

Field experiments of success-breeds-success dynamics

Arnout van de Rijt^{a,b,1}, Soong Moon Kang^c, Michael Restivo^d, and Akshay Patil^e

Departments of ^aSociology and ^cComputer Science, State University of New York, Stony Brook, NY 11794; ^bInstitute for Advanced Computational Science, Stony Brook, NY 11794; ^dDepartment of Management Science and Innovation, University College London, London WC1E 6BT, United Kingdom; and ^eDepartment of Sociology, State University of New York, Geneseo, NY 14454

Edited by Karen S. Cook, Stanford University, Stanford, CA, and approved March 28, 2014 (received for review September 10, 2013)

Seemingly similar individuals often experience drastically different success trajectories, with some repeatedly failing and others consistently succeeding. One explanation is preexisting variability along unobserved fitness dimensions that is revealed gradually through differential achievement. Alternatively, positive feedback operating on arbitrary initial advantages may increasingly set apart winners from losers, producing runaway inequality. To identify social feedback in human reward systems, we conducted randomized experiments by intervening in live social environments across the domains of funding, status, endorsement, and reputation. In each system we consistently found that early success bestowed upon arbitrarily selected recipients produced significant improvements in subsequent rates of success compared with the control group of nonrecipients. However, success exhibited decreasing marginal returns, with larger initial advantages failing to produce much further differentiation. These findings suggest a lesser degree of vulnerability of reward systems to incidental or fabricated advantages and a more modest role for cumulative advantage in the explanation of social inequality than previously thought.

Matthew effect | preferential attachment | scale-free networks | rich-get-richer effects | power law

Social scientists have long debated why we often see similar persons experience diverging trajectories of accomplishment, with some accumulating long strings of successes and others failing repeatedly. One explanation is that subtle variation along hard-to-observe dimensions of ability equips individuals with unequal a priori chances that gradually are revealed through differential achievement (1–5). A competing hypothesis states that “success breeds success” (4, 6–12). This hypothesis claims that the ultimate success of select persons may be born out of small, random initial advantages that grow ever larger through runaway positive feedback. Such cumulative advantage has been argued to produce significant, and arbitrary, inequality in many domains of human achievement (12–15). These two theoretical positions on the origins of societal inequities regularly meet in academic and public debate about whether great success is an accurate indicator of great talent (1–3, 5, 16–19).

Determining the origins of success in empirical studies is made difficult by the confounding of exogenous factors with endogenous processes. For instance, although some scholars have taken the extreme variance of success distributions as a tell-tale sign of cumulative advantage (8, 11, 16, 20), critics have pointed out that various other generative mechanisms, such as the existence of a convex correspondence between fitness and success (21, 22), can generate the same empirical regularities (17, 23–28). Further, in longitudinal records of success, unobserved dimensions of fitness generate apparent bias toward past winners (3, 4, 12, 15). In these cases, the higher success rates of talented, privileged, and well-connected individuals give rise to temporal correlations between successes, which may be erroneously interpreted as a causal effect of past on future success.

This problem of empirical confounding may be overcome through randomized experiments. Prior studies have used experimental methods to identify positive social feedback (13, 29–33). While these studies confirm the operation of reinforcement processes, they provide limited insight into the degree to which

these processes distort the allocation of resources to individuals in various reward systems. First, the success-breeds-success hypothesis covers a much wider variety of types of success than previous experiments have investigated. In this paper we evaluate the presence of cumulative advantage by consistently applying the same experimental intervention across a diverse range of reward systems. The systems we study vary in the degree to which the rewards transferred carry immediate monetary value, affect the social status of recipients, or are of entirely ideological nature. Second, the degree to which cumulative advantage can disrupt meritocracies depends critically on whether greater initial advantages breed proportionately greater amounts of subsequent success. In our experiments we systematically vary the magnitude of the initial advantage to quantify the marginal effects on the size of the ultimate success gap.

