An experimental study of the generation of inequality and unpredictability in an artificial
academic literature market

Nicholas Fox^{1,2}, Steven Mattia², Nate Honeycutt², & Lee Jussim²

¹Center for Open Science, Charlottesville VA USA

²Rutgers, The State University of New Jersey, New Brunswick, NJ USA

Abstract

Academic publishing has changed substantially in the past 30 years due to the advent of the internet. Unlike print publications, digital publications have the ability to provide additional information to visitors about the publication and previous readers via metadata like download counts. This study investigated the effect of this metadata on the development of download inequality and unpredictability of success in an experimental academic literature marketplace. We found that presence of an accurate download count increased inequality in article downloads, meaning fewer papers accumulated a larger share of the total download count. We also found that the presence of download count increased the unpredictability of success, meaning across identical instances, different papers became the most popular. Finally, an exploratory analysis found papers were more highly rated when download counts were present. Together, these findings provide insight into how the download behaviors of previous academic readers may influence literature choice.

The first academic journals (*Journal des Sçavans* and *Philosophical Transactions* of the Royal Society of London) were both published in 1665. For over 300 years, physical journals were the most efficient way for scientists to distribute their research findings to other scientists, journalists, and lay readers (Larivière et al., 2015). The late 1990s saw the beginning of the digital era of academic publishing, with more publishers producing digital versions of their printed journals, and some new journals being exclusively digital. We are currently in the unique in-between era where scientists utilize both print and digital journals to access and disseminate research findings. Importantly, many digital outlets for scientific articles can also present additional data to the would-be reader. This could include supplemental material from the authors, such as more detailed descriptions of methodological or analytical techniques used in the paper. Digital journal articles can also include “metadata”, or data about the paper itself, which could include how many times it has been viewed or downloaded, the number of citations the paper has, or how many people are talking about it in blogs or on social media platforms (*Value All Research Products*, 2013).

In the process of doing science, it is necessary for researchers to stay current on the work being done by their peers. This is to help understand new findings in their own research, as well as to discover which scientific questions remain unanswered (Subramanyam, 2013). While the rate of scientific papers published has been increasing, the average number of articles read by scientists per month on average (22 papers) has remained unchanged since 2005 (Van Noorden, 2014), leading to an increased burden to find and read the most appropriate scientific papers for the amount of time available.

There have been a number of approaches adopted by scientists aimed at finding the most important papers to read for a given line of research. Before the digital era, when academic

literature was only in print, researchers utilized trained science librarians to help guide their reading (Chen, 1974). Now, some use digital services that rely on keywords or trends to alert the reader to new published research in an area of interest (i.e. “e-mail me whenever a paper is published using the term ‘implicit bias’” or “e-mail me papers others are talking about on Twitter”), or rely on behavioral strategies to search for the most relevant papers (Pain, 2016). All of these strategies may use article metadata as an input in the decision-making process to read a given research article.

A particularly important type of metadata are those that indicate the behavior of other scientists. The number of views, the number of downloads, and the number of citations an article has all provide a potential reader information on how other scientists have approached the article in the past. Previous research on the relationship between the download counts of biomedical papers and their subsequent citation count five years later found a moderately strong positive correlation (Pearson’s $r = 0.5$) (Perneger, 2004), suggesting higher rates of downloading (and presumably reading) a paper leads to higher use in future scientific research. Perneger (2004) ends his article by stating: “Online readers judge the scientific value of an article from the title and the abstract, and if this assessment is favorable, they access the full paper”. However, it is possible that these metadata constitute a form of social influence that acts on the potential reader, biasing their decision to engage with an article away from just the presented scientific value.

Solomon Asch’s line-length experiments clearly demonstrate the majority effect on individuals and the willingness to submit to existing social pressure (Asch, 1951). Here, individuals had the tendency to yield to the judgement of the group, even when that judgement was objectively false. In a debrief interview, one of Asch’s participants stated, “If I’d been the

first, I probably would have responded differently”, identifying the effect of having to judge the length of a line after seeing the judgement decisions made by a group of peers. Download counts on articles may generate the same effect – a signal of judgement on an academic paper by a group of peers who participated previously.

