Self-Correcting Dynamics in Social Influence Processes[1]

Arnout van de Rijt
Utrecht University

Social influence may lead individuals to choose what is popular over what is best. Whenever this happens, it further increases the popularity advantage of the inferior choice, compelling subsequent decision makers to follow suit. The author argues that despite this positive feedback effect, discordances between popularity and quality will usually self-correct. Reanalyzing past experimental studies in which social information initially heavily favored inferior options, the author shows that in each experiment superior alternatives gained in popularity. This article also reports on a new experiment in which a larger number of subject choices allowed trials to be run to convergence and shows that in each trial the superior alternative eventually achieved popular dominance. To explain the persistent dominance of bestsellers, celebrities, and memes of seemingly questionable quality in everyday life in terms of social influence processes, one must identify conditions that render positive feedback so strong that self-correcting dynamics are prevented.

Over the past decades, consumers, voters, employers, and investors, when faced with some choice among products, candidates, or ideas, have increas-

[1] I thank Vincent Buskens, Damon Centola, and Eran Shor for in-depth discussion; Jerker Denrell, Vincenz Frey, Jason Jones, Chengwei Liu, Werner Raub, Matthew Salganik, Jeroen Weesie, and four anonymous reviewers for helpful comments; and Michael Claffey for implementation of the experiment. This work was supported by the National Science Foundation, grant SES-1340122 to the author. Direct correspondence to Arnout van de Rijt, Department of Sociology, Utrecht University, Padualaan 14, Utrecht, the Netherlands, 3584CH. E-mail: a.vanderijt@uu.nl

0002-9602/2019/12405-0004$10.00

ingly been able to draw on information about the choices made by others before them. Through likes and follower counts, bestseller lists, election polls, citation indices, and music charts, modern communication technology provides continuously updated information on the aggregate popularity of choice alternatives in nearly every domain of industry and culture. In these environments, it can happen that an option of inferior quality or utility has an incidental popularity advantage over a superior alternative (e.g., a better product, a more qualified candidate, a smarter idea, or a more truthful account). This presents decision makers with a choice between what is popular and what would appear to be best. Many experimental studies have shown that people, when placed in such scenarios where many others have made a seemingly bad choice before them, are often persuaded to nonetheless make that same choice (Asch 1951; Cialdini and Goldstein 2004; Salganik, Dodds, and Watts 2006; Sorenson 2007; Anderson and Holt 2008; Salganik and Watts 2008; Davis, Bowers, and Memon 2011; Margetts et al. 2011; Muchnik, Aral, and Taylor 2013; Hosanagar et al. 2014; van de Rijt et al. 2014; Lynn et al. 2016*a*; Lynn, Walker, and Peterson 2016*b*). By also choosing the inferior option, individuals increase its popularity, compelling yet others to follow suit.

This article is concerned with the choice dynamics that ensue in this scenario: What happens to the relative popularity of the inferior option vis-à-vis the superior alternative? Prior studies have suggested that in social influence processes the idiosyncratic choices of early decision makers may set the tone for later others and thus remain dominant among later decision makers (Bikhchandani, Hirshleifer, and Welch 1992; Salganik et al. 2006; Lynn, Podolny, and Tao 2009; Krumme et al. 2012; van de Rijt et al. 2014; Frey and van de Rijt 2016). That is, as individuals are influenced toward already popular options, their biased choices extend the lead of these options over alternatives, which in turn further skew the decisions of yet later generations of choosers, and so on. A process of cumulative advantage (Merton 1968; Granovetter 1978; Allison, Long, and Krauze 1982; Barabási and Albert 1999; DiPrete and Eirich 2006; Denrell and Liu 2012) then leads the early popularity advantage for the inferior object over the superior alternative to be perpetuated. In such situations, social influence bias may allow popularity to become decoupled from quality (Lynn et al. 2009, 762; Manzo and Baldassarri 2015, 346; Correll et al. 2017), rendering long-term outcomes dependent on initial conditions (e.g., Goldstone 1998). These *self-reinforcing* social influence dynamics have been used to explain large-scale investment in bad financial assets (Bikhchandani and Sharma 2001; Lorenz et al. 2011), the inability of experts to predict success in cultural markets (Salganik et al. 2006; Keuschnigg 2015), bubbles in rating and reputation systems (Muchnik et al. 2013; van de Rijt et al. 2014), and the arbitrary success and failure of political and financial campaigns (van de Rijt et al. 2014; Gonzalez-Vaillant et al. 2015). More generally, they provide microfoundations for the social construction of valuations and beliefs (Watts 2007; Zuckerman 2012).

Here I emphasize a second theoretical possibility: many individuals may choose the most popular option, while substantial numbers of others resist the social influence and choose what is best. Their deviant choices then become reason for others to also deviate. As a result, the relative popularity of the superior alternative increases, leading later generations of decision makers to choose it in greater numbers, and so on, until it gains the lead. Such a *self-correcting* dynamic (Elster 1989, 373–75; Goeree et al. 2007) may drive the superior alternative all the way from a small minority position to a majority position, recovering its natural dominance in popularity. In this dynamic regime, social influence bias does not allow the popularity of entities to remain durably decoupled from their quality and long-term popularity ranks are independent of initial conditions. Gould's model of status hierarchies (Gould 2002; Lynn et al. 2009, 761; Manzo and Baldassarri 2015) and recent models of long-term scientific impact (Wang, Song, and Barabási 2013; Sinatra et al. 2016) fall into this regime. They predict that despite the operation of cumulative advantage processes, in equilibrium, status ranks will closely follow quality ranks.

Whether social influence processes are self-reinforcing or self-correcting is crucial, as only in the former case can social influence bias at the microlevel explain collective aberration as a stable macrolevel outcome. I propose that whether the popularity of inferior options perpetuates or is corrected critically depends on the magnitude of the social influence effect. It is determined by whether a majority choosing a bad thing can lead an even greater majority to subsequently choose that bad thing. As the degree of social influence will vary across contexts, depending on the mechanisms driving it, errors may cascade in some but be corrected in others.

This article focuses on a setting that is commonly considered in the theoretical and empirical study of social influence processes in sociology, social psychology, political science, and economics. In this setting, members of a group or population must, one at a time, choose from among a number of discrete alternatives, after first observing what other members before them chose. This "social information" may be provided as a historical sequence of choices or as tallies for each alternative. Choice alternatives are of different quality, and individuals have a sense of what is good and what is bad but are uncertain about their quality assessments. Social information may present normative pressure to conform or may lead individuals to think their own judgment may be off and instead trust that of others. Apart from these social influence effects there is no other mechanism that prevents individuals from considering unpopular alternatives in this scenario. Access to or knowledge of minority choices is not restricted, as it would be under differential media attention, network diffusion, or limited shelf space (van de Rijt et al. 2013). Their lack of popularity does not render them less valuable, as it would in contexts of economic or technological complementarity (Schelling 1978; Arthur 1989; DiMaggio and Garip 2011). And minority deviants are not sanctioned (Centola, Willer, and Macy 2005; Centola and Macy 2007; Willer, Kuwabara, and Macy 2009).

I argue that under these conditions in which feedback is limited to the conformity and uncertainty-reducing effects of social information, accidental majority support for an inferior option will often self-correct. I reanalyze data from six experimental studies of the above scenario in which subjects were confronted with inferior alternatives that had been popularized through social information. Each established a significant effect of social information on individual decision-making. Yet reanalysis of these studies shows that crowds corrected themselves. For example, the Asch (1951) experiments, while widely cited as some of the first controlled studies showing that people can be pressured into agreeing with an obvious falsehood (even if Asch himself emphasized the independent mindedness shown by many subjects; e.g., Friend, Rafferty, and Bramel 1990), provide strong evidence that social influence fizzles out, with decision-making quickly converging on the correct choice. In reanalysis of the Musiclab experiments (Salganik et al. 2006; Salganik and Watts 2008), which served to demonstrate the unpredictability of success in cultural markets, I find that discordances between quality and popularity were largely corrected during trials and this corrective process still seemed ongoing when trials were ended. I show that even in stylized choice situations created in the laboratory where the theory of information cascades (Banerjee 1992; Bikhchandani et al. 1992) predicts that rational individuals act on the available social information and ignore their personal inclination, subjects pay too little attention to what others have done for a bad decision to propagate. In every study, the initial discrepancy between quality and popularity was on a path of convergence on the superior alternative. That is, the superior alternatives were recovering from the initial setback. However, the experiments did not include long enough strings of subject choices to see this process fully unfold.