We constructed an experimental design in which we explicitly control the allocation of success (*Materials and Methods*). In this setup, we bestow early successes upon randomly selected members of a population, thereby ensuring that the expectations of success before intervention are equal for recipients and nonrecipients. To allow a robust test of cumulative advantage in multiple contexts, we deployed this design in four naturally occurring systems, representing distinct forms of personal success—financial gain, endorsement, social status, and social support. First, in the financial domain, we applied the design to the crowdfunding website kickstarter.com, where creators of projects in the areas of technology, arts, and entertainment compete for donations from the general public. We sampled 200 new, unfunded projects and donated a percentage of the funding goal to 100

Significance

Social scientists have long debated why similar individuals often experience drastically different degrees of success. Some scholars have suggested such inequality merely reflects hard-to-observe personal differences in ability. Others have proposed that one fortunate success may trigger another, thus producing arbitrary differentiation. We conducted randomized experiments through intervention in live social systems to test for success-breeds-success dynamics. Results show that different kinds of success (money, quality ratings, awards, and endorsements) when bestowed upon arbitrarily selected recipients all produced significant improvements in subsequent rates of success as compared with the control group of nonrecipients. However, greater amounts of initial success failed to produce much greater subsequent success, suggesting limits to the distortionary effects of social feedback.

Author contributions: A.v.d.R., S.M.K., M.R., and A.P. designed research; A.v.d.R., S.M.K., M.R., and A.P. performed research; A.v.d.R., S.M.K., M.R., and A.P. analyzed data; and A.v.d.R. and S.M.K. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: arnout.vanderijdt@stonybrook.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1316836111/-DCSupplemental.

randomly chosen projects. Second, on the website epinions.com reviewers are paid for posting written evaluations of new products, and those evaluations subsequently are rated by website visitors as “very helpful,” “helpful,” “somewhat helpful,” or “not helpful.” Reviewers are paid more for evaluations that are considered more helpful. We sampled 305 new, unrated reviews that we evaluated as being very helpful and gave a random subset of these reviews a “very helpful” rating. Our third application involved the encyclopedia website wikipedia.org, where highly productive editors receive status awards from community members in recognition of their dedication (34). We sampled 521 editors who belonged to the top 1% of most productive editors and conferred an award to a randomly chosen subset of these editors. Fourth, on the petition website change.org individuals seek support from the general public for social and political goals through signature campaigns that can be signed electronically by any named or anonymous supporter. We sampled 200 early-stage campaigns and granted a dozen signatures to 100 randomly chosen petitions. In each experiment, we kept a daily record of subsequent donations, ratings, awards, and signatures given by third parties after the treatment in both the experimental and control condition. These four interventions thus represent a range of types of success, covering resource transfers in which both source and recipients are financially affected (kickstarter.com), transfers in which the recipient benefits materially without the source incurring a cost (epinions.com), conferrals of social status (wikipedia.com), and expressions of ideological support (change.org).

Results

To isolate the effect of our experimentally induced success on the rate of success accumulation in each study, net of any exacerbating or counteracting second-order effects that successive successes may have had on one another, we first calculated separately for both the experimental and control conditions the proportion of individuals who experienced at least one more success during the observation period. In all four domains, the experimental treatment produced significant increases in rewards for the treated individuals. Fig. 1 shows that in each study the artificial contribution of success had a positive effect on the rate of success. In the control condition of the crowd-funding study, 39% of project initiators received subsequent funding by




Fig. 1. Percentage of cases with posttreatment success. From left to right: percentage of crowd-funding project creators who collected subsequent funding; percentage of reviewers who subsequently received positive ratings; percentage of Wikipedia editors who subsequently received awards from third parties; and percentage of petitioners whose petitions were subsequently signed by others. The difference between conditions at the end of the observation period is statistically significant for funding ($\chi^2 = 19.4$; $P = 0.000$), ratings ($\chi^2 = 9.54$; $P = 0.002$), awards ($\chi^2 = 4.72$; $P = 0.030$), and signatures ($\chi^2 = 4.05$; $P = 0.044$).