This idea has been tested more recently in both real and artificial cultural markets. In studying how cultural items such as songs, books, or movies become popular, sociologists have focused on markers that confer social preference. Muchnik et al. (2013) experimentally manipulated up-voted and down-voted comments on the popular social news aggregation web site Reddit and tracked social herding over a 5-month period. Comments that were up-voted once by the researchers generated accumulated herding that increased comments’ final ratings by 25% relative to control (unvoted by the researchers) comments, suggesting a role for social influence in the accumulation of up-votes.

Salganik, Dodds, and Watts (2006) developed a digital music jukebox that let participants listen to and download different songs. They found that, when download counts for each song were displayed to participants, download inequality went up – meaning those songs that accumulated downloads continued to accumulate downloads at the expense of other songs. This effect was enhanced when they increased the saliency of social influence by ordering the songs in descending order by download count. When the songs downloaded the most were early in the list of songs, inequality increased further.

Both the studies performed by Muchnik et al. (2013) and Salganik et al. (2006) demonstrate an effect of metadata on decision-making. This decision-making process is thought to have two-steps (Krumme et al., 2012). First, the participant needs to choose to attend to an item; to read a comment, to listen to a song, or to read an academic paper's title and abstract. This first decision is made based on appeal, position, and any other potential social preference markers, like current download count, or number of up-votes. The second decision, in this case to download a song or paper, or to vote on a comment, is based on the perceived quality of the item. Social influence, in the form of metadata, acts on the first step of this two-step process, potentially by increasing appeal, by changing accessibility via list position, or by some other means (Krumme et al., 2012). See Figure 1 for an illustration of this model.

Figure 1. Deciding to download an academic paper is a two-step process. The first decision is to choose to read an abstract, which is affected by the paper title's appeal, location/accessibility, and markers conferring social influence. The second decision is to choose to download the full text paper, which is affected by perceived quality of the abstract. Illustration and explanation adapted from Krumme et al., 2012.

Metadata on journal articles is becoming increasingly common. It is easy to find the citation count for papers in popular academic article search engines such as Google Scholar and PubMed. Download counts have been adopted by article databases such as ScienceDirect and

PsyArXiv, and altmetrics (“alternative metrics”, like shares on Twitter or mentions on Facebook) are becoming more common as scientists share their work via blogs and social media.

For this reason, it is important to determine how scientists choose which articles they will read. There is significant evidence that metadata that signals peer group judgement can result in herding effects, though this has not been demonstrated in a scientific community. It is possible that scientists, as a highly trained group of professionals, approach the scientific literature in a way that allows the highest quality research to become the most popular research. Alternatively, the behavior of some scientists may influence others to engage with the literature in a process independent of the quality of the published research.

The current work seeks to determine the effect of social influence on the decision to download scientific literature from an artificial academic market. We hypothesize that social influence will increase the inequality of downloads, meaning more downloads will accrue on a smaller number of articles. We also hypothesize that social influence will increase unpredictability of success, meaning across identical versions of the experimental academic market, different papers will become the most downloaded.

Methods

Web tool development and distribution. To collect data for this study, we developed a bespoke web tool called “AbstractFindr”, which was published online in August, 2018 at <http://www.abstractfindr.com>.

The main purpose of this tool was to present participants with the title and abstracts of academic papers. When participants clicked on a paper title, they were presented with the abstract of that paper, along with a place to rate the abstract on a one-to-five-star rating system.

Figure 2. (top) The main page of AbstractFindr. Paper titles were displayed in a single column. In the influence conditions, download counts were displayed below the title of each paper, and paper titles were listed in descending order from most to least downloaded. If two papers had the same number of downloads (including zero downloads), titles were displayed in a randomized order. When a paper title was clicked on, a pop-up with the title and abstract text was presented, along with an input for rating the abstract (one-to-five-stars) *(bottom)*. Once rated, a download decision could be made, where a click on “yes” downloaded a pdf file of the full paper to the visitor’s computer. Participants had to rate the abstract and choose to download or not for the pop-up to disappear.