I report on a novel experiment in which I similarly popularized inferior alternatives, but with longer sequences of subject choices than in previous experiments. In this experiment the social influence process fully corrected itself, with the superior alternative reaching stable dominance in every trial. Together these results suggest that in order to explain the persistent dominance of bestsellers, celebrities, and memes of seemingly questionable quality in everyday life in terms of social influence processes, one must identify conditions that render feedback so strong that self-correcting dynamics are prevented.

SELF-CORRECTING DYNAMICS IN PAST EXPERIMENTAL STUDIES

A large number of experimental studies present evidence for significant effects of social information on choice under controlled experimental conditions. Below I review studies that popularize an inferior choice option and that demonstrate that doing so biases subject choices in favor of this option. For each of these studies I evaluate whether the effect of social information is strong enough for a self-reinforcing dynamic in which inferior options remain dominant or instead gives rise to a self-correcting dynamic in which superior alternatives recover from the popularity setback.

The studies I review were selected because in each (1) a choice alternative A that is no better or strictly worse than its alternative B was popularized through social information, (2) popularity did not impact the material costs or benefits of alternatives or their accessibility, and (3) the necessary data for reanalysis were available in the article or made available by the authors. The nature of alternatives greatly varied across studies, including right and wrong answers, petitions for causes of variable importance, crowdfunding projects of variable promise, good and bad wines, and good and bad songs. In each study B is of higher quality, by which I will mean that under equal past popularity of A and B or in the absence of information on past choices, at least half of subjects would opt for B. B may but need not be intrinsically or objectively of greater merit or truthfulness. This definition of quality allows an investigation of the conditions under which a population of sequential decision makers can produce a majority choice that they would not actually have favored if choosing independently. Consider individuals labeled by the order in which they choose, where 1 is the first decision maker. Inferior alternative A is more popular than B among early decision makers. A focal individual i must make a decision knowing that among all prior individuals $j \in \{1,\ldots i-1\}$ a majority chose A over B. The question becomes, Will A's popularity advantage over B be sustained in subsequent decision-making by i and those deciding after i, or will quality increasingly determine popularity and correct the choice process?

An analytical answer can be found in the strength of social influence. Individual i's choice is denoted by $x_i \in \{A, B\}$. I assume that the propensity for i to choose A over B is a monotonically increasing function of the proportion of past decision makers that chose A over B. Then, for the population to lock in on inferior option A, the process must reach a point where individuals still favor A and select A over B with a probability $P(x_i = \{A\})$ that is greater than or equal to the fraction that selected A in the past:

$$P(x_i = \{A\}) \geq \frac{\sum_{j<i}(x_j = \{A\})}{i-1}. \quad (1)$$

In figure 1, where the x-axis measures the percentage of past choices of option A and the y-axis the probability that the next decision maker chooses option A, the shaded area shows where condition (1) is satisfied. Outside the shaded area, the probability of selecting the inferior alternative is smaller than the percentage of past choices. A sequence of choices that by chance initially favors the inferior alternative but remains outside the shaded area is self-correcting: the initial popularity advantage evokes a smaller relative advantage among later adopters, in turn yielding a yet weaker response in subsequent generations.

As an example of this self-correcting dynamic, imagine that only 15% of people would in the absence of information about the prior choices of others choose A over B, making B the higher quality alternative by our definition of quality. However, by accident, initial decision makers disproportionately represent the 15% minority that would naturally prefer A over B. Let's say 95% chose A. In scenarios of extreme social influence, where such a head start leads subsequent decision makers to nearly always follow suit, A's popularity advantage could be sustained. But if individuals balance their personal preferences against the social influence signal, they will choose A, say, 55% of the time (halfway between 15% and 95%). Condition (1) is not satisfied. The result is that yet later decision makers face a reduced social influence signal of only $95\%/2 + 55\%/2 = 75\%$ in favor of A. Their response to the reduced signal will then favor A even less, choosing A less than half the time, leading the social influence signal for yet later choosers to be further reduced and so on. The process then converges to dominance of superior alternative B.

We can now evaluate the experiments that artificially popularized an inferior choice in terms of the location of their social influence magnitude estimates in figure 1. Panels 1 and 2 of figure 1 show for each the initial percentage of subjects choosing A and the subsequent percentage of subjects choosing A in response to that initial popularity signal. Note that for each study in panel 1 of figure 1, higher initial popularity levels are associated with higher subsequent popularity levels, supporting the monotonicity assumption. (Panel 2 does not permit an evaluation of monotonicity because here initial popularity is correlated with quality.)

The diamonds in panel 1 of figure 1 represent data points from pioneering experiments on social influence by Asch (1956). Subjects were shown a card with a vertical line and another card with three lines of variable length, one of which quite clearly matched the line on the initial card. Before being asked to give the right answer (B), subjects first heard five to eight others (confederates) give one of the two wrong answers (A). Subjects chose this incorrect answer A 37% of the time under unanimous support by confederates (corresponding to the diamond placed at 100% popularity and 37%

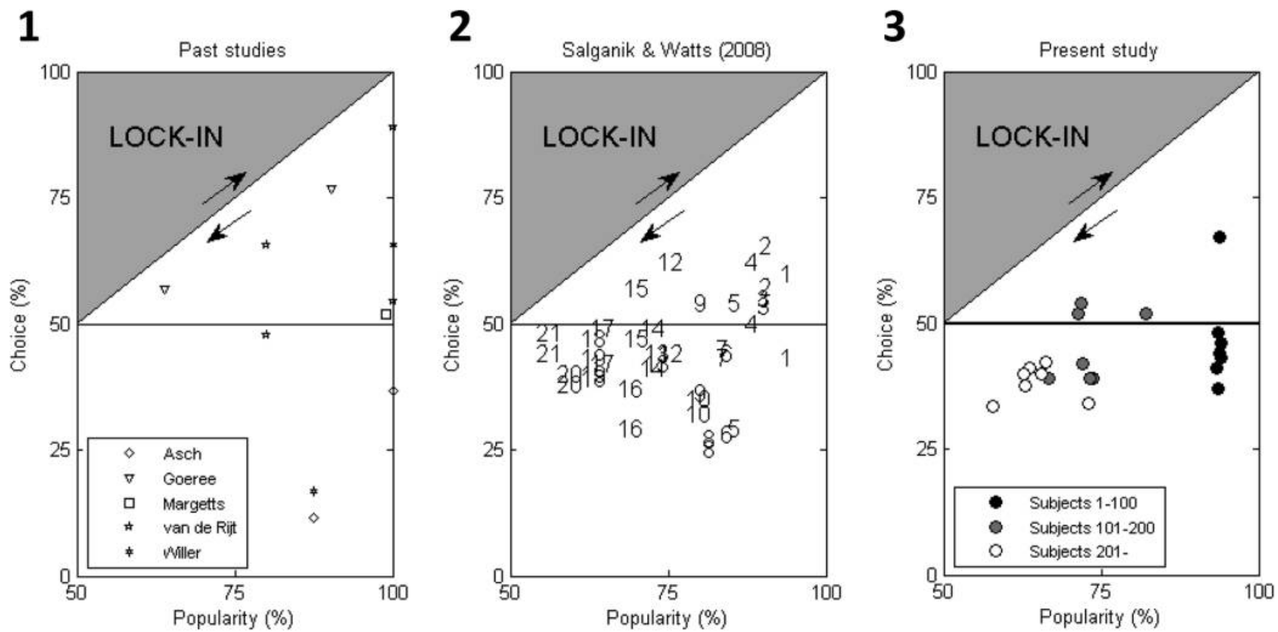

FIG. 1.—Choice of the inferior alternative by its popularity in past social information experiments (panels 1 and 2) and in the present study (panel 3). In each panel the vertical axis measures choice as the percentage of experimental subjects that chose the inferior option over the superior alternative, while the horizontal axis measures popularity as the percentage of prior choices for the inferior option experimental subjects faced. In the shaded lock-in area choice equals or exceeds popularity. In panel 1 the different symbols distinguish five experimental studies. In panel 2 each number represents a pair of songs in Salganik and Watts (2008) whose popularity counts were swapped. Pairs in which the superior song was better ranked after swapping (11, 22, 23, and 24) are omitted from panel 2. In panel 3 the seven choices in the new experiment described in the main text are shown separately for the first 100 subjects, the second 100 subjects, and for all later subjects.

choice in panel 1). This percentage of 37% provides strong support for the idea that under pressure people can be made to conform to something that is obviously wrong. The point I wish to make is that the Asch results show that these errors cannot be sustained in a population. The percentage conformists decreased to 9% in a variation of the original experiment in which one confederate or subject gave the right answer (which is the diamond placed in panel 1 at the coordinate with popularity 87.5% and choice 9%). In other words, under the most favorable initial conditions for the wrong answer, only 37% deviated from the right answer to begin with, a point Asch himself emphasized (e.g., Friend et al. 1990), and only a single choice for the right answer was sufficient to reduce subsequent choices for the wrong answer down to 9%. The evidence strongly suggests that any sequence of decisions in this scenario converges to the correct answer. While Asch (1951) is commonly cited as one of the best-known early studies showing controlled experimental evidence of social influence, the present analysis shows that it at the same time provides clear evidence that individual responses to social pressure are too weak to perpetuate a false choice. The Asch study has been replicated many times with various modifications of the original setup. Across 133 conformity studies reviewed by (Bond and Smith 1996), estimates of the conformity percentage roughly follow a normal distribution with the original Asch result reported in figure 1 positioned toward the high end (pp. 132–36). These replications indicate that the claim that social influence in the Asch setting fizzles out readily generalizes to related settings and different subject populations in many countries.

In an experiment reported in Goeree et al. (2007; triangles in panel 1), trials involved 20 or 40 subjects in a laboratory positioned behind networked computers. The computers showed two urns filled with red and blue balls. Urn 1 contained more red balls, while urn 2 contained more blue balls. One urn was then chosen at random from which balls were repeatedly drawn with replacement. Each subject's draw was private and upon seeing the ball the subject made a guess about the urn it came from. Before their turn subjects were informed about all prior guesses and in which order they had occurred. In this scenario, when facing a unanimous choice by several prior subjects, it is rational to ignore one's own ball and copy earlier subjects. Because the same is true for the next subject, rational decision makers should then theoretically form "information cascades" (Banerjee 1992; Bikhchandani et al. 1992) in which everyone chooses the same urn, which may well be the wrong one. Information cascades are commonly cited as a key mechanism for herd behavior in the adoption of products or technologies under quality uncertainty (Hung and Plott 2001; Hirshleifer and Teoh 2003). Panel 1 of figure 1 shows the fraction of subjects choosing the wrong ball as a function of the fraction of prior subjects choosing the wrong ball. Subject choices in panel 1 are split into two sets of data points, namely, choices made by subjects facing a popularity of the wrong ball between 75% and 100% of prior subjects and situations in which subjects faced a popularity of the wrong ball between 50% and 75%. (Other binning decisions yield the same findings.) The results show a positive relationship, confirming that individuals were influenced by the choices of others, a finding that has been replicated across a range of information cascade experiments (Anderson and Holt 2008). Nonetheless, the estimates in panel 1 lie well outside the shaded lock-in area. Subjects were paid for correct guesses, yet despite this incentive, they still paid too little attention to the choices of others than they should have rationally, thereby failing to give the information cascades the predicted self-perpetuating character. Exposed to a large majority of earlier choices for the wrong urn, a smaller majority of subjects followed suit (in panel 1, popularity 90%, choice 77%), and exposed to a yet smaller majority of choices for the wrong urn an even smaller fraction of subjects chose the wrong urn (in fig. 1, popularity 64%, choice 57%). Again, the evidence is consistent with the proposed self-correcting dynamic through which populations, despite a false start, converge on the correct choice.

Margetts et al. (2011) presented subjects with six petitions for good causes through a custom-built web interface. Subjects were asked if they were willing to sign these petitions. Petitions were either presented without social information or with a low, medium, or high number of prior signatories. Margetts et al. found subjects to be significantly more willing to sign petitions in the high category than in the other categories. However, the effect of showing a petition with over a million signatures instead of one with fewer than 100 signatures was an increase in the probability of signing from 62% to 67% on average (squares in panel 1 of fig. 1: popularity 99%, choice $100\% * 67/(62 + 67) = 52\%$). With 52%, choice is only marginally in favor of the alternative that was initially vastly more popular, close to the 50% level of zero social influence. In other words, social influence bias was minimal even after an extreme social information advantage, so that right in the next generation of choosers the original bias would have been mostly corrected. Gonzalez-Vaillant et al. (2015) and van de Rijt et al. (2016) replicate these minimal effects of social information on petition signing. Our analysis of this experiment is again consistent with the notion that the influence effect of social information leads initial discrepancies between popularity and quality to be corrected.

Van de Rijt et al. (2014) sampled projects from the crowdfunding platform kickstarter.com that had not yet received any donations and randomly assigned these to three conditions in which they made zero, one, or four financial donations. Van de Rijt et al. also sampled product reviews from the platform epinions.com that had not yet received any ratings. High-quality reviews were given zero, one, or four positive ratings. The study identified significant feedback through randomized experimental intervention in familiar real-world settings. As with the other studies I reviewed, a significant social influence effect was found. However, the feedback effects that were found fall well outside the lock-in area (pentagrams in panel 1). Campaigns with four initial donations accumulated 8.16 times the number of donations given to campaigns with zero initial donations (in panel 1: popularity $= 100\% * 4/(0 + 4) = 100\%$, choice $= 100\% * 8.16/(1 + 8.16) = 89\%$), but only 1.91 times that of campaigns with one initial donation (popularity $= 100\% * 4/(1 + 4) = 80\%$, choice $= 100\% * 1.91/(1 + 1.91) = 66\%$). This last difference between the one and four donations conditions was not found to be statistically significant and is consistent with full correction to equal popularity (50%). Moreover, van de Rijt et al. did not find a difference between conditions in the number of projects that reached their funding goal, indicating that initial successes did not result in noticeably greater long-term success. Product reviews with four initial positive ratings accumulated 1.19 times the number of positive ratings by third parties than product reviews with zero initial positive ratings (popularity 100%, choice $=100\% * 1.19/(1 + 1.19) = 54\%$), and only 0.91 times the number of positive ratings than product reviews with one initial positive rating (popularity 80%, choice $100\% * 0.91/(1 + 0.91) = 48\%$). In other words, an advantage of 80% was again fully corrected to equal popularity.

Willer et al. (2009) conducted a wine-tasting experiment in which subjects tasted and rated three wines 1, 2, and 3 (labeled A, B, and C in the original study) using letter grades ranging from A (best) to F (worst). Subjects were always in fifth position and were given the (fake) ratings of four prior subjects who rated wine 1 highest and wine 2 lowest. Subjects did not know that wines 1, 2, and 3 were in actuality the same wine nor that wine 3 had been tainted with vinegar. Panel 1 shows the percentage of $1 > 2$ over the joint number of $1 > 2$ and $2 > 1$ raters, splitting ties. As Willer et al. report, subjects exhibited significant susceptibility to social influence, rating wine 1 on average higher than 2 despite equal quality. However, the percentage of subjects rating 1 above 2 dropped from the stimulus 100% to 66%, splitting ties (hexagram in panel 1, popularity 100%, choice 66%). Note that as in the Margetts et al. and van de Rijt et al. studies the popularized choice option was of equal quality to the alternative, making it easier for subjects to conform as they did not have to violate their personal inclination. Subjects still did not follow the emerging norm enough to perpetuate a false belief. While the panel 1 analysis does not prove a dynamic process would have fully eliminated the popularity advantage of wine 1, the response being much weaker than the signal strongly suggests A’s advantage would have kept shrinking among later subjects. Panel 1 also compares wines 2 and 3, which were of different quality, as wine 3 was tainted with vinegar. Wine 3 received better scores than wine 2 from only 17% of participants (splitting ties), despite the four stimulus ratings consistently rating 3 higher than 2 (hexagram in panel

1, popularity 88%, choice 17%), showing that quality dominated social information in the determination of subjects' evaluation of the wines. The inferior option's initial popularity advantage was fully corrected toward majority support for the superior alternative. I stress that Willer et al.'s aim was not at all to investigate lock-in through social influence but rather to test whether subjects in the next phase of the experiment would publicly criticize those with deviant wine evaluations, which they indeed found to be the case. I reanalyze the data from the first phase of the experiment here because they provide compelling empirical evidence from yet another context regarding the ability of social information to perpetuate bad decisions in a population. Again I find that discrepancies between popularity and quality are corrected.

In Salganik and Watts (2008), subjects were presented with a screen showing links to 48 previously unknown songs. Subjects could listen to, rate, and download these songs. In two trials (referred to by the authors as "worlds"), songs were accompanied by counts of downloads made by prior subjects in that same world, which after a start-up period were suddenly inverted: the most popular song so far was displayed to the next subject at the bottom of the song list with the download count of the least popular song, while, vice versa, the least popular song was shown to be most popular, the second most popular song was given the download count and rank of the 47th most popular song, and so on. Hence there were 24 pairs whose popularity counts were swapped by the inversion, which I index here by the highest rank (e.g., 1 denoting the pair consisting of the best- and worst-ranked song). A comparison with an "independent" condition in which subjects made downloading decisions without social information shows that in 20 of these 24 pairs the more popular song before the inversion was of greater quality than the less popular song before the inversion, so that the inversion accomplished popularization of the inferior alternative. Panel 1 of figure 1 shows for these 20 pairs along the horizontal axis the percentage of downloads of the lower quality song right after the inversion and along the vertical axis the subsequent percentage of downloads. The remaining four pairs (indexed 11, 22, 23, and 24) for which the inversion popularized the superior quality song are not included in the analysis as they cannot help evaluate inferior lock-in. Salganik and Watts (2008) showed that the inversion impacted download behavior in favor of the lower quality songs, which were initially downloaded more often than their quality would warrant. However, panel 2 in figure 1 shows that subjects nonetheless downloaded good songs over bad songs at a rate exceeding their relative popularity. As a result, later subjects were confronted with improved relative popularity scores for good songs, which in turn further increased their download rates. The experiment was terminated while this corrective process was still ongoing, and we cannot definitively conclude from figure 1 that the process would not in some cases have entered the upper diagonal area. However, panel 2 shows that without exception, regardless of the magnitude of the advantage for the inferior option, the advantage always shrank during the experiment toward dominance of the superior alternative. Thus, regardless of whether uncertainty about relative quality was high or low, the better song consistently grew in relative popularity. Moreover, panel 2 shows that for most pairs the higher quality song was downloaded more often than the lower quality song after the inversion, despite the popularity information favoring the latter. In all these cases the natural popularity ordering had thus been restored. In other cases, the inferior song was in total downloaded more often after the inversion, but the superior song had already become the most popular song among the latest participants in the experiment. That is, the good songs were also already dominating the bad songs in popularity among later subjects in these cases; the cumulative counts displayed in panel 2, which include the earlier subjects, just had not caught up yet. For example, even though song pair 1 is in one case placed above the horizontal 50% line in panel 2, while in the other case it is placed below it, in both worlds it was already downloaded more often than its popularized counterpart among the second half of subjects (namely, 75 downloads of the best song vs. 74 downloads of the worst song in one world and 99 vs. 58 downloads in the other world). Thus, even though the best song had to overcome the extreme popularity setback of being displayed last at the bottom of a scrollable screen instead of on top, it had also already regained its natural dominance by the end of the experiment.

I finish this review of past studies by reanalyzing the data of perhaps the best-known sociological study of social influence (Salganik et al. 2006), which aimed to test the hypothesis that social influence renders success in cultural markets unpredictable. Salganik et al. (2006) conducted two experiments similar to the one reported in Salganik and Watts (2008) with the key difference that no inversion took place. In both experiments, eight social influence worlds all presented the same 48 songs along with true counts of song downloads by previous subjects in the same world. Because, as in Salganik and Watts (2008), download counts only included downloads made by subjects in the same world, the eight worlds evolved independently of one another. There was again also an independent condition in which no download counts were shown, which served to provide a measure of song quality. There are two differences between experiments 1 and 2: (1) songs in experiment 1 were shown in a 3×16 grid, while songs in experiment 2 were shown in a single list where subjects had to scroll down to see the end of the list (as in Salganik and Watts 2008); and (2) songs in the social influence worlds of experiment 1 were presented in random order, while in the social influence worlds of experiment 2 they were listed in descending order of popularity. In both experiment 1 and experiment 2, Salganik et al. found that the eight worlds exhibited different popularity rankings by the end of the experiment.

To evaluate whether the relationship between popularity and quality grew stronger during the experiments, as would be the case in a self-correcting dynamic, I split downloads in each world chronologically into two equal halves. I compare the rank correlations between the social influence worlds and the independent world for the first half of song downloads with those for the second half of song downloads. Table 1 shows the results for both experiments 1 and 2. Correlations with the independent condition increased in all cases in experiment 1 and in all cases but one in experiment 2, indicating that worlds were moving toward a ranking that reflected the underlying distribution of quality. The exception of world 3 in experiment 2 is worth mentioning. Close inspection of the data shows that in this world, song 31, which was not particularly popular in other worlds, was downloaded more frequently in both halves than song 25, which was the most popular song in all other worlds. It is impossible to tell whether this is signal or noise, but it should be noted as being in conflict with the self-correction claim.

Furthermore, it is true that the experiments were terminated when worlds were still somewhat different from one another, with some pairs of songs of similar quality being ranked one way in one world and the opposite way in another. This makes it impossible to conclude that worlds would have fully converged on a distribution of downloads according to quality. Nonetheless, it is possible to assess how far the self-correcting dynamic had progressed by

TABLE 1
CHANGES IN THE CORRELATION BETWEEN POPULARITY AND QUALITY IN EXPERIMENTS 1 AND 2 REPORTED IN SALGANIK ET AL. (2006)

World	Experiment 1			Experiment 2		
	First Half	Second Half	Change	First Half	Second Half	Change
1	.58	.74	+.18	.56	.72	+.16
2	.65	.80	+.15	.66	.72	+.06
3	.61	.79	+.18	.52	.46	−.06
4	.57	.80	+.23	.53	.74	+.21
5	.62	.81	+.19	.61	.73	+.12
6	.69	.73	+.04	.57	.68	+.11
7	.68	.80	+.12	.59	.64	+.05
8	.62	.74	+.12	.51	.61	+.10
Benchmark	.83	.79	−.04	.83	.75	−.08

NOTE.—Shown in the second and fifth (third and sixth) columns are the rank correlations between the first (second) half of song downloads in the eight social influence worlds and quality, averaged across 10 random samples. The fourth and seventh columns show the change from the first to the second half. The last row shows a benchmark correlation, which is calculated as the rank correlation between a random sample from the first or second half of downloads in the independent world and quality. The size of each benchmark sample equals the average number of downloads across the eight social influence worlds.

the end of the experiment. Table 1 additionally shows for both experiments the rank correlation between the total number of downloads in the independent condition (quality) and download totals in benchmark samples. These numbers show how strong correlations would be if subjects downloaded songs on the basis of quality only, as in the independent condition. The benchmark samples are constructed by randomly sampling downloads from the independent world. The size of each benchmark sample matches the average size of the social influence worlds. The benchmark correlations in both experiment 1 and experiment 2 decrease somewhat from half 1 to half 2. In the first half of each experiment, the rank correlations consistently lie well below the benchmark, showing that quality was an imperfect predictor of popularity. In the second half of each experiment, however, the correlations are nearly as high as the benchmark, showing that the self-correcting social influence dynamic was already close to completion. Initially popular songs of inferior quality had mostly lost their advantage. Together then these results suggest that under the social influence conditions created in the experiments, long-run success rankings are largely insensitive to initial conditions and rather predictable from quality.[2]

In each of the six studies I have reviewed, choice was significantly impacted by popularity. Yet our analysis shows that the percentage of choices for the popular option consistently fell short of the experimentally induced percentage of prior choices, giving rise to a self-correcting process in which superior alternatives became increasingly dominant. The evidence overall then is rather unanimous in finding social information capable of nudging fence-sitters to one side but unable to sustain a popularity advantage of an inferior alternative. Taking positive result bias for statistically significant influence effects into account further strengthens this conclusion.

CONVERGENCE IN A NEW EXPERIMENT

While each of the experiments I reviewed showed a self-correcting dynamic toward dominance of the superior alternative, in none of them was the choice process iterated to a state of convergence. Panels 1 and 2 in figure 1

[2] One may wonder why in the Salganik et al. (2006) experiments there were consistently initial discrepancies between quality and popularity only to be resolved in a later stage. Why did self-correcting forces not instead achieve concordance at once? The answer may lie in the number of choice alternatives in these experiments. Subjects faced with 48 previously unknown songs had too many options to try out, so they sampled a few songs and downloaded the ones they liked (Krumme et al. 2012). Early subjects, lacking social information, did not know what to sample, so many of them missed the best songs and could download only mediocre songs. Later subjects had good social information at their disposal so they mostly sampled popular songs, which tended to be of high quality. Social information was increasingly correlated with song quality, leading later subjects to download better songs.

show that even though later choices were consistently more in favor of the superior alternative than earlier choices, in many cases a majority of experimental subjects still chose the inferior option. This leaves open the possibility that the corrective process would have locked in at a point before dominance of the superior alternative. Moreover, populations may respond in nonmonotonic ways to social information signals, for example, if they are partly composed of individuals seeking to deviate from the majority. Therefore, even in cases where the majority of subjects was found to respond to a strong social information signal favoring the inferior option by choosing the superior alternative, we cannot infer that a weaker social information signal would also have produced superior dominance.

I conducted a novel large-scale internet experiment (details in the appendix; data and code in supplemental materials) in which I sought to reach a stable fraction of subjects choosing one alternative over the other. The experiment was carried out following a research protocol approved by the Stony Brook University's Committee on Research Involving Human Subjects (CORIHS no. 2015-3001-R1). The study was posted as a Human Intelligence Task (HIT) on the Amazon Mechanical Turk website (www.mturk.com), available to any interested subject of at least 18 years of age.

The Amazon Mechanical Turk subject population has well-known limitations. There are bots, some responders are hasty, as they seek to maximize profit at the expense of data quality, and most subjects come from the United States and India (Mason and Suri 2012). The gender composition is balanced, and the median age is about 30. The key advantage vis-à-vis earlier studies is that this subject population provides the scale necessary to allow a self-correcting dynamic process to converge.

Subjects who selected the HIT were routed to the study website. Upon informed consent, subjects were randomly assigned to one of the experimental conditions, which determined the mode in which questions were presented to them. Subjects were presented with the seven binary choices between pairs of cultural objects involving pictures, song fragments, and texts shown in figure 2. The seven questions covered a range of contexts, from elections to arts to collective intelligence. The subjectivity of answers ranges from pure matters of taste (houseware, wallpaper, modern art, music) to intellective tasks (McGrath 1984, 61) in which a problem with a correct answer is solved (visual test) and hybrid cases that combine subjective preferences with real differences in suitability (elections) or technological promise (green ideas). We can expect social influence in the case of the intellective and hybrid tasks to more heavily rest on an informational mechanism than a normative one. The use of different tasks allows a more robust test of the self-correction thesis.

As before, I define an alternative's quality as the relative frequency by which it is chosen by independent decision makers. Alternatives were inten-

Category	Question	Alternative A	Alternative B
Houseware	Here are two designs of an item for the home. Please select the design that you like better.		
Wallpaper	Which of these tropical pictures would you prefer as wallpaper on your personal computer?		
Visual test	Is the majority of the rectangular area black or white? 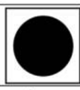	 Black	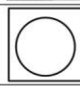 White
Modern art	Here are two works of art. Please select the one that you like best.		
Music	Here are fragments from two songs. Please select the fragment that you like best.		
Elections	Which of these two Hollywood actors would you rather see become U.S. president?	Jack Nicholson	Meryl Streep
Green ideas[a]	Here are two ideas for creating greener cities. Please select the idea that you like best.	Implement congestion pricing.	Plug ships into electricity grids so they don't idle in city ports, significantly reducing emissions.

FIG. 2.—Seven questions asked of subjects

NOTE.—A statistically significant majority chose B over A in the control condition ($P < .001$).

[a]This question was modified from a wiki survey reported in Salganik and Levy (2015).

tionally selected to be difficult to choose between, so that the subject population could be expected to be divided on the issues. Furthermore, the resulting high degree of uncertainty faced by subjects can be expected to increase the potential for social influence to impact their decision-making (Hedström 1998; Lynn et al. 2009; Azoulay, Stuart, and Wang 2013; Correll et al. 2017). Had I instead opted for pairs of alternatives that clearly differed in sensibility, veracity, or beauty, it would have been easier to confirm that the influence bias from social information does not typically lock in. Both the order in which choices were presented and the order in which the two competing alternatives were placed on the screen were randomized. Participants were asked to choose one of the two answers shown. Upon clicking their answers, participants were shown the next question. Subjects were paid one U.S. dollar upon completion of the HIT, regardless of the answers they chose or the speed with which they chose them. A total of 5,068 subjects participated.

In the control condition, subjects saw the two alternatives without information on their popularity among prior participants. For each pair, one alternative (A) was chosen significantly less often than the other (B) ($P < .001$), receiving between 31% and 41% of the vote (fig. 1, panel 2). By our definition of quality, B was therefore always the superior option. In the information condition, the alternatives were accompanied by popularity counts and bars visualizing the relative counts. I implemented an extreme social information signal in favor of inferior alternative A by choosing an artificial starting value of A's popularity count between 100 and 138 and a starting value for B between 7 and 9, yielding in each case an initial two-digit advantage of about 93%. To increase the chances that the information condition would reach a stable fraction of choices for A and B, I disproportionately routed subjects to that condition, with random assignment probabilities set to, respectively, 1/8 (control) and 7/8 (information). A total of 27,859 choices were recorded.

RESULTS

To assess the long-term dynamics of influence through social information I examined for each A-B pair A's average popularity as a percentage of all subject choices (fig. 1, panel 3) and A's evolving popularity over the course of the experiment in the information condition (fig. 3, panel 2). In the information condition, subjects were initially more likely to choose A than their counterparts in the control condition. Nonetheless, the percentage of choices for A fell far short of A's experimentally induced popularity advantage (fig. 1, panel 3). Consistent with the theory, this produced a self-correcting trajectory toward a majority B percentage. The first 100 subjects chose A on average about half the time, depending on the question (black circles in fig. 1, panel 3). The next 100 subjects, confronted with a less extreme social information signal in favor of A, chose A somewhat less often (gray circles in fig. 1, panel 3). The remaining subjects, confronted with mixed social information, consistently chose B over A (white circles in fig. 1, panel 3). Indeed, in all seven cases the initial 93% advantage for A was fully eliminated over the course of the experiment and converted into majority support for B (fig. 3, panel 2). These results reinforce the conclusion drawn from the earlier analysis of previous experiments that the influence effects of social information are generally insufficient for lock-in on an inferior option. In contrast to the previous experiments I reviewed, however, I reach this conclusion not through extrapolation from unconverged estimates (fig. 1, panels 1 and 2) but by observing the actual convergence of social influence dynamics over the course of many subject choices (fig. 3, panel 2). The finding that all trials, covering distinct choice contexts across which the degree of subjectivity of the right answer

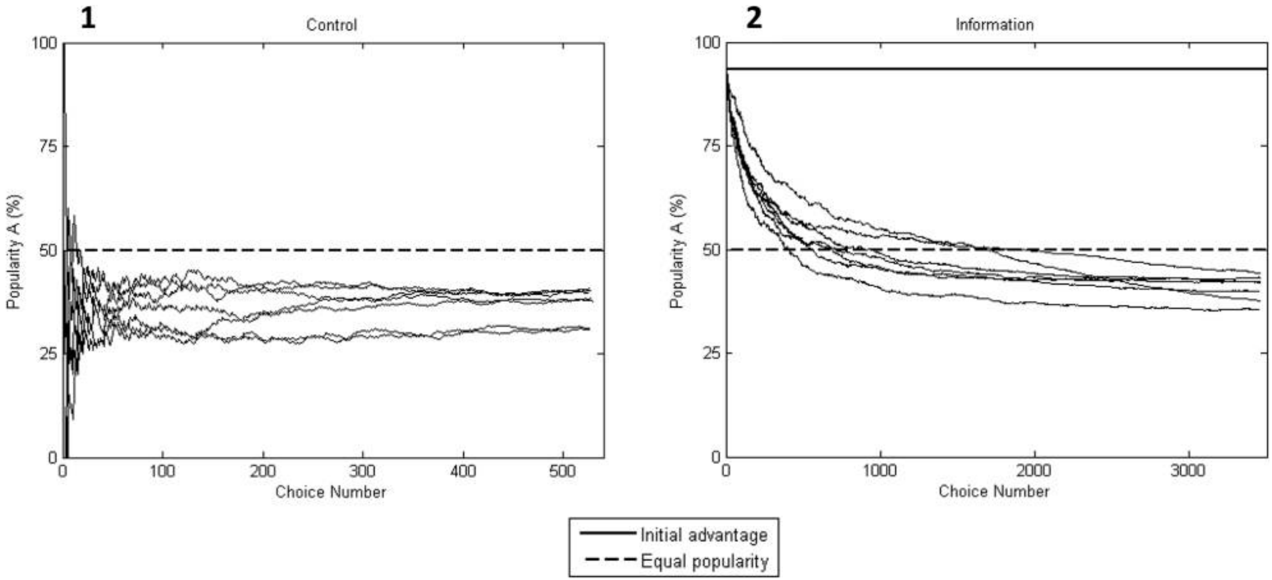

FIG. 3.—Popularity of choice alternative A over the course of the experiment. Panels 1 and 2 show results for, respectively, the control condition in which subjects were not informed about the choices of prior subjects and the social information condition in which they saw counts of prior choices for A and B. In each panel the vertical axis measures the cumulative percentage of choices for the inferior option over the superior alternative, while the horizontal axis measures choices in the order by which they were made. The dashed horizontal line indicates the 50% choice level above which the inferior alternative dominates in popularity. The solid horizontal line in panel 2 shows the popularity advantage given to inferior alternative A at the start of the social information trials.

varied, showed a self-correcting dynamic converging to dominance of the superior alternative suggests that the main result is robust across the choice contexts considered.

DISCUSSION

I conclude that the social influence effect of social information is often unable to decouple popularity from quality. Even extreme popularity advantages that are at odds with quality are resolved through a self-correcting process in which better alternatives gradually recover. The findings do not imply that the collective choices achieved in real life—the cultural practices that become standard, the celebrities that continue to draw public interest, the technologies that end up mass adopted, or the ideas that prevail—reflect a fixed majority preference or taste in the population. They may well be very different from the collective choices that would have resulted from the aggregation of independent decisions in the absence of social influence or from a slightly different start position. What the results instead suggest is that where outcomes are unpredictable, social information is not likely the responsible mechanism. The theoretical analysis shows that a self-sustaining feedback spiral generated by social information would require a strong degree of herding. Available evidence shows that various populations in diverse contexts when confronted with social information consistently exhibited a much weaker form of herding, with a large number of individuals refusing to consider inferior quality options just because of their popularity. In each of these cases this allowed a self-correcting process to propel superior quality alternatives from an initial popularity disadvantage back to popular dominance. This was even found in situations characterized by high uncertainty about quality, which are theoretically most prone to decoupling (Hedström 1998; Gould 2002; Podolny 2001; Lynn et al. 2009; Krumme et al. 2012; Azoulay et al. 2013; Lynn et al. 2016*a*; Lynn, Walker, and Peterson 2016*b*; Correll et al. 2017).

Lynn et al. (2009) make an important distinction between weak and strong social construction. Throughout the article I have exclusively focused on strong social construction: situations in which the popularity order of two alternatives does not match the quality order. The weak form characterized by Lynn et al. (2009) pertains to situations in which popularity exaggerates a quality difference but preserves the ordering of the choice alternatives. While I have refuted strong social construction across a number of experimental studies, the weaker form does occur in many of these studies. For example, Salganik et al. (2006) show that social influence increases inequality in market share, which confirms weak social construction. They also report that social influence rendered song success unpredictable. I do not dispute this result, as the measure of unpredictability used in their analysis captures the average difference in the quantity of success of a song across worlds, not in the success ranking of that song, and as such does not measure strong social construction. My main argument, however, does qualify this unpredictability claim: when social influence is driven by social information, long-run popularity tends to preserve an ordering based on underlying quality. Quality differences may be much smaller than emergent popularity differences suggest (weak social construction), but in this scenario superior alternatives ultimately become more popular than inferior alternatives (no strong social construction). Unpredictability measures that capture rank reversals, such as the rank correlation measures employed in the present manuscript, should converge to zero unpredictability.

In many real-world settings, populations of decision makers may still lock in on an inferior product, idea, or behavior, not because of imitation, but rather because structural forces sometimes heavily constrain individual choice. For example, the greater exposure of audiences to popular options and the inaccessibility of unpopular alternatives in many scenarios may naturally lead people to follow the crowd even in the absence of any inclination to herd. This happens in situations where individuals automatically learn about alternatives when they come in contact with others who previously chose it, such as in casual conversation with acquaintances or through diffusion in networked communication systems (Hedström 1994; Valente 1995; Strang and Soule 1998; Denrell and Le Mens 2007; van den Bulte and Joshi 2007; Centola 2010; Liu, King, and Bearman 2010; Aral and Walker 2012; Bond et al. 2012; DiMaggio and Garip 2012; Wang and Soule 2012; Banerjee et al. 2013; Lewis, Gray, and Meierhenrich 2014; Rossman 2014). These interaction patterns strongly favor the spread of already popular memes, enterprises, and commodities, greatly restricting access to potentially better alternatives. Another context is mass media environments where already popular individuals and topics receive disproportionate coverage (Tuchman 1973; Fishman 1980; Gans 1980; Lamont 1987; Oliver and Maney 2000; Andrews and Caren 2010; Seguin 2016) and discussions often focus on who or what is trending. Similarly, in science, the work of famous scholars is readily encountered in textbooks and invited lectures, and their highly cited articles are much more easily discovered than the work of their obscure counterparts. In the appendix I present results of an additional condition I ran in which besides providing social information I also introduced a strong form of such self-reinforcing exposure: subjects were exposed to alternatives in proportion to their popularity. Most decision makers in this condition simply could not choose the superior option because they were not exposed to it due to its lack of popularity. In this condition in which individual choice is highly constrained by the availability of alternatives we observe several cases of lock-in on an inferior option.

There are also less constrained situations in which more popular alternatives enjoy just a weak exposure advantage from being more prominently displayed. Bestseller lists and music charts have been found to produce a modest uptick in sales by further increasing exposure to already widely known books and songs (Sorenson 2007; Hosanagar et al. 2014). Alternatively, in e-markets, in news coverage, on crowdfunding and petition websites, and in knowledge-sharing communities it is common practice to present large number of options together with up-to-date popularity quantifiers sorted in descending order. Individuals may then preferentially sample popular choices but have ready access to unpopular alternatives. Experiment 2 reported in Salganik et al. (2006) and Salganik and Watts (2008) that we reviewed presented songs in such a dynamically sorted list, leading better-placed songs to be listened to more frequently. However, the near-equal exposure associated with neighboring placements allowed better songs to receive more downloads and gain on initially more popular songs and swap places. The tendency for discrepancies between popularity and quality to be ultimately corrected in these prior experiments suggests that more generally popular lists may often succeed at directing audiences toward superior quality options, despite the significant social influence effects identified in field experiments across a range of such platforms (Hanson and Putler 1996; Muchnik et al. 2013; van de Rijt et al. 2014; Margetts et al. 2015).

In settings in which feedback is too weak to sustain long-run dominance of an inferior object so that the discordance between status and quality is ultimately resolved, there may still be a lengthy out-of-equilibrium period that presents significant profit and loss possibilities. To paraphrase a quote by Shilling (1993), suggested to me by an anonymous reviewer, that is commonly attributed to John Maynard Keynes, “Markets can remain irrational a lot longer than you and I can remain solvent.” My analysis thus suggests that on the one hand a powerful actor or media organization could not manufacture consent (Lippmann 1922) simply by broadcasting an initial position and letting a social influence process driven by a social information mechanism run its course. On the other hand, it leaves open the possibility that those who are in a position to manipulate social information can profit by producing a temporary perturbation in collective behavior (Salganik and Watts 2008; Muchnik et al. 2013; van de Rijt et al. 2014).

Lock-in on an inferior alternative may also come about under strong social contingencies in the value of choices, as in bank runs, the adoption of technology with increasing returns, and high-stakes coordination problems (Arthur 1989; Corten and Buskens 2010; DiMaggio and Garip 2012; Correll et al. 2017) where individuals may face insurmountable barriers to deviating from the majority. Yet differently, costly institutional or peer sanctions against deviance may perpetuate an inefficient norm or behavior (Centola, Willer, and Macy 2005; Willer, Kuwabara, and Macy 2009; DellaPosta and Nee 2017).

What each of these lock-in scenarios has in common is that the widespread adoption of an inferior choice occurs not because people readily follow the masses but as a result of structural factors that constrain independent decision-making. The present analysis therefore does not suggest that social realities should be viewed as the inevitable product of objective truths, innate qualities, or universal preferences, with self-correcting feedback dynamics ultimately recovering their dominance from any perturbed start. Instead it suggests that where outcomes from choice processes are unpredictable, this unpredictability will often not be a consequence of an overwhelming urge of individuals to take cues from others about what to consider or what to like. The origins of lock-in in culture and markets should predominantly be sought in structural constraints blocking free individuals from trying out less popular alternatives, leaving them little choice but to follow the pack, thereby creating the illusion of herd mentality to an outside observer.

APPENDIX

Details of the Experimental Design

Construction of Choice Alternatives

To minimize the impact of influence processes prior to the experiment I only used objects that were fictive (e.g., movie stars running for U.S. president), were unknown (e.g., art works from unknown artists), or required expertise to judge (e.g., environmental measures). To maximize the possibility for social influence to produce decoupling, I deliberately constructed pairs of alternatives with no obvious preference order. It was necessary to gauge before the start of the experiment which alternative would probably be chosen by the majority of subjects in the control condition, because this alternative had to be given a popularity disadvantage in the initial counts in the information condition. To this end I conducted a pilot study between April 3, 2015, and April 13, 2015, and 1,005 subjects completed the pilot study. Subjects were asked a total of 12 questions, one question per screen, in random order, without popularity information. For five of these questions I asked subjects to rank five alternatives from 1 (most preferred) to 5 (least preferred). Based on subjects' answers I then selected for our main study a pair of alternatives from among these five that were closest in rank—thus maximizing the potential for social influence to impact choice in the main study—yet still statistically distinguishable. In the remaining seven cases I gave subjects one of three versions of a yes/no question with answers of different degrees of certainty ranging from 1 (definitely) to 2 (probably) to 3 (probably not) to 4 (definitely not) across which some feature was varied. From among the three versions I then chose for our main study the question that produced the largest minority—again maximizing the potential for social influence to impact choice in the main study—while the percentages choosing 1 or 2 were still statistically different from the percentages choosing 3 or 4, so that inferior and superior could be established. For seven of the 12 questions the majority choice in the control condition in the main study replicated the majority found in the pilot study and was statistically significant ($P < .001$ in a one-sample two-tailed proportion test) so that the modal population preference could be established with near certainty. These are the seven cases analyzed in the main text.

Experimental Conditions

The main text discusses two conditions to which subjects were randomly assigned. I collected data on two other conditions to which 1/5 of subjects were assigned. The "information without advantage" condition was identical to the information condition except that no initial advantage was given to alternative A. Both alternatives started at zero choices. In the "information and exposure" condition, choices were also presented with popularity counts and bars and the same initial advantage was given as in the information condition, but this time exposure to alternatives was limited to the options referred by network neighbors. I embedded subjects in a social network through a uniform attachment process (Krapivsky, Redner, and Leyfraz 2000; Fotouhi and Rabbat 2013): each consecutive node was given two incoming network ties from two nodes sampled uniformly at random from among all preceding nodes, including the fictive subjects representing the starting values. The focal node was exposed to the alternatives chosen by these two prior nodes along with counts of prior choices. In cases where the two sampled alternatives were both A or both B instead of presenting a subject a choice with a guaranteed outcome I automated the choice. After subjects' choices, popularity counts were updated accordingly, and a new node was linked to two prior nodes. This implementation of the proportional exposure regime ensures that exposure is exactly proportional to popularity, while at the same time ensuring that every subject always makes a binary choice between A and B, rendering subjects' experiences across conditions identical in all respects but the availability of popularity information.

Data Analysis

Subjects Choices by Condition

Table A1 shows the percentages of subjects choosing alternative A over B, by condition. In the information condition the percentage of A choices is below 50%, leading superior alternative B to gain popular dominance. Only in the information and exposure condition in the case of elections and green ideas does the percentage of A choices exceed 50%. This produces lock-in in these two cases as the proportion of subjects that gets the option to choose B shrinks toward 0.

TABLE A1
Percentage of Subjects Choosing A over B, by Experimental Condition

	Percentage of Subjects Choosing A instead of B			
Category	Control	Information without Advantage (Not in Main Text)	Information	Information and Exposure
Houseware	38.3	37.0	37.9	42.4
Wallpaper	40.2	44.8	41.1	41.3
Visual test	31.2	33.1	33.7	40.8**
Modern art	37.6	43.5	40.2	34.6
Music	40.7	39.3	40.3	43.4
Elections	39.6	37.4	42.6	53.6***
Green ideas	31.0	23.5**	35.5*	55.1***
No. of observations	529	534	3,457	550

Note.—Asterisks indicate results of a two-sided exact test for a difference with the control condition.
* $P < .05$.
** $P < .01$.
*** $P < .001$.

REFERENCES

Allison, Paul, Scott J. Long, and Ted K. Krauze. 1982. "Cumulative Advantage and Inequality in Science." *American Sociological Review* 47:615–25.

Andersen, Lisa R., and Charles A. Holt. 2008. "Information Cascade Experiments." *Handbook of Experimental Economics Results* 1:335–43.

Andrews, Kenneth T., and Neal Caren. 2010. "Making the News: Movement Organizations, Media Attention, and the Public Agenda." *American Sociological Review* 75:841–66.

Aral, Sinan, and Dylan Walker. 2012. "Identifying Influential and Susceptible Members of Social Networks." *Science* 337:337–41.

Arthur, W. Brian. 1989. "Competing Technologies, Increasing Returns, and Lock-In by Historical Events." *Economic Journal* 99:116–31.

Asch, Solomon E. 1951. "Effect of Group Pressure upon the Modification and Distortion of Judgments." Pp. 177–90 in *Groups, Leadership and Men: Research in Human Relations*, edited by H. Guetzkow. Pittsburgh: Carnegie Press.

———. 1956. "Studies of Independence and Conformity: I. A Minority of One against a Unanimous Majority." *Psychological Monographs: General and Applied* 70 (9): 1.

Azoulay, Pierre, Toby Stuart, and Yanbo Wang. 2013. "Matthew: Effect or Fable?" *Management Science* 60 (1): 92–109.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107:797–817.

Banerjee Abhijit V., Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Micro-Finance." *Science* 341:1236498.

Barabási, Albert-László, and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286 (5439): 509–12.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades.” *Journal of Political Economy* 100:992–1026.

Bikhchandani, Sushil, and Sunil Sharma. 2001. “Herd Behavior in Financial Markets.” *IMF Staff Papers* 47 (3): 279–310.

Bond, Rod, and Peter B. Smith. 1996. “Culture and Conformity: A Meta-Analysis of Studies Using Asch’s (1952b, 1956) Line Judgment Task.” *Psychological Bulletin* 119 (1): 111.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D. Kramer, Cameron Marlow, Jamie E. Settle, and James H. Fowler. 2012. “A 61-Million-Person Experiment in Social Influence and Political Mobilization.” *Nature* 489:295–98.

Centola, Damon. 2010. “The Spread of Behavior in an Online Social Network Experiment.” *Science* 329:1194–97.

Centola, Damon, and Michael Macy. 2007. “Complex Contagions and the Weakness of Long Ties.” *American Journal of Sociology* 113:702–34.

Centola, Damon, Robb Willer, and Michael W. Macy. 2005. “The Emperor’s Dilemma: A Computational Model of Self-Enforcing Norms.” *American Journal of Sociology* 110:1009–40.

Cialdini, Robert B., and Noah J. Goldstein. 2004. “Social Influence: Compliance and Conformity.” *Annual Review of Psychology* 55:591–621.

Correll, Shelley J., Cecilia L. Ridgeway, Ezra W. Zuckerman, Sharon Jank, Sara Jordan-Bloch, and Sandra Nakagawa. 2017. “It’s the Conventional Thought That Counts: How Third-Order Inference Produces Status Advantage.” *American Sociological Review* 82:297–327.

Corten, Rense, and Vincent Buskens. 2010. “Co-Evolution of Conventions and Networks: An Experimental Study.” *Social Networks* 32:4–15.

Davis, Colin J., Jeffrey S. Bowers, and Amina Memon. 2011. “Social Influence in Televised Election Debates: A Potential Distortion of Democracy.” *PLoS ONE* 6:e18154.

DellaPosta, Daniel, and Victor G. Nee. 2017. “Endogenous Dynamics of Institutional Change.” *Rationality and Society* 29 (1): 1–44.

Denrell, Jerker, and Chengwei Liu. 2012. “Top Performers Are Not the Most Impressive When Extreme Performance Indicates Reliability.” *Proceedings of the National Academy of Sciences of the USA* 109:9331–36.

Denrell, Jerker, and Gail Le Mens. 2007. “Interdependent Sampling and Social Influence.” *Psychological Review* 114:398–422.

DiMaggio, Paul, and Filiz Garip. 2011. “How Network Externalities Can Exacerbate Intergroup Inequality.” *American Journal of Sociology* 116:1887–1933.

———. 2012. “Network Effects and Social Inequality.” *Annual Review of Sociology* 38:93–118.

DiPrete, Thomas, and Gregory M. Eirich. 2006. “Cumulative Advantage as a Mechanism for Inequality.” *Annual Review of Sociology* 32:271–97.

Elster, Jon. 1989. *Nuts and Bolts for the Social Sciences*. Cambridge: Cambridge University Press.

Fishman, Mark. 1980. *Manufacturing the News*. Austin: University of Texas Press.

Fotouhi, Babak, and Momeni G. Rabbat. 2013. “Network Growth with Arbitrary Initial Conditions: Degree Dynamics for Uniform and Preferential Attachment.” *Physics Review E* 88:062801.

Frey, Vincenz, and Arnout van de Rijt. 2016. “Arbitrary Inequality in Reputation Systems.” *Scientific Reports* 6:38304.

Friend, Ronald, Yvonne Rafferty, and Dana Bramel. 1990. “A Puzzling Misinterpretation of the Asch ‘Conformity’ Study.” *European Journal of Social Psychology* 20 (1): 29–44.

Gans, Herbert J. 1980. *Deciding What’s News: A Study of CBS Evening News, NBC Nightly News, Newsweek and Time*. New York: Vintage.

Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey. 2007. "Self-Correcting Information Cascades." *Review of Economic Studies* 74:733–62.

Goldstone, Jack A. 1998. "Initial Conditions, General Laws, Path Dependence, and Explanation in Historical Sociology." *American Journal of Sociology* 104:829–45.

Gonzalez-Vaillant, Gabriela, Juhi Tyagi, Idil Afife Akin, Fernanda Page Poma, Michael Schwartz, and Arnout van de Rijt. 2015. "A Field-Experimental Study of Emergent Mobilization in Online Collective Action." *Mobilization* 20 (3): 281–303

Gould, Roger V. 2002. "The Origins of Status Hierarchies: A Formal Theory and Empirical Test." *American Journal of Sociology* 107:1143–78.

Granovetter, Mark S. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420–42.

Hanson, Ward A., and Daniel S. Putler. 1996. "Hits and Misses: Herd Behavior and Online Product Popularity." *Marketing Letters* 7:297–305.

Hedström, Peter. 1994. "Contagious Collectives: On the Spatial Diffusion of Swedish Trade Unions, 1890–1940." *American Journal of Sociology* 99:1157–79.

———. 1998. "Rational Imitation." In *Social Mechanisms: An Analytical Approach to Social Theory*, edited by P. Hedström and R. Swedberg. Cambridge: Cambridge University Press.

Hirshleifer, David, and Siew H. Teoh. 2003. "Herd Behavior and Cascading in Capital Markets: A Review and Synthesis." *European Financial Management* 9:25–66.

Hosanagar, Kartik, Daniel M. Fleder, Dokyun Lee, and Andreas Buja. 2014. "Will the Global Village Fracture into Tribes: Recommender Systems and Their Effects on Consumers." *Management Science* 60:805–23.

Hung, Angela A., and Charles R. Plott. 2001. "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." *American Economic Review* 91:1508–20.

Keuschnigg, Marc. 2015. "Product Success in Cultural Markets: The Mediating Role of Familiarity, Peers, and Experts." *Poetics* 51:17–36.

Krapivsky, Paul L., Sidney Redner, and Francois Leyfraz. 2000. "Connectivity of Growing Random Networks." *Physics Review Letters* 85:4629–32.

Krumme, Coco, Manuel Cebrian, Galen Pickard, and Sandy Pentland. 2012. "Quantifying Social Influence in an Online Cultural Market." *PLoS ONE* 7(5):e33785.

Lamont, Michèle. 1987. "How to Become a Dominant French Philosopher: The Case of Jacques Derrida." *American Journal of Sociology* 93:584–622.

Lewis, Kevin, Kurt Gray, and Jens Meierhenrich. 2014. "The Structure of Online Activism." *Sociological Science* 1:1–9.

Lippmann, Walter. 1922. *Public Opinion*. New York: Harcourt, Brace and Company.

Liu, Ka-Yuet, Marissa King, and Peter S. Bearman. 2010. "Social Influence and the Autism Epidemic." *American Journal of Sociology* 115:1387–1434.

Lorenz, Jan, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. "How Social Influence Can Undermine the Wisdom of Crowd Effect." *Proceedings of the National Academy of Science of the USA* 108:9020–25.

Lynn, Freda B., Joel Podolny, and Lin Tao. 2009. "A Sociological Deconstruction of the Relationship between Status and Quality." *American Journal of Sociology* 115:755–804.

Lynn, Freda B., Brent Simpson, Mark H. Walker, and Colin Petersen. 2016*a*. "Why Is the Pack Persuasive? The Effect of Choice Status on Perceptions of Quality." *Sociological Science* 3:239–63.

Lynn, Freda B., Mark H. Walker, and Colin Peterson. 2016*b*. "Is Popular More Likeable? Choice Status by Intrinsic Appeal in an Experimental Music Market." *Social Psychology Quarterly* 79 (2): 168–80.

Manzo, Gianluca, and Delia Baldassarri. 2015. "Heuristics, Interactions, and Status Hierarchies." *Sociological Methods and Research* 44:329–84.

Margetts, Helen, Peter John, Tobias Escher, and Stéphane Reissfelder. 2011. "Social Information and Political Participation on the Internet: An Experiment." *European Political Science Review* 3:321–44.

Margetts, Helen, Peter John, Scott Hale, and Taha Yasseri. 2015. *Political Turbulence: How Social Media Shape Collective Action*. Princeton, N.J.: Princeton University Press.

Mason, Winter, and Siddharth Suri. 2012. "Conducting Behavioral Research on Amazon's Mechanical Turk." *Behavioral Research Methods* 44:1–23.

McGrath, Joseph E. 1984. *Groups: Interaction and Performance*. Englewood Cliffs, N.J.: Prentice-Hall.

Merton, Robert K. 1968. "The Matthew Effect in Science." *Science* 159:56–63.

Muchnik, Lev, Sinan Aral, and Sean J. Taylor. 2013. "Social Influence Bias: A Randomized Experiment." *Science* 341:647–51.

Oliver, Pamela E., and Gregory M. Maney. 2000. "Political Processes and Local Newspaper Coverage of Protest Events: From Selection Bias to Triadic Interactions." *American Journal of Sociology* 106:463–505.

Podolny, Joel M. 2001. "Networks as the Pipes and Prisms of the Market." *American Journal of Sociology* 107:33–60.

Rossman, Gabriel. 2014. "The Diffusion of the Legitimate and the Diffusion of Legitimacy." *Sociological Science* 1:49–69.

Salganik, Matthew J., Peter S. Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311:854–57.

Salganik, Matthew J., and Karen E. C. Levy. 2015. "Wiki Surveys: Open and Quantifiable Social Data Collection." *PLoS ONE* 10:e0123483.

Salganik, Matthew J., and Duncan J. Watts. 2008. "Leading the Herd Astray: An Experimental Study of Self-Fulfilling Prophecies in an Artificial Cultural Market. *Social Psychology Quarterly* 71:338–55.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: Norton.

Seguin, Charles. 2016. "Cascades of Coverage: Dynamics of Media Attention to Social Movement Organizations." *Social Forces* 94:997–1020.

Shilling, A. Gary. 1993. "Scoreboard." *Forbes* 151 (4): 236.

Sinatra, Roberta, Dashun Wang, Pierre Deville, Chaoming Song, and Albert-László Barabási. 2016. "Quantifying the Evolution of Individual Scientific Impact." *Science* 354:aaf5239.

Sorensen, Alan T. 2007. "Bestseller Lists and Product Variety." *Journal of Industrial Economics* 55:715–38.

Strang, David, and Sarah A. Soule. 1998. "Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills." *Annual Review of Sociology* 24 (1): 265–290.

Tuchman, Gaye. 1973. "Making News by Doing Work: Routinizing the Unexpected." *American Journal of Sociology* 79:110–31.

Valente, Thomas W. 1995. *Network Models of the Diffusion of Innovations*. Cresskill, N.J.: Hampton.

van den Bulte, Christophe, and Yogesh V. Joshi. 2007. "New Product Diffusion with Influentials and Imitators." *Marketing Science* 26:400–421.

van de Rijt, Arnout, Idil Afife Akin, Robb Willer, and Matthew Feinberg. 2016. "Success-Breeds-Success in Collective Political Behavior." *Sociological Science* 3:940–50.

van de Rijt, Arnout, Soong Moon Kang, Michael Restivo, and Akshay Patil. 2014. "Field Experiments of Success-Breeds-Success Dynamics." *Proceedings of the National Academy of Sciences of the USA* 111:6934–39.

van de Rijt, Arnout, Eran Shor, Charles Ward, and Steven Skiena. 2013. "Only 15 Minutes? The Social Stratification of Fame in Printed Media." *American Sociological Review* 78 (2): 266–89.

Wang, Dan J., and Sarah Soule. 2012. "Social Movement Organizational Collaboration: Networks of Learning and the Diffusion of Protest Tactics, 1960–1995." *American Journal of Sociology* 117:1674–1722.

Wang, Dashun, Chaoming Song, and Albert-László Barabási. 2013. "Quantifying Long-Term Scientific Impact." *Science* 342:127–32.

Watts, Duncan J. 2007. "The Collective Dynamics of Belief." Pp. 241–72 *On Capitalism*, edited by V. Nee and R. Swedberg. Redgewood, Calif.: Stanford University Press.

Willer, Robb, Ko Kuwabara, and Michael W. Macy. 2009. "The False Enforcement of Unpopular Norms." *American Journal of Sociology* 115 (2): 451–90.

Zuckerman, Ezra W. 2012. "Construction, Concentration, and (Dis) Continuities in Social Valuations." *Annual Review of Sociology* 38:223–45.