one or more donors. In contrast, 70% of the individuals in the experimental condition received contributions from third parties, indicating that the mere presence of an initial donation made recipients about twice as likely to attract funding ($\chi^2 = 19.4$; $P = 0.000$). In the endorsement study, the baseline likelihood of success was much higher than in the other studies, with 77% of untreated product reviews receiving at least one “very helpful” rating during the 14 d immediately following the treatment. This percentage rose further to 90% in the treatment condition ($\chi^2 = 9.54$; $P = 0.002$), and this increase did not come at the expense of a parallel increase in less favorable ratings (*SI Results*). In the control condition of the Wikipedia study, 31% of the editors received a status award during the observation period. In comparison, 40% of editors who received their first award through our experiment received one or more other awards from fellow editors ($\chi^2 = 4.72$; $P = 0.030$). Finally, in the signature study, 52% of the individuals in the control condition received at least one more signature toward their petition goal during our observation period, whereas 66% of petitioners in the experimental condition subsequently accumulated additional signatures ($\chi^2 = 4.05$; $P = 0.044$). The effect of signatures contributed through our experiment on petitioners’ yield of subsequent signatures suggests that social reinforcement effects are operative even for expressions of ideological support (35). Consistent with earlier results, we find a causal link between past and future rewards in four distinct substantive domains, providing robust evidence for positive feedback operating on arbitrary early advantages in the allocation of resources to individuals.

Cumulative Advantage Dynamics. To assess whether the effect of our treatment was only transient or instead had an enduring impact on success accumulation, we calculated in each system the average number of posttreatment successes accumulated as a function of time (Fig. 2). All the posttreatment measures of success shown exclude the success applied through the treatment. In each study the arbitrary gap in subsequent success between the recipients and nonrecipients of early success persisted throughout the study. In the funding study, our donation increased the average number of subsequent donations from 1.11 in the control condition to 2.49 in the experimental condition. This difference between conditions is statistically highly significant (signed-rank test; $z = 3.95$; $P = 0.000$). In the endorsement study, 14 d after our ratings were applied, the number of subsequent positive ratings given by third parties still differed significantly, with a total of 11.4 in the control condition and 14.9 in the experimental condition (rank-sum test; $z = 3.213$; $P = 0.001$). In the awards study, 1 mo after our intervention, editors in the control condition had accumulated noticeably fewer awards on their user pages by fellow editors than editors in the experimental condition ($z = 2.635$; $P = 0.008$), and this difference still remained noticeable after 3 mo ($z = 1.982$; $P = 0.048$), at 0.17 and 0.28 awards per person, respectively. Finally, at 2 wk after intervention in the signature study, Fig. 2 shows a small gap remaining between the posttreatment signature yields of campaigns in the control condition, which had accumulated another 1.74 signatures on average, and those in our experimental condition, which had recruited an average of 2.32 additional signatures ($z = 1.759$; $P = 0.079$). In combination, these findings indicate that, despite qualitative differences in the nature of success across the four reward systems, an early advantage consistently drives a sustained difference between individuals with equal initial likelihood of success.

Marginal Returns of Success. Although the impact of the initial advantage in each study demonstrates the susceptibility of these reward systems to arbitrary and self-reinforcing differentiation between ex ante equivalent individuals, it tells us little about the extent to which inequalities can be affected. If more sizeable




Fig. 2. The success-breeds-success effect over time. The curves represent running numbers of donations (blue), positive ratings (red), awards (yellow), and campaign signatures (green) in the experimental condition (solid lines) and the control condition (dashed lines). The horizontal axis is normalized so that 0 marks the time of experimental intervention, and 1 marks the end of the observation period. The vertical axis is normalized so that for each system a value of 1 equals the maximum across time and conditions.

initial differences were introduced between individuals, how much more severely would the subsequent allocation of resources be impacted?

To test the effects of larger initial endowments on cumulative returns, we subsequently varied the strength of the treatment in both the funding study and the rating study. In the funding study we included funding goals of up to \$5,000 and withheld a donation, donated 1% of the funding goal through one donor, or donated 4% of the funding goal through four separate donors. By holding the per-donor contribution level constant across treatment conditions, we neutralized any social influence effects that the size of the average prior contribution may exert on followers. In the rating study we again sampled previously unrated reviews, and when we found them to be of high quality, we left them unrated, rated them as “very helpful” once, or rated them as “very helpful” four times by four different raters.