Once the participant rated the abstract, they were then asked if they would like to download the full article. If they clicked “yes”, a pdf file of the full paper was downloaded to their computer, and then they were returned to the list of paper titles. If they clicked “no”, they were directly returned to the list of paper titles. Participants had to both rate and make a download decision

before they could choose another paper title. Participants were blocked from rereading an abstract after they had made a download decision, either “yes” to download or “no” to not download. See Figure 2 for an example of what AbstractFindr looked like.

Forty-eight papers were used as stimuli in this study. The complete list of papers is available in [Appendix Table A1 ‘downloads’](#). This number of papers was reached based on previous social influence literature (Salganik et al., 2006), the desire to make a list of papers comparably long to what is typically encountered by readers on a single webpage (i.e. on Google Scholar or PubMed), and consideration of the time burden for participants. All papers used as stimuli in this study were preprints publicly available on PsyArXiv, a digital platform for the distribution of non-peer reviewed academic preprints in psychology ([www.http://psyarxiv.com](http://psyarxiv.com)). We considered papers for this study if they met the following criteria:

1. The paper had fewer than 150 downloads on PsyArXiv on the acquisition date (May 9th, 2018).
2. The paper made no indication that it was either published in an academic journal or was submitted for peer review at an academic journal.
3. The paper made no indication that the authors did not want further distribution of the preprint.

Papers were chosen to represent a variety of topics within the field of psychology. See [Appendix Table A1 ‘downloads’](#) for the complete details of each paper.

The population of interest for this study was all individuals who consider themselves researchers in psychology. This included graduate students, post-doctoral researchers, lab managers, professors of all tenure levels, non-tenure track lecturers, and other titles that may be

used in other countries. As our population of interest is broad and dispersed geographically, we distributed the web tool digitally.

First, the web tool was packaged inside a Qualtrics survey for digital distribution. This was to provide another point of informed consent for all participants. After providing consent, participants were asked to visit AbstractFindr with a provided link. Invitations to participate in this study were distributed over Twitter, through the “PsychMAP” Facebook group, to members and member laboratories of the Psychological Science Accelerator via their network-wide email list, and via direct email to all American tenured or tenure-track faculty associated with a PhD-granting psychology department. Participants were also encouraged to share the participation link to their peers, students, and friends within the population of interest.

Procedure. Participants were first asked to accept the End User Agreement, which provided details about ownership of the website, ownership of the academic papers hosted on the website, how collected data will be used, and who has access to collected data. This agreement served as a second level of informed consent.

After accepting the End User Agreement, participants were then asked to provide some information about themselves: how they heard about AbstractFindr, the university or institution they are associated with, the subfield of psychology they most closely associate themselves with, their gender, how they typically access academic literature, and their position at their university or institution. No other identifying information was collected from participants. Participants were then asked to click a radio button to affirm that they knew they could only participate once. With the radio button clicked, they could proceed to the main page that contained the 48 academic papers.

At this point, participants were randomized into either the control condition (independent of social influence) or the social influence condition. See Figure 3 for the complete randomization path, which is further detailed in the following section.

Figure 3. Complete randomization path potential for each participant. Value above each path represents the probability of being randomized into the subsequent category.

Social influence. Participants were presented with academic paper titles in a 48x1 column in the center of their web browser (See Figure 2, top panel). In the experimental condition, the download count for each paper was displayed beneath each title. This download count accurately displayed how many times that paper was downloaded by previous participants. Additionally, in the experimental condition, paper titles were ordered in descending order by download count. In other words, the paper that had the most downloads was presented first in the list, followed by the paper with the second most downloads, and so on. If multiple papers had the same number of downloads (including zero downloads), those titles were presented in a random order per participant within that segment of the list. In the control condition, the

download count was not displayed to participants, and all paper titles were ordered randomly in the list per participant.

Participants were randomized into either the control condition (40%) or the experimental condition (60%). Those randomized into the experimental condition were then further randomized into either World 1, World 2, or World 3 – three parallel, independent versions of the experimental condition. Participants were free to navigate the tool as they liked, and could click, rate, and download as many titles as they'd like.

Early participants generated the social influence cues used by later participants and thus did not experience strong social cues themselves. In other words, the download decisions made by the first participants of each world created the experimental conditions experienced by later participants. The download decisions made within a world (i.e., World 1) were exclusive to that world, and did not influence participants in other worlds.

Design statement. This study is conceptually a 1x2 design (with 2 levels of social influence: no influence and influence), but was conducted as a 1x4 design, with 2 levels of social influence across 4 worlds (World 0, World 1, World 2, and World 3). Worlds 1, 2, and 3 are identical social influence replicates of each other. World 0 was the control world, independent of social influence.

This study had two outcomes of interest: download inequality, which measured how equal or unequal downloads were distributed across all possible papers, and unpredictability of success, which measured how stable a paper's popularity is across different identical worlds. If social influence affects how researchers decide to download academic papers, both download inequality and the unpredictability of success will be higher in social influence worlds (Worlds 1,

2 and 3) compared to the control world (World 0). We did not hypothesize a difference between experimental worlds, as they were exact replicates of each other. Descriptions of the analyses for download inequality and unpredictability of success follow below.

Bootstrapped Confidence Intervals. All intervals reported were estimated using the percentile bootstrap, which is a non-parametric test. In bootstrapping, the original sample is treated as a miniature representation of the larger population sampled. Observations in the sample are resampled with replacement, until a new sample of the original size is obtained. From this new sample, a new estimate is calculated. These steps are repeated 10,000 times to obtain a distribution for each estimate. Values at the 2.5th percentile and 97.5th percentile represent the lower and upper bound of the 95% CI for each estimate. All CIs were calculated using the ‘boot’ package in the statistical program R (Ripley, 2021).

Inequality of downloads. Download inequality was measured by calculating the Gini coefficient per world for the population of 48 papers. The Gini coefficient represents the expected difference in market share between two randomly chosen papers, scaled so that a coefficient of 0 represents complete equality (every paper has the same number of downloads) and a coefficient of 1 represents complete inequality (only one paper has every download) (Atkinson, 1970; Salganik et al., 2006). Gini coefficients were calculated in the statistical program R using the “ineq” package (Zeileis & Kleiber, 2015).

To determine if social influence affected how papers were downloaded, bootstrapped 95% confidence intervals were generated for each world. Gini coefficients for each world were bootstrapped. Differences between worlds were considered significantly different if the bootstrapped confidence interval for one world did not contain the estimate for another world.

Unpredictability of success. To determine if social influence changes the predictability that a particular paper becomes popular, or highly downloaded, this study collected data in three different “worlds”. The 60% of web tool visitors randomized into the experimental condition were further randomized into either World 1, World 2, or World 3 (0.33 probability per world, so that 20% of total visitors were randomized into each experimental world). Visitors only saw download counts that corresponded to the downloads of earlier visitors in their assigned world. Similarly, the list of papers was only sorted by the downloads of earlier visitors in that same world. For example, the 100th participant assigned to World 1 only saw the download counts corresponding to the previous 99 participants (and their download behavior) in World 1. They did not see the downloads of the other 100 participants that were randomized into World 2 or World 3. In this way, this study allowed the direct observation of the evolution of paper popularity in three independent, but identical, environments.

Unpredictability of success is defined as the sum total difference in download market share for an individual paper across a pair of experimental worlds (Salganik et al., 2006b), for all possible world pairs. If a paper receives the same market share in World 1 and World 2 and World 3, its unpredictability is zero. Unpredictability of each paper (u_i) was calculated across World 1, World 2, and World 3 as,

$$u_i = \frac{|m_{i,1} - m_{i,2}| + |m_{i,1} - m_{i,3}| + |m_{i,2} - m_{i,3}|}{3}$$

where $m_{i,1}$ is the market share for paper i in World 1, $m_{i,2}$ is the market share for paper i in World 2, and $m_{i,3}$ is the market share for paper i in World 3. For the control condition (World 0), data was randomly split into two subgroups, and unpredictability was calculated across these two subgroups.

The unpredictability of a condition (U) (either the “no social influence” condition or the “social influence” condition) was calculated as the average unpredictability for that condition,

$$U = \frac{\sum u_i}{48}$$

To determine if social influence affected unpredictability of success, bootstrapped 95% confidence intervals were generated for experimental and control unpredictability (U). The difference between experimental and control influence conditions was considered statistically significant if the bootstrapped confidence interval for one condition did not contain the estimate for the other condition.

Results

Participants and web tool usage. The AbstractFindr webtool was visited by 215 unique users between August 17th, 2018 and October 18th, 2018 (a total of 63 days). Participant descriptive data can be seen in Tables 1 and 2. Following the randomization plan in Figure 3, 71 participants were randomized into World 0, 56 into World 1, 43 into World 2, and 45 into World 3. A total of 906 abstracts were rated and 270 full text papers were downloaded across all worlds. A vast majority of participants typically access scientific literature digitally, and most participants learned about the study by direct email invitation.

Table 1. Reported gender of participants.

Gender	Participants	Proportion
Male	116	54.0%
Female	89	41.4%
Unanswered	10	4.7%
Total	215	100%

Table 2. *Reported gender of participants by social influence condition.*

Condition	Male	Female	Unanswered	Total
Independent	35	30	6	71
Influence	81	59	4	144

Table 3. *Psychological disciplines reported by participants.*

Area	Participants	Proportion
social	50	23.3%
cognitive	33	15.3%
clinical	28	13.0%
developmental	17	7.9%
neuroscience	14	6.5%
I/O	12	5.6%
evolutionary	8	3.7%
metascience	8	3.7%
other	8	3.7%
personality	6	2.8%
undefined	5	2.3%
educational	4	1.9%
health	4	1.9%
psycholinguistics	4	1.9%
psychophysics	3	1.4%
applied behavior	2	0.9%
decision making	2	0.9%
forensic	2	0.9%
media	2	0.9%
religious	2	0.9%
history	1	0.5%
Total	215	100%

Participants varied in psychological discipline. The greatest proportion of participating psychologists identified as social psychologists. Cognitive and clinical psychologists were also well represented. The full list of psychological disciplines of the participants can be seen in Table 3.

Download inequality. There was higher download inequality in all three influence conditions worlds compared to the independent world (see Figure 4). World 0, the control

condition which had no indicators of social influence, had a Gini Coefficient of 0.38, and a 95% bootstrapped confidence interval of [0.30, 0.45]. World 1 had a Gini Coefficient of 0.66 [0.55, 0.75]. World 2 had a Gini Coefficient of 0.66 [0.56, 0.76]. World 3 had a Gini Coefficient of 0.60 [0.48, 0.70].

Figure 4. Download inequality, measured by Gini Coefficients, per world. World 0 (*pink*) had no indicators of social influence displayed to participants. Worlds 1, 2, and 3 (*blue*) each displayed accurate download counts beneath each paper title, and paper titles were sorted in descending order by download count. Estimated Coefficient with 95% bootstrapped CIs.

The confidence interval for World 0 does not contain any of the Gini Coefficients for the social influence worlds (1, 2, or 3), meaning the Gini Coefficient for World 0 was significantly smaller than the Gini Coefficients for World 1, World 2, and World 3. The confidence interval for World 1 contains the Gini Coefficient estimates for World 2 and 3, the confidence interval for World 2 contains the Gini Coefficient estimates for World 1 and 3, and the confidence interval for World 3 contains the Gini Coefficient estimate for World 1 and 2. We concluded there was

no significant difference in the Gini Coefficient estimates between these three social influence worlds.

Abstract rating. Participants were given the ability to rate each abstract they chose to read, on a one-to-five-star scoring system (see Figure 2, *bottom*). Participants needed to rate the abstract before they could choose to either download or not download the full text of the paper they had selected.

Figure 5. Participant ratings of abstracts by influence condition. Participants in the social influence worlds rated most abstracts higher compared to those in the control condition. Points above the line ($y = x$) represent abstracts rated higher in the influence condition. From 906 total ratings on 48 papers.

The average rating for an abstract in the control condition (World 0) was 3.25 stars. The average rating for an abstract in the influence condition (Worlds 1, 2, and 3) was 3.50 stars (see Appendix Table A2'ratings'). In an exploratory 1x2 ANOVA, the average rating of abstracts in the influence condition was significantly larger than the average rating of abstracts in the control condition, $F(1, 904) = 9.83$, $p = 0.002$. Participants in the social influence condition, on average, rated the quality of abstracts as higher than those in the control condition. This analysis was considered exploratory since abstract ratings were nested within participant, and the method of data collection did not allow us to identify which ratings came from which participants, denying us the ability to perform a multi-level model.

At a per-abstract level, 34 out of 48 papers were rated higher in the influence condition compared to the control condition. Only 13 abstracts were rated lower when influence was present, and one abstract received the same rating in both influence and control conditions. Figure 5 depicts the rating for each abstract by influence condition. Points above the line represent abstracts rated higher in the influence condition.

Unpredictability of success. There was higher unpredictability of success in the social influence condition compared to the control condition (see Figure 6). The unpredictability of success for the control condition (World 0) was 0.008 [0.006, 0.011]. The unpredictability of success for the influence condition (Worlds 1, 2, and 3) was 0.05 [0.037, 0.064].

Figure 6. Unpredictability of success by social influence condition. The independent condition (World 0, *pink*) had no indicators of social influence displayed to participants. The influence condition (Worlds 1, 2, and 3, *blue*) displayed accurate download counts beneath each paper title, and paper titles were sorted in descending order by download count. Estimated unpredictability of success with 95% bootstrapped CIs.

The confidence interval for the independent condition does not contain the calculated unpredictability of success for the social influence condition, meaning the unpredictability of success for the independent condition is significantly smaller than that of the influence condition.

Discussion

In this study, there were two outcome measures that spoke to the behavior of scientists when they transact with the published literature, and one outcome measure that examined the effect of that behavior.

The first measure on scientist behavior was the Gini Coefficient, which is a measure of variability among values of a frequency distribution, first described by Corrado Gini in 1912 (Sen, 1973). In this study, a Gini Coefficient of 1 represents maximum inequality, where one academic paper receives all the downloads from participants. Conversely, a Gini Coefficient of 0 represents an equal number of downloads across all papers.

As seen in Figure 4, World 0, the control condition where no social influence was present, had a significantly lower Gini Coefficient compared to the three worlds where social influence was present. There was insufficient evidence to detect a difference in Gini Coefficient between the three influence condition worlds, suggesting similar levels of inequality across all three replicate worlds.

These data support the stated hypothesis. In the presence of social influence markers (an accurate download count per paper and a popularity-sorted list of papers), a smaller pool of papers accrued a larger share of downloads compared to the control condition, reflected in a larger Gini Coefficient. This was true for all three replicate worlds. This result indicates that researchers take a narrower approach to the published literature in the presence of social influence indicators, like download count.

The second outcome measure on scientist behavior was the subjective rating score. After selecting a paper title, participants were presented with the abstract to that paper, and then asked to rate the abstract on a one-to-five-star scale (see Figure 2, *bottom*). This rating was left

intentionally without specification, as different readers may value different qualities in an abstract (such as writing clarity, novelty of methods or findings, general interest, or specific interest to one's own research). This type of rating scale was chosen as it is a relatively common way to assess general quality and is used extensively in many domains (i.e., rating service establishments on Yelp!, or hotels on TripAdvisor).

As seen in Figure 5 and Table A2, participants on average rated abstracts significantly higher in the social influence condition compared to the control condition where no social influence markers were present. All 48 abstracts were rated in both conditions, and of them, 34 abstracts were more highly rated when social influence was present. The average rating increase for a paper from the independent to the influence condition was 5.16%.

It is important to note that participants were not presented with any indication of how previous participants rated abstracts at any point in their use of the webtool. The only social influence presented was at the paper title level, where paper titles were ordered based on how many times the full text of that paper was downloaded, and the current count of full text downloads.

It is possible that participants used download count as a proxy measure of quality. In this scenario, one would predict an abstract's rating would be correlated with the number of downloads it has received (a paper of higher quality would receive more downloads). We did not observe a statistically significant effect of influence condition to conclude that this is the case. The correlation between an abstract's rating in the influence condition and the download count in the influence condition was $r = 0.11$, $p = 0.44$.

The results from these two outcomes suggest that social influence effects both the first and second steps in the two-step decision-making process put forth by Krumme et al. (2012). First, social influence markers are presented at the presentation of paper titles. This may influence which titles are clicked on for subsequent reading of the abstract. Although no social influence markers were present at the abstract level, readers in the influence condition rated abstracts higher in quality, potentially shifting the decision to download the full text of the paper.

The third outcome measure of this study was unpredictability of success, which measures an outcome of scientists' behaviors when transacting with the literature. This unpredictability was observed by utilizing the multi-world design of this study (Salganik et al., 2006b). Worlds 1, 2, and 3 were designed to be identical in all ways. The only difference was how downloads were accumulated in each world based on the download behavior of the participants randomized into them. The number of downloads per paper were set to zero at the onset of the study in every world. As participants were randomized into a world, their download behavior was recorded to that world only. For example, if Participant A downloaded Paper 15 in World 1, Paper 15 would show 1 download in World 1, but zero downloads in World 2 and World 3. In this way, these three separate worlds could generate three unique distributions of downloads across the corpus of 48 papers.

As seen in Figure 6, unpredictability of success was significantly higher in the social influence condition compared to the control condition. What this means is that there was observed instability in the rank position of a paper across worlds. For example, the paper "Doing Good vs. Doing Bad in Prosocial Choice: A Refined Test and Extension of the Morality Preference Hypothesis" tied with another paper in receiving the most downloads in World 1 (6 downloads), received 2 downloads in World 2, and 0 downloads in World 3 (see [Appendix Table](#)

A1'downloads'). Although the structure of the worlds was the same, the interaction between the participants and the literature developed an environment where this paper could be very popular in one world, and not at all popular in another. This instability is captured by the measure of unpredictability of success.

Strengths, limitations, and future directions. This study is unlike a real academic cultural market in several ways. First, much of the information typically associated with an academic article was stripped away so that participants were only using a very specific set of information in this study. The two biggest pieces of information not available to participants were the paper's authors and the name of a publishing journal.

All articles used in this study were pre-prints available on PsyArXiv.org, and none of the chosen articles indicated that they had been published in an academic journal. Scientists have long used journals as indicators of topic interest and quality (Chen, 1974), and this influence may be even more important now given the limitations of researchers to adequately sample the growing corpus of published literature (Van Noorden, 2014). To limit the influence of journal name on the behavior of participants, only pre-prints were used as full-text articles. An important next step will be to investigate how journal prestige potentially changes the perceived quality and importance of published academic literature.

Authorship was removed from the papers used in this study as an attempt to prevent participants from engaging their scientific social network while participating in this study. Co-authorship rates have risen in the past 20 years (Henriksen, 2016), and the average American psychologist's academic social network may exceed 180 people (Fox, Honeycutt, and Jussim 2021). While it is possible that participants could possibly identify the originating lab based on the content of the title or abstract, by removing the authors' names, direct recognition of

individuals was eliminated. Future studies focusing on the relationship between authorship and reader's choice of academic literature will help to elucidate how social circles within the research community shape which findings become widely read.

The real academic cultural market also includes other potential markers of social influence, such as mentions on Twitter or Facebook, emails from influential journals (see Figure 7), and word-of-mouth from academic peers. The environment in which researchers judge and read academic papers is rich with social cues. Although this study reports effects of the presence of social cues as operationalized by download count, this may not have been the most salient cue from the actual cultural market. If researchers rely more heavily on recommendations from academic peers, or from Twitter, they may be more relevant social cues to participants. In the future, studies on the effect of social influence on researchers should investigate how different social cues are used.

Read the most popular articles of 2018 from the journals of the Association for Psychological Science

Today at 11:32 AM

[View as webpage](#) | [Email alerts](#)

Share:

Explore these top-read articles from the journals of
the Association for Psychological Science

SAGE is proud to publish the journals of the [Association for Psychological Science \(APS\)](#). Read
on for leading research, free for a limited time!*

Figure 7. An email from SAGE publishing (which publishes the journals of the Association for Psychological Science, or APS) promoting the most-read academic papers of 2018.

There is evidence that the presence of social influence can modify the way participants subjectively rate the value of stimuli via modulated engagement of the nucleus accumbens and orbitofrontal cortex. Zaki et al. (2011) found this effect 30 minutes after presenting their social influence condition, suggesