poorly with the mostly teenage participants in Experiment 2—it was in the bottom half of popularity in six of eight influence worlds. However, in Experiment 3, with an older, more international population of participants, “Life’s Mystery” fared much better, coming in sixth and fifth in the two influence worlds. In general, if we calculate \bar{K}_i , the average market rank of song i in all of the social influence worlds in an experiment, the correlation between this measure in Experiments 2 and 3 was 0.52. Thus, similarity of results at the aggregate level between Experiments 2 and 3 was coupled with some change in the ordering of the songs.

A comparison between Experiments 1 and 2 reveals a different pattern. In these experiments, there was little change at the song level; the correlation between average market rank in the social influence worlds of Experiments 1 and 2 was 0.93.¹⁸ In other words, songs that did well in Experiment 1 tended to do well in Experiment 2, as one might suspect given that participants in both experiments were similar. However, changing the song menu led to large changes at the aggregate level with the Gini coefficient increasing from $G \sim 0.34$ to $G \sim 0.50$. Together, therefore, these experiments show that one of our interventions, increasing the salience of the popularity of the songs (Experiments 1 and 2), led to stability in the ordering of the songs coupled with changes in the aggregate features of the market; whereas our other intervention, changing the pool of participants (Experiments 2 and 3), led to stability in the aggregate features of the market coupled with changes in the ordering of the songs.

7. Predicting success

The relationship between appeal and success shown in Fig. 6 suggests that one might reasonably think of outcomes that are observed in the social world as arising out of a combination of both individual preferences and aggregation dynamics. Predicting outcomes, therefore, requires one to understand not only the distribution of individual preferences in the target population but also to intuit how all those preferences will aggregate when everyone is also paying attention to what other people are doing. We have already argued that it is this second component to success that is inherently unpredictable; but our comparison of Experiments 2 and 3 raises a related question—namely, how much can one learn from the results of one experiment that could be applied to another? This question has a practical analog in business. For example, it is often the case that an executive contemplating the launch of an object X (e.g., a book or a TV show) in market B can observe the success of the same object in some other market A in which it was launched previously. Given that the two markets are presumably characterized by different distributions of preferences, the question, then, is how much weight should the executive place on X ’s success in market A relative to his/her own premarket audience testing in market B ?

Our experiment was not designed to answer this particular question, and so it can offer no definitive answer. Nevertheless, we can examine our data in order to gain a qualitative sense of the relative predictive value of these two kinds of information in our experiment. We begin by calculating the prediction error,

$$\varepsilon = \sum_{s=1}^{48} |m_{sw} - \hat{m}_{sw}|$$

defined as the total (over all songs) difference between the predicted and actual market shares, where m_{sw} is the market share of song s in world w , and \hat{m}_{sw} is the predicted market share of song s in world w . In this case, the predicted market share will be defined to be the market share that has been observed; in the hypothetical example, this could be the market share in either market A or in premarket testing in market B. Unfortunately, ε is hard to interpret directly, ranging from a minimum of 0, if the prediction is perfectly accurate, to a maximum of 2. To improve understanding, therefore, we can give these results a proportional reduction in error interpretation by comparing the prediction using our method and a given piece of information to the prediction that would be made using no information. In this case, the prediction with no information would be that each song would earn the average market share,

$$\varepsilon_{\text{naive}} = \sum_{s=1}^{48} |m_{sw} - 1/48|$$

By comparing the errors with different prediction methods, we can give our predictor a proportional reduction in error interpretation (Costner, 1965),

$$\text{PRE} = \frac{\varepsilon_{\text{naive}} - \varepsilon}{\varepsilon_{\text{naive}}}$$

This measure varies from $-\infty$ to 1, where a negative value means that naive predictions are better than our predictions, a value of 0 means that both predictions have equal error, and a value of 1 means that our predictions have no error. We can illustrate the measure with our hypothetical example of a business executive who wishes to predict market share in a social influence world in Experiment 2.¹⁹ Using the naive prediction, his/her error would be 0.77;²⁰ however, if he/she also has access to extensive premarket testing (analogous to knowing the distribution of popularity in the Experiment 2 independent condition), his/her expected error is 0.65. The expected error, in other words, has been reduced by about $(0.77 - 0.65)/0.77 = 0.16$ or 16%, leaving the remaining 84% of the error unexplained. Even knowing the full distribution of popularity in the independent condition is therefore of limited value to someone attempting to predict outcomes in the presence of social influence—a result consistent with our general claim that some portion of the unpredictability of market success cannot be eliminated, even with perfect information about the preexisting preferences of individuals.

Having absorbed the importance of social influence on individual decision making, one might suspect that evidence of ‘‘real world’’ success (or failure) in a different market would be more informative than an independent sampling of the audience in question. Yet in fact, the opposite is true: The proportional reduction in error from using outcomes from a social influence world in Experiment 3 to predict success in a social influence world from Experiment 2 is a paltry 2%. Performing the same exercise with Experiment 3, moreover, reveals the same result, but with an even more striking contrast. Whereas knowing the distribution of popularity in the independent condition of Experiment 3 yields a reduction in error of 38% when predicting outcomes in a social influence world in Experiment 3, using a social influence world from Experiment 2 to predict success in a social influence world in Experiment 3 yields an *increase* in error of 3% relative to the naive prediction.²¹ In other words, in this situation, our hypothetical executive would do just as well simply guessing.

That past performance does not guarantee future returns may not seem like an overly surprising result, and no doubt many examples can be found where a book, movie, or song that has been successful in one market had been turned down by executives on the basis that ‘‘it wouldn’t work over here.’’ Indeed, it was presumably on this basis that the mega-hit *American Idol* was rejected by TV executives in the United States in spite of the popularity of its progenitor, *Pop Idol*, in the United Kingdom. On the other hand, following on the American success of the British game show *Who Wants to be a Millionaire?* NBC eagerly aired *The Weakest Link*, another game show popular in the United Kingdom, which went on to flop (Carter, 2006).

In light of the many complexities involved, we wish to avoid placing too much emphasis on these results, which are purely suggestive.²² Nevertheless, the result that one could learn virtually nothing from one market outcome that could be applied in another is, we feel, a surprising result. Indeed, as indicated in Fig. 11 (which summarizes all the comparisons




Fig. 11. Proportional reduction in error in predicting success. Having information about the success in a social influence world in one experiment is essentially useless in predicting success in a social influence world in the other experiment.

above), even the independent conditions from the “wrong” experiment were more informative than the social influence outcomes (compare the diagonal arrows with the vertical arrows on the right-hand side). The results from Fig. 6, however, provide at least some intuition for these findings. If we consider the outcomes of a social influence world to be noisily related to appeal, then it stands to reason that one noisy outcome cannot predict another noisy outcome, especially when the underlying signal itself may be changing. We hope that future modeling or experiments will address these issues in more detail.

8. Discussion and next steps

The experiments presented here provide additional support for the results of Salganik et al. (2006), suggesting that both inequality and unpredictability in cultural markets inevitably arise from an individual-level process of social influence. Using a pool of participants with different demographics, behavior, and preferences, the results of Experiment 3 were qualitatively similar to those from Experiment 2. However, a comparison of the two experiments also reveals that our qualitative understanding of the process is insufficient to explain the differences between Experiments 2 and 3. There are several things that are different about the participants and we are not sure which of these lead to the similarities and differences in collective-level results that we observed. Emphasizing the preliminary nature of our conclusions, therefore, we end by outlining some possible next steps.

One natural next step would be to formulate a mathematical or agent-based model of this process. Agent-based models are useful tools for studying collective behavior, but they are often limited by a lack of empirical foundations at either the individual or collective level (Epstein & Axtell, 1996; Goldstone & Janssen, 2005; Macy & Willer, 2002).²³ We believe that these experimental results are ideal for guiding agent-based modeling because they provide both detailed individual-level data that can be used to choose between alternative microfoundations *and* detailed collective-level patterns that can then be used to validate the model. Once researchers are able to formulate a simulation model that can reproduce the observed collective behavior (both dynamics and steady state), they can then explore how different parameters affect the observed collective dynamics. An example of this approach of linking group experiments to agent-based modeling is the work carried out on foraging patterns by Goldstone and colleagues (Goldstone & Ashpole, 2004; Goldstone et al., 2005; Roberts & Goldstone, 2006). With a fully calibrated model one can ask many questions. How would the results have been different if the social influence was increased? What if the songs were more similar (or different) in terms of their appeal? What if there were 48,000 songs instead of 48? These are all important questions because the experiments presented here represent only three data points in the set of all possible experiments with this design.²⁴ Careful simulation modeling could suggest interesting predictions about other portions of this parameter space that can be tested with further experiments, and some work in this direction has already begun (Borghesi & Bouchaud, 2007).²⁵

In addition to other experiments using similar designs, we suspect that many future experiments in large-group collective dynamics will also take place on the Web. These

experiments will pose specific challenges above and beyond the challenges of Web-based experiments in general (Kraut et al., 2004; Nosek, Banaji, & Greenwald, 2002a; Reips, 2002; Skitka & Sargis, 2006). The experiments performed here allowed for a very limited form of interaction in the sense that participants made one-time decisions based on the prior decisions of others. By contrast, many social situations exhibit more complex interactions where individuals are constantly adjusting their behavior based on the behavior of others. In an experimental situation when this decision making is happening in continuous time, all experimental subjects must be participating at exactly the same time. This demand for simultaneity, which was not present in the experiments described here, introduces considerable implementation challenges. Some of these difficulties can be overcome by placing all subjects physically in the same lab at the same time—for example, in the conformity experiments of Latané and L’Herrou (1996), and the graph coloring experiments of Kearns, Suri, and Montfort (2006)—but then the difficulty of scaling beyond small groups again becomes an obstacle. Lifting the scale constraints, therefore, requires a true Web-based “platform” for conducting experiments, but it is not yet clear how such a platform would handle complex coordination of subjects, participant dropout, and other issues that are relatively easily resolved in the close physical proximity of a physical lab.²⁶

Finally, another limitation to scaling to truly large groups is that most experiments currently run by social scientists are not intrinsically enjoyable for participants; thus, subjects are generally paid in exchange for their participation. Although paying people is certainly reasonable, at just \$8 per subject, the three experiments presented here—which collectively involved 17,271 subjects—would have required about \$138,000 *just for subject compensation*. To run large experiments on the Web, therefore, researchers will need to develop new models of compensation, such as micropayments on “crowd-sourcing” platforms like Amazon’s Mechanical Turk (Mason & Watts, 2009), or, as we have done here, by designing experiments that are intrinsically rewarding for participants (von Ahn & Dabbish, 2004; Dodds et al., 2003; Nosek, Banaji, & Greenwald, 2002b; Srivastava, John, Gosling, & Potter, 2003). It remains to be seen what types of incentive schemes will emerge for Web-based experiments, but social scientists have proven adept at experimental design in the past, and we suspect that there will be a number of clever experiments deployed over the Web in the next few years. Regardless, the Web offers exciting possibilities for the experimental study of collective social dynamics, and we hope these preliminary results will encourage other researchers to explore the Web not only as an object of sociological interest in itself but as a tool for doing social science on a scale only recently imaginable (Watts, 2007).

Notes

1. For those who find this failure rate unbelievably high, it is useful to remember that, in cultural markets, successes receive much more attention than failures. One way to attempt to overcome this sampling bias is to review the extremely long list of failed movies, which is presented in Appendix B of the book *Fiasco: A History of Hollywood’s Iconic Flops* (Parish, 2006).

2. It is an interesting historical note that this very same issue has puzzled researchers for some time, suggesting that it may not be specific to contemporary cultural markets. None other than Gustave Le Bon, one of the grandfathers of the study of collective social dynamics, briefly addressed this puzzle in his classic, *The Crowd* (Le Bon, [1895] 1969). He wrote, “It is often impossible on reading plays to explain their success. Managers of theaters when accepting pieces are themselves, as a rule, very uncertain of their success, because to judge the matter it would be necessary that they should be able to transform themselves into a crowd.” Unfortunately, he did not explain further.
3. This argument is reminiscent of the debates over path dependence and lock-in in the economics literature. That is, while some scholars have argued that inferior technologies—the QWERTY keyboard and VHS VCR—became standards due to contingent historical events (Arthur, 1994; David, 1985), others have argued that these technologies were not in fact inferior (Liebowitz & Margolis, 1994).
4. Details about song and band selection are available in the study of Salganik et al. (2006).
5. In some respects, our design is similar to many other experiments in which participants are randomly assigned to one of several groups (e.g., Sherif, 1936). However, our experiment does differ in two important respects: First, the size of our “groups” is many times larger than in traditional lab experiments, comprising over 700 people per group; and second, the outcome we are measuring—the full distribution of market share of 48 songs—is considerably more complex. For these reasons, and our desire to connect with the motivating thought experiment, we prefer the term “worlds” in place of the more familiar “groups,” but the meaning is not otherwise different.
6. Complete details about the website, including the text of the consent form, survey, and instructions as well as screenshots, are available in the study of Salganik (2007).
7. These experimental procedures were approved by the Columbia University Institutional Review Board (protocol: IRB-AAAA5286).
8. In the experiments presented here, the download counts were updated honestly, but in another experiment we explored the possibility that manipulated download counts could create self-fulfilling prophecies (Salganik & Watts, 2008).
9. We note that there was a large change in the percentage of females between Experiments 1 and 2. Subjects in both experiments were mainly drawn from <http://www.bolt.com>, but they were drawn from different parts of the site. A majority of the subjects in Experiment 1 were likely drawn from the “music” and “free stuff” sections, while a majority of the subjects in Experiment 2 were likely drawn from a “BoltNote” sent to all users and banner ads. Ideally, these differences would not have occurred, but we have no reason to believe that they had a substantial effect on our findings.
10. One pattern in Experiment 2, but not present in Experiment 1, was that participants were also slightly more likely to listen to the least popular songs. This pattern could represent a form of anticonformist behavior (Berger & Heath, 2008; Simmel, 1957), or it might merely be an artifact of the list format in the song menu, similar to ballot

order effects in elections (Meredith & Salant, 2007). This issue is considered further in the study of Salganik (2007).

11. Other measures of inequality such as the Herfindahl index, coefficient of variation, and total market share of the top five songs give essentially the same results (Salganik, 2007).
12. Another difference between the results from Experiments 1 and 2—one that we did not anticipate—was the decrease in inequality in the independent condition from $G = 0.24$ to $G = 0.19$. We believe that this decrease was caused by a decrease in inequality in the number of listens per song in Experiment 2, which occurred as a result of changing the song menu from a grid to a list (Salganik, 2007).
13. A careful inspection of Figs. 4C and 4D reveals that one of the worlds in Experiment 1 experienced an extreme drop in inequality followed by a gradual rise, a pattern qualitatively different from the others (Salganik, 2007). The reason for this difference is that in this world, the seventh participant downloaded almost every song causing a large drop in inequality, but one that was eventually washed out over the course of the experiment. By the end of the experiment, this world—world 1—ended up with a Gini coefficient of 0.33, consistent with the other social influence worlds in Experiment 1.
14. This nonlinear relationship between appeal and success is thought to be a feature of real cultural markets as has been noted by Rosen (1981) and Frank and Cook (1995).
15. Ideally, we would have made a complete replication with eight social influence worlds and one independent world, but that would have required many more participants than our scaled-down replication.
16. The two songs that had higher batting averages in Experiment 3 were “Gnaw” by Silverfox and “Life’s Mystery” by Dante. These two songs could be described as two of the more “mellow” or “easy-listening” of the available songs.
17. In general, the self-reported ratings were consistent with participants’ download decisions (Salganik, 2007). The probability that a song would be downloaded and given a listen (i.e., the batting average) was highly correlated to the average rating of that song in Experiment 1 ($r = .87$), 2 ($r = .91$), and 3 ($r = .91$). It was also the case in all three experiments that the higher rating a subject gave to a song, the more likely that subject would download that song.
18. Some readers may find this correlation between average market rank in Experiments 1 and 2 surprisingly high given the unpredictability described previously. However, it is important to recall that our definition of unpredictability was based on market share, a key measure in cultural markets, but to be consistent with Lieberman (1997) here we compare market rank. Further, as was shown in Fig. 6, averaging across the eight worlds as we are doing to construct \bar{K}_i obscures a lot of variability.
19. Of course, executives rarely, if ever, attempt to predict the entire distribution of market share, or transplant multiple objects from one market to another simultaneously. In this sense, therefore, our proposed prediction exercise is more difficult than what real-world executives typically face. Nevertheless, we feel that the results still provide some insight.

20. When there are multiple worlds involved (e.g., predicting the success in a social influence world in Experiment 2), we averaged the prediction errors over all worlds.
21. The proportional reduction in error when predicting success in Experiments 2 and 3 is not symmetric because the error from the naive prediction, $\varepsilon_{\text{naive}}$, differs in Experiments 2 and 3.
22. Again, we want to emphasize that our experiment is clearly unlike real cultural markets in a number of ways. First, the media (e.g., MTV, radio, magazines) were not included in any way. Also, we had only 48 songs and a small fraction of participants listened to and downloaded almost all of them. This behavior is clearly not possible in real cultural markets where the number of products is overwhelming. Another difference is that the 48 songs used were fixed at the beginning of the experiment, whereas in practice new products are continually introduced.
23. There are of course some agent-based models with empirical foundations—for example, Axtell et al. (2002) and Bruch and Mare (2006)—but such models are rare.
24. A fourth experiment described elsewhere provides one additional, but slightly different, data point (Salganik & Watts, 2008).
25. These model-based predictions, however, would need to be tested with additional experiments because it is not clear that the individual behavioral parameters that were estimated from this set of experiments would carry over to other experiments.
26. When the interaction in the experiment happens in discrete rounds, the subjects no longer need to be participating at exactly the same time (e.g., a game of chess by mail). However, in practice, subject dropout probably means that simultaneous participation will be required.

Acknowledgments

We thank Peter Hausel, Dan Goldstein, 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. 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 study are publicly available from the data archive of the Office of Population Research at Princeton University (<http://opr.princeton.edu/archive>).

References

- Adler, M. (1985). Stardom and talent. *American Economic Review*, 75(1), 208–212.
- von Ahn, L., & Dabbish, L. (2004). Labeling images with a computer game. *ACM conference on human factors in computing systems, CHI 2004* (pp. 319–326), Vienna, Austria.

- Aizen, J., Huttenlocher, D., Kleinberg, J., & Novak, A. (2004). Traffic-based feedback on the web. *Proceedings of the National Academy of Sciences of the United States of America*, 101(Suppl. 1), 5254–5260.
- Allison, P. D. (1978). Measures of inequality. *American Sociological Review*, 43(6), 865–880.
- Allport, F. H. (1924). The group fallacy in relation to social science. *American Journal of Sociology*, 29(6), 688–706.
- Anderson, P. W. (1972). More is different. *Science*, 177(4947), 393–396.
- Ariely, D. (2008). *Predictably irrational*. New York: Harper Collins.
- Ariely, D., Lowenstein, G., & Prelec, D. (2003). Coherent arbitrariness: Stable demand curves without stable preferences. *Quarterly Journal of Economics*, 118(1), 73–105.
- Arthur, W. B. (1994). *Increasing returns and path dependence in the economy*. Ann Arbor, MI: University of Michigan Press.
- Asch, S. E. (1956). Studies of independence and submission to group pressure: I. A minority of one against a unanimous majority. *Psychological Monographs*, 70.
- Axelrod, R. (1997). *The complexity of cooperation*. Princeton, NJ: Princeton University Press.
- Axtell, R. L., Epstein, J. M., Dean, J. S., Gumerman, G. J., Swedlund, A. C., Harburger, J., Chakravartya, S., Hammond, R., Parker, J. and Parker, M. (2002). Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl. 3), 7275–7279.
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3), 797–817.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512.
- Barton, A. H. (1968). Bringing society back in: Survey research and macro-methodology. *The American Behavioral Scientist*, 12(2), 1–9.
- Becker, H. S. (1982). *Art worlds*. Berkeley, CA: University of California Press.
- Bentley, R. A., Lipo, C. P., Herzog, H. A., & Hahn, M. W. (2007). Regular rates of popular culture change reflect random copying. *Evolution and Human Behavior*, 28, 151–158.
- Berger, J., & Heath, C. (2008). Who drives divergence? Identity-signaling, outgroup dissimilarity, and the abandonment of cultural tastes. *Journal of Personality and Social Psychology*, 95(3), 593–607.
- Bielby, W. T., & Bielby, D. D. (1994). “All hits are flukes”: Institutionalized decision making and the rhetoric of network prime-time program development. *American Journal of Sociology*, 99(5), 1287–1313.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashions, customs, and cultural change as information cascades. *Journal of Political Economy*, 100(5), 992–1026.
- Bond, R., & Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using Asch’s (1952b, 1956) Line Judgement Task. *Psychological Bulletin*, 119(1), 111–137.
- Borghesi, C., & Bouchaud, J.-P. (2007). Of songs and men: A model for multiple choice with herding. *Quality and Quantity*, 41, 557–568.
- Bourdieu, P. (1984). *Distinction* (R. Nice, Trans.). Cambridge, MA: Harvard University Press.
- Bowles, S. (1998). Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature*, 36, 75–111.
- Bruch, E. E., & Mare, R. D. (2006). Neighborhood choice and neighborhood change. *American Journal of Sociology*, 112(3), 667–709.
- Carter, B. (2006). *Desperate networks*. New York: Doubleday.
- Caves, R. E. (2000). *Creative industries: Contracts between art and commerce*. Cambridge, MA: Harvard University Press.
- Chung, K. H., & Cox, R. A. K. (1994). A stochastic model of superstardom: An application of the Yule distribution. *The Review of Economics and Statistics*, 76(4), 771–775.
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591–621.
- Coleman, R. (1989). *The man who made the Beatles*. New York: McGraw-Hill.
- Coleman, J. S. (1990). *Foundations of social theory*. Cambridge, MA: Harvard University Press.

- Coleman, J. S., Katz, E., & Menzel, H. (1957). The diffusion of innovation among physicians. *Sociometry*, 20(4), 253–270.
- Costner, H. L. (1965). Criteria for measures of association. *American Sociological Review*, 30(3), 341–353.
- Coulter, P. B. (1989). *Measuring inequality: A methodological handbook*. Boulder, CO: Westview Press.
- Crutchfield, R. S. (1955). Conformity and character. *American Psychologist*, 10, 191–198.
- Cutting, J. E. (2003). Gustave Caillebotte, French Impressionism, and mere exposure. *Psychonomic Bulletin & Review*, 10(2), 319–343.
- David, P. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2), 332–337.
- De Vany, A. (2004). *Hollywood economics: How extreme uncertainty shapes the film industry*. London: Routledge.
- De Vany, A., & Lee, C. (2001). Quality signals in information cascades and the dynamics of the distribution of motion picture box office revenues. *Journal of Economic Dynamics and Control*, 25, 593–614.
- De Vany, A., & Walls, W. D. (1996). Bose-Einstein dynamics and adaptive contracting in the motion picture industry. *The Economic Journal*, 106, 1493–1514.
- Denisoff, R. S. (1975). *Solid gold*. New Brunswick, NJ: Transaction Books.
- Deutsch, M., & Gerard, H. B. (1955). A study of normative and informative social influences upon individual judgement. *Journal of Abnormal Social Psychology*, 51, 629–636.
- DiMaggio, P. (1987). Classification in art. *American Sociological Review*, 52(4), 440–455.
- DiPrete, T. A., & Eirich, G. (2006). Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments. *Annual Review of Sociology*, 32, 271–297.
- Dodds, P. S., Muhamad, R., & Watts, D. J. (2003). An experimental study of search in global social networks. *Science*, 301, 827–829.
- Dodds, P. S., & Watts, D. J. (2004). Universal behavior in a generalized model of contagion. *Physical Review Letters*, 92, 218701.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Boca Raton, FL: Champan & Hall.
- Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom-up*. Washington, DC: Brookings Institution/MIT Press.
- Faulkner, R. R., & Anderson, A. B. (1987). Short-term projects and emergent careers: Evidence from Hollywood. *American Journal of Sociology*, 92(4), 879–909.
- Frank, R. H., & Cook, P. J. (1995). *The winner-take-all society*. New York: The Free Press.
- Gans, H. (1974). *Popular culture and high culture*. New York: Basic Books.
- Gitlin, T. (1983). *Inside primetime*. New York: Patheon.
- Goldman, W. (1983). *Adventures in the screen trade*. New York: Warner Books.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90.
- Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, 11(3), 508–514.
- Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, 12(1), 81–87.
- Goldstone, R. L., & Janssen, M. A. (2005). Computational models of collective behavior. *Trends in Cognitive Science*, 9(9), 424–430.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6), 1420–1443.
- Hamlen, W. A. (1991). Superstardom in popular music: Empirical evidence. *The Review of Economics and Statistics*, 73(4), 729–733.
- Hanson, W. A., & Putler, D. S. (1996). Hits and misses: Herd behavior and online product popularity. *Marketing Letters*, 7(4), 297–305.
- Harrington, B., & Fine, G. A. (2000). Opening the “black box”: Small groups and twenty-first-century sociology. *Social Psychology Quarterly*, 63(4), 312–323.
- Hartwell, L. H., Hopfield, J. J., Leibler, S., & Murray, A. W. (1999). From molecular to modular cell biology. *Nature*, 402, C47–C52.

- Hedström, P. (1998). Rational imitation. In P. Hedström & R. Swedberg (Eds.), *Social mechanisms: Analytical approach to social theory* (pp. 306–327). Cambridge, England: Cambridge University Press.
- Hedström, P. (2005). *Dissecting the social*. Cambridge, England: Cambridge University Press.
- Hedström, P. (2006). Experimental macro sociology: Predicting the next best seller. *Science*, 311, 786–787.
- Hirsch, P. M. (1972). Processing fads and fashions: An organization-set analysis of cultural industry systems. *American Journal of Sociology*, 77(4), 639–659.
- Katz, E., & Lazarsfeld, P. F. (1955). *Personal influence*. Glencoe, IL: Free Press.
- Kearns, M., Suri, S., & Montfort, N. (2006). An experimental study of the coloring problem on human subject networks. *Science*, 313, 824–827.
- Kelley, H. H., & Thibaut, J. W. (1978). *Interpersonal relations: A theory of interdependence*. New York: John Wiley & Sons.
- Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2), 117–136.
- Kraut, R., Olson, J., Banaji, M., Bruckman, A., Cohen, J., & Couper, M. (2004). Psychological research online: Report of board of scientific affairs' advisory group on the conduct of research on the Internet. *American Psychologists*, 59(2), 105–117.
- Krueger, A. B. (2005). The economics of real superstars: The market for rock concerts in the material world. *Journal of Labor Economics*, 23(1), 1–30.
- Ku, G., Galinsky, A. D., & Murnighan, J. K. (2006). Starting low but ending high: A reversal of the anchoring effect in auctions. *Journal of Personality and Social Psychology*, 90(6), 975–986.
- Latané, B., & L'Herrou, T. (1996). Spatial clustering in the conformity game: Dynamic social impact in electronic groups. *Journal of Personality and Social Psychology*, 70(6), 1218–1230.
- Lawless, J. (2005). The interview: Nigel Newton: Is there life after Harry Potter? You bet your Hogwarts there is. *Independent (London)*, July 3, 2005.
- Le Bon, G. (1969). *The crowd: A study of the popular mind*. New York: Ballantine Books.
- Leibenstein, H. (1950). Bandwagon, snob and Veblen effects in the theory of consumer demand. *Quarterly Journal of Economics*, 64(2), 183–207.
- Lieberson, S. (1997). Modeling social processes: Some lessons from sports. *Sociological Forum*, 12(1), 11–35.
- Liebowitz, S. J., & Margolis, S. E. (1994). Network externality: An uncommon tragedy. *Journal of Economic Perspectives*, 8(2), 131–150.
- Luce, R. D., & Raiffa, H. (1957). *Games and decisions: Introduction and critical survey*. New York: Wiley.
- Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28, 143–166.
- Mason, W. A., & Watts, D. J. (2009). Financial incentives and the “performance of crowds.” *Proceedings of the Human Computation Workshop*. Paris: ACM, June 28, 2009.
- Meredith, M., & Salant, Y. (2007). The causes and consequences of ballot order-effects. Cambridge, MA: Department of Political Science, MIT.
- Moretti, E. (2008). Social learning and peer effects in consumption: Evidence from movie sales. NBER Working Paper No. 13832. Available at <http://www.nber.org/papers/>. Accessed on June 1, 2009.
- Northcraft, G. B., & Neale, M. A. (1987). Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes*, 39, 84–97.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002a). E-Research: Ethics, security, design, and control in psychological research on the Internet. *Journal of Social Issues*, 58(1), 161–176.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002b). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics*, 6(1), 101–115.
- Ordonez, J. (2002). Behind the music: MCA spent millions on Carly Hennessy—Haven't heard of her? *The Wall Street Journal*, Feb. 26, A1, A10.
- Parish, J. R. (2006). *Fiasco: A history of Hollywood's iconic flops*. Hoboken, NJ: John Wiley & Sons.
- Park, C. W., & Lessig, V. P. (1977). Students and housewives: Difference in susceptibility to reference group influence. *Journal of Consumer Research*, 4(2), 102–110.

- Pasupathi, M. (1999). Age differences in response to conformity pressure for emotional and nonemotional material. *Psychology and Aging, 14*(1), 170–174.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology, 43*, 87–131.
- Payne, J. W., Bettman, J. R., & Schkade, D. A. (1999). Measuring constructed preferences: Towards a building code. *Journal of Risk and Uncertainty, 19*, 243–270.
- Reips, U.-D. (2002). Standards for Internet-based experimenting. *Experimental Psychology, 49*(4), 243–256.
- Roberts, M. E., & Goldstone, R. L. (2006). EPICURE: Spatial and knowledge limitations in group foraging. *Adaptive Behavior, 14*(4), 291–313.
- Rosen, S. (1981). The economics of superstars. *American Economic Review, 71*(5), 845–858.
- Salganik, M. J. (2007). *Success and failure in cultural markets*. PhD thesis, New York: Department of Sociology, Columbia University.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science, 311*, 854–856.
- Salganik, M. J., & Watts, D. J. (2008). Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly, 71*(4), 338–355.
- Schelling, T. C. (1978). *Micromotives and macrobehavior*. New York: W. W. Norton.
- Seabrook, J. (2000). *Nobrow*. New York: Vintage Books.
- Sherif, M. (1936). *The psychology of social norms*. New York: Harper & Brothers.
- Simmel, G. (1957). Fashion. *American Journal of Sociology, 62*(6), 541–558.
- Simon, H. A. (1955). On a class of skew distribution functions. *Biometrika, 42*(3/4), 425–440.
- Simonson, I., & Drolet, A. (2004). Anchoring effects on consumers' willingness-to-pay and willingness-to-accept. *Journal of Consumer Research, 31*, 681–690.
- Skitka, L. J., & Sargis, E. G. (2006). The Internet as psychological laboratory. *Annual Review of Psychology, 57*, 529–555.
- Slovic, P. (1995). The construction of preference. *American Psychologist, 50*, 364–371.
- de Solla Price, D. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science, 27*(5), 292–306.
- Sorensen, A. T. (2007). Bestseller lists and product variety. *Journal of Industrial Economics, LV*(4), 715–738.
- Srivastava, S., John, O. P., Gosling, S. D., & Potter, J. (2003). Development of personality in early and middle adulthood: Set like plaster or persistent change? *Journal of Personality and Social Psychology, 84*(5), 1041–1053.
- Thibaut, J. W., & Kelley, H. H. (1959). *The social psychology of group*. New York: Wiley.
- Tucker, N. (1999). The rise and rise of Harry Potter. *Children's Literature in Education, 30*(4), 221–234.
- Tversky, A., & Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131.
- UNDP (2004). *Human development report*. New York: United Nations Development Programme.
- Vogel, H. L. (2004). *Entertainment industry economics*. Cambridge, England: Cambridge University Press.
- Watts, D. J. (2002). A simple model of global cascades in random networks. *Proceedings of the National Academy of Sciences of the United States of America, 99*(9), 5766–5771.
- Watts, D. J. (2007). A twenty-first century science. *Nature, 445*, 489.
- Yule, G. U. (1925). A mathematical theory of evolution based on the conclusions of Dr. J. C. Willis, F.R.S. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 213*, 21–87.
- Zelditch, M. (1969). Can you really study an army in the laboratory? In A. Etzioni and E. N. Lehman (Eds.), *A sociological reader on complex organization* (pp. 528–539). New York: Holt, Rinehart, and Winston.