Among subjects in the zero-donor condition, 32% attracted subsequent funding from one or more donors, whereas 74% of the subjects in the one-donor condition and 87% of the subjects in the four-donors condition collected subsequent funds. The difference between the one-donor condition and the control condition is statistically significant ($\chi^2 = 11.0$; $P = 0.001$), as is the difference between the four-donors condition and the control condition ($\chi^2 = 19.4$; $P = 0.000$). However, the increase in the size of the initial advantage as represented by the difference between the one-donor and four-donors conditions did not result in a significantly higher chance of one or more donations ($\chi^2 = 1.65$; $P = 0.199$). In the endorsements study, reviewers who wrote high-quality reviews but received no positive rating from us exhibited a 77% chance of receiving one or more positive ratings, compared with 90% of reviewers who received one positive review from us and 94% of reviewers who received four positive reviews from us. Again, the treatment effects are positive in both experimental conditions ($\chi^2 = 9.54$; $P = 0.002$ and $\chi^2 = 9.38$; $P = 0.002$) but do not differ from one another ($\chi^2 = 0.926$; $P = 0.336$). Together, these patterns of one-by-one comparisons between conditions suggest decreasing marginal returns: Each additional unit increase in input yields a progressively smaller increase in output. Indeed, in each experiment an increase in input from zero to one produces a significant increase in per-unit output, whereas the additional increase in input from one to four never yields a noticeable increase in per-unit output.

To quantify these marginal returns, we calculated average posttreatment success as a function of the number of successes

applied through treatment, shown in Fig. 3. Fig. 3A displays the average total dollar amount raised by the number of donations bestowed. Fig. 3B displays the average number of donations accumulated by the number of donations made. Fig. 3C displays the number of positive ratings received by the number of positive ratings experimentally bestowed. The averages reported in each panel exclude the dollars, donations, and ratings applied through our experimental intervention. Consistently across all panels, the average marginal returns of an increase from zero to one exceed the average marginal returns of an increase from one to four. The average return on a single donation (on average \$24.52) is \$191.00, but the additional three donations are estimated to bring in only \$89.57 each (Fig. 3A). Accordingly, the difference in the amount of dollars raised between the zero-donations and one-donation conditions is significant (signed-rank test; $z = 3.02$; $P = 0.003$), and so is the difference between the zero-donations and the four-donations conditions ($z = 3.61$; $P = 0.000$), but the one-donation and four-donations conditions do not deviate significantly ($z = 1.70$; $P = 0.090$). A single donation raises the number of subsequent third-party donations by 4.3, whereas each of the additional three donations brings in only 1.7 subsequent third-party donations (Fig. 3B). Indeed, the difference in the number of donations elicited in the zero-donations and one-donation conditions is significant ($z = 3.20$; $P = 0.001$), as is the difference between the zero-donations and four-donations conditions ($z = 4.16$; $P = 0.000$), whereas the difference between the one-donation and four-donations conditions falls just short of statistical significance ($z = 1.95$; $P = 0.051$). Finally, a single “very helpful” rating given to a product reviewer increases the number of subsequent third-party “very helpful” ratings by 3.48, but awarding an additional three positive ratings does not appear to increase the expected number of “very helpful” ratings further, as indicated by a slightly negative marginal effect of -0.43 (Fig. 3C). The difference between the zero-ratings and one-rating conditions is significant (rank-sum test; $z = 3.21$; $P = 0.001$), but the four-ratings condition differs from neither the zero-ratings condition ($z = 1.83$; $P = 0.067$) nor the one-rating condition ($z = 1.07$; $P = 0.288$).

Discussion

Our findings reveal the presence of a noticeable feedback effect in each of the distinct settings that we investigated, in that initial arbitrary endowments create lasting disparities in individual success. These results suggest that the inadvertent magnification

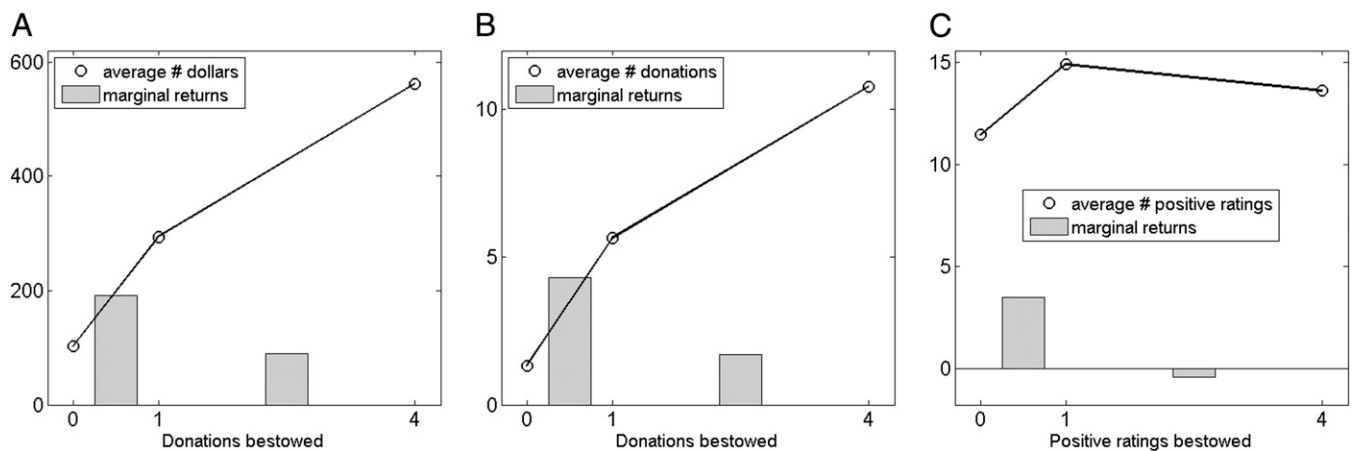


Fig. 3. Marginal returns of success. The horizontal axis measures the number of donations (in *A* and *B*) or ratings (in *C*) applied through experimental intervention, namely, none, one, or four. The circles measure the average dollar amount (*A*), number of donations (*B*), and positive ratings (*C*) obtained in each condition, excluding the treatment. Shaded bars measure the marginal returns, which are calculated as the slopes of the lines connecting the averages. In each panel, marginal returns decrease with the size of the treatment.

of arbitrary differences between individuals of comparable merit may be a common feature of many types of social reward systems. At the same time, our experimental demonstration of decreasing marginal returns to success may suggest bounds to the degree to which the natural allocation of resources can be disrupted by social feedback effects. Without a priori differentiation in quality or structural sources of advantage, cumulative advantage alone may not be able to generate the extreme kinds of runaway inequality that so commonly are attributed to it (4–16, 29, 30). The vulnerability of meritocracies to biases from success-breeds-success effects thus may be more limited than generally assumed.

The deliberate allocation of success in our experiments demonstrates that cascades of positive reinforcement can be initiated intentionally by a strategic actor. This form of purposive action presents the possibility of perverse effects, such as profit-seeking entities offering loans, positive reviews, and endorsements in exchange for the pecuniary equivalent of the anticipated ripple effect. It also raises the possibility of a philanthropic entity jumpstarting support through first-mover loans to underappreciated projects as a social policy instrument for counteracting nonmeritocratic disparities in populations. However, these opportunities for manipulation are offset by the decreasing marginal returns of success identified in our study, which suggests limits on the scale of such purposive intervention. Both the crowd-funding study and the rating study suggest that the reinforcement value of a single initial success is much larger than that of additional successes. The per-donor effect of a single donation by a single donor on fundraising success was greater than that of four donations by four separate donors. Similarly, a single positive rating of an unrated review increased the number of positive ratings by more than the increase resulting from each of the four positive ratings by four separate raters. Strategic contributions aimed at steering dynamics in a more positive direction thus may be less effective when made to campaigns that already have garnered some minimal degree of support. Hence, the susceptibility of reward systems to deliberate manipulation may be restricted mostly to interventions favoring those individuals who cannot muster any initial success otherwise.

Materials and Methods

The four studies (ID numbers 373335, 366647, 230771, and 442574) were approved by the Stony Brook University Human Subjects Committee and were conducted in compliance with the terms of use of kickstarter.com, epinions.com, wikipedia.org, and change.org.

Study 1. Kickstarter.com. Site. Kickstarter.com is a crowd-funding website launched in April 2009 for projects in diverse categories ranging from music, film, and video games to innovative products. The Kickstarter platform facilitates the gathering of monetary donations from the general public. Project creators choose a deadline and set a goal of raising a minimum amount of funds. If the chosen goal is not reached by the deadline, no funds are collected, and all donations are returned to the donors, or “backers.” Kickstarter takes 5% of the funds raised. Payments are made through the online retailer Amazon.com, which charges an additional 3–5%. As of January 13, 2014, 128,887 projects had been launched, with \$943 million raised (www.kickstarter.com/help/stats). For our experiments, we sampled from projects created by United States residents only.

Design and procedure. Fundraising goals in Kickstarter projects range from less than \$100 to more than \$1 million. To mitigate negative effects of this large variance on the statistical power of our study, we limited our sampling frame to projects with a low funding goal (\$1,000 in round 1 and \$5,000 in round 2) and matched projects across conditions according to their goal amount. To allow for comparative temporal analysis, we selected only projects that had to be funded within 28–30 d. Projects exhibit idiosyncratic funding behavior during their first 5–6 d, with some projects suddenly spiking in funding activity because of differentiated campaign efforts by project creators on social networking sites such as Facebook.com. To avoid this variability in the data, we selected only projects that had not yet been funded at 24 d before the end of the funding period. When a project fulfilling the selection criteria became available, we randomly assigned it to the experimental or the control condition. The next project with a similar goal amount that became available was assigned to the alternate condition. Because our donations were relatively small, the kickstarter.com website did not display the projects we invested in at a more prominent location than those in the control condition; thus we avoided a treatment effect caused by a difference in the visibility of the matched projects. We kept a daily record of donations until the funding deadline was reached. We collected a total of 293 projects in two rounds. Data collection took place between July 2012 and February 2014 for round 1 and between September 2013 and March 2014 for round 2.

Round 1. The treatment in round 1 of data collection involved the donation of 1% or 10% of a funding goal of up to \$1,000. No donations were made to projects in the control condition.

Round 2. The treatment in round 2 involved the donation of either 1% by one donor or a total of 4% by four donors of a funding goal up to \$5,000. No donations were made to projects in the control condition.

Study 2. Epinions.com. Site. Epinions.com is a popular general consumer review site that was established in 1999. In 2013, it had ~1 million visitors per month. The epinions platform gives users the ability to write reviews evaluating all types of products (www.epinions.com). Other users can rate reviews as being “very helpful,” “helpful,” “somewhat helpful,” or “not helpful.” The amount users are paid for writing reviews depends on the ratings their reviews receive from other users. [Epinions specifies that these earnings are determined by an undisclosed algorithm that provides

greater financial compensation to authors of better-rated reviews (www.epinions.com/help/faq/show/~faq_earnings.)

Design and procedure. When an unrated review became available, we read the review to determine its quality, classifying it as either “high” or “low.” We then randomly assigned high-quality reviews to either the experimental or control condition. High-quality reviews in the experimental condition were rated “very helpful,” and high-quality reviews in the control condition were left unrated. After 14 d, we counted the number of ratings and the type of ratings that each of the selected reviews had accumulated. We collected a total of 481 cases in two rounds. Data collection took place between October 2012 and August 2013 for round 1 and between September 2013 and January 2014 for round 2.

Round 1. In round 1 of data collection, the treatment involved the application of a single rating. High-quality reviews received a single “very helpful” rating in the experimental condition and were left unrated in the control condition.

Round 2. In round 2 of data collection the treatment involved the application of either a single rating or four ratings on behalf of four members of the research team. High-quality reviews in the experimental condition received one or four “very helpful” ratings. High-quality reviews in the control condition were left unrated.

Study 3. Wikipedia.org. Site. Wikipedia is a collaboratively written encyclopedia, started in the United States in 2001, that as of January 13, 2014 encompassed 31 million articles in 285 languages (<http://en.wikipedia.org/wiki/Wikipedia>About>). Wikipedia is created through voluntary contributors, or “editors,” who actively generate, update, and modify its content. The website poses very low barriers to participation by allowing any individual to edit its articles under a self-chosen pseudonym. For this research, we restricted our experiment to the English-language Wikipedia, which as of January 2014 had 4.4 million articles, 21 million registered accounts, and 118,082 active editors who had made at least one edit in the last month.

Design and procedure. On Wikipedia, editors can grant other editors a virtual award (<http://en.wikipedia.org/wiki/Wikipedia:Awards>) by posting such an award on the target editor’s user page for public display. To ensure that we acted in accordance with informal Wikipedia rules to grant awards only to individuals who have contributed significantly to the project, we sampled only top contributors. We ranked the population of editors by their total number of edits in the preceding month and included only editors in the top 1% by edit count. We also eliminated from our target population anonymous editors and editors with special authority within the project.

We randomly placed 208 individuals in the experiment condition and anonymously gave them an award, which editors in the broader Wikipedia editing community could see when making their own decisions about whom to thank or reward. To do so, we posted a customized award on the focal editor’s user page. In a prior experiment conducted 2 y before the present study, we used a similar study design but instead granted generic “barnstar” awards (<http://en.wikipedia.org/wiki/Wikipedia:Barnstars>) and focused on effects of these awards on productivity (34). The awards we gave out could be seen only on the editors’ user pages, and nowhere did Wikipedia

overview award recipients or sort them by popularity, so that recipients did not enjoy greater exposure than nonrecipients as a result of our treatment. After 90 d, we reviewed the editor pages for all 521 subjects in our sample and counted how many additional third-party awards an editor had received. Data collection spanned the period February 2012 to July 2012. Note that we collected data only on a sample of the target population, because the award data from digital historical records had to be collected manually (automatic data collection proved unreliable and error-prone).

Study 4. Change.org. Site. Change.org is a popular online petition website. Individuals who are passionate about a cause can initiate a signature campaign with only a few clicks after they have created an account with a functional e-mail and postal address. Visitors to the website can sign petitions either anonymously or by name. As of January 13, 2014 the website had hosted petitions initiated by more than 30 million people (www.change.org/). Petitions range in purpose from the mundane—“ban ‘x’ from our multiplayer gaming platform”—to topics that dominate the national news. These extreme discrepancies in the scope and importance of petitions create large natural variation in the potential for signatures. Indeed, most campaigns yield only a handful of sympathizers, but a few reach millions of supporters.

Design and procedure. We sampled from the population of new petitions whose creator had acquired at most 15 signatures. We could not sample petitions below the minimum of five signatures required for public posting on the website. The selection of small petitions prevented the petitions we studied from appearing on a list of popular petitions, thus precluding this sorting mechanism from mediating any treatment effect. We further selected petitions that had been initiated less than 14 d earlier, ensuring that the petitions had not lost their relevance. From among these recent petitions we selected those that had not been signed in the past 24 h, thus making sure that our treatment did not co-occur with an ongoing surge of petition signing. Two hundred petitions satisfied these criteria and passed a screening test against mal-intended campaigns that sought to do harm to an individual or group.

We randomly assigned 100 petitions to the experimental condition and 100 petitions to the control condition. We added 12 anonymous signatures to petitions in the experimental condition and withheld signatures from petitions in the control condition. We kept a daily record of the number of signatures the petitioners received for a period of 2 wk, after which interval most campaigns had stopped accumulating signatures. Data collection took place between July and August of 2012.

Further details on design and data analysis can be found in *SI Materials and Methods* and *SI Results*.

ACKNOWLEDGMENTS. We thank Damon Centola for helpful discussions, the editor and two anonymous reviewers for useful comments, and Idil Akin, Gabriella Gonzalez, Hua Mo, Fernanda Page, and Juhi Tyagi for assistance with data collection. This work was supported by National Science Foundation Grants SES-1340122 (to A.v.d.R.) and SES-1303522 (to M.R. and A.v.d.R.).

- Herrnstein RJ, Murray C (1994) *The Bell Curve: Intelligence and Class Structure in American Life* (Free, New York).
- Simonton D (1997) Creative productivity: A predictive and explanatory model of career trajectories and landmarks. *Psychol Rev* 104(1):66–89.
- Huber J (1998) Cumulative advantage and success-breeds-success: The value of time pattern analysis. *J Am Soc Inf Sci* 49(5):471–476.
- Allison P (1980) Estimation and testing for a Markov model of reinforcement. *Sociol Methods Res* 8(4):434–453.
- Mankiw NG (2013) Defending the one percent. *J Econ Persp* 27(3):21–34.
- Merton RK (1968) The Matthew effect in science. *Science* 159(3810):56–64.
- Barabási AL (2012) Network science: Luck or reason. *Nature* 489(7417):507–508.
- de Solla Price D (1976) A general theory of bibliometric and other cumulative advantage processes. *J Am Soc Inf Sci* 27(5):292–306.
- Reskin B (1977) Scientific productivity and the reward structure of science. *Am Sociol Rev* 42(3):491–504.
- Allison P, Long J, Krauze T (1982) Cumulative advantage and inequality in science. *Am Sociol Rev* 47(5):615–625.
- Barabási AL, Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512.
- DiPrete TA, Eirich GM (2006) Cumulative advantage as a mechanism for inequality. *Annu Rev Sociol* 32:271–297.
- Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762):854–856.
- Simcoe TS, Waguespack DM (2010) Status, quality and attention. *Manage Sci* 57(2):274–290.
- Petersen AM, Jung WS, Yang JS, Stanley HE (2011) Quantitative and empirical demonstration of the Matthew effect in a study of career longevity. *Proc Natl Acad Sci USA* 108(1):18–23.
- Simon H (1955) On a class of skew distribution functions. *Biomet* 42(3-4):425–440.
- Mandelbrot B (1959) A note on a class of skew distribution functions: Analysis and critique of a paper by H. A. Simon. *Inform. & Contr* 2(1):90–99.
- Murray C (2003) *Human Accomplishment: The Pursuit of Excellence in the Arts and Sciences, 800 BC to 1950* (HarperCollins, New York).
- Denrell J, Liu C (2012) Top performers are not the most impressive when extreme performance indicates unreliability. *Proc Natl Acad Sci USA* 109(24):9331–9336.
- Yule GU (1925) A mathematical theory of evolution. *Phil. Trans. B* 213:21–87.
- Rosen S (1981) The economics of superstars. *Am Econ Rev* 71(5):845–858.
- Frank RH, Cook PJ (1996) *The Winner-Take-All Society* (Free, New York).
- Newman ME (2005) Power laws, Pareto distributions and Zipf’s law. *Contemp Phys* 46(5):323.
- Huberman BA, Adamic LA (1999) Growth dynamics of the World-Wide-Web. *Nature* 401(6749):131.
- Papadopoulos F, Kitsak M, Serrano MA, Boguñá M, Krioukov D (2012) Popularity versus similarity in growing networks. *Nature* 489(7417):537–540.
- D’Souza RM, Borgs C, Chayes JT, Berger N, Kleinberg RD (2007) Emergence of tempered preferential attachment from optimization. *Proc Natl Acad Sci USA* 104(15):6112–6117.
- Stumpf MPH, Porter MA (2012) Mathematics. Critical truths about power laws. *Science* 335(6069):665–666.
- Adamic L (2011) Complex systems: Unzipping Zipf’s law. *Nature* 474(7350):164–165.
- Hanson WA, Putler DS (1996) Hits and misses: Herd behavior and online product popularity. *Mark Lett* 7(4):297–305.

30. Salganik MJ, Watts DJ (2008) Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Soc Psychol Q* 74(4):338–355.
31. Muchnik L, Aral S, Taylor SJ (2013) Social influence bias: A randomized experiment. *Science* 341(6146):647–651.
32. Margetts H, John P, Escher T, Reissfelder S (2011) Social information and political participation on the internet: An experiment. *Eur. Pol. Sc. Rev* 3(3):321–344.
33. Ginsburgh VA, van Ours JC (2003) Expert opinion and compensation: Evidence from a music competition. *Am Econ Rev* 93(1):289–296.
34. Restivo M, van de Rijt A (2012) Experimental study of informal rewards in peer production. *PLoS ONE* 7(3):e34358.
35. Centola D, Macy M (2007) Complex contagions and the weakness of long ties. *Am J Sociol* 113(3):702–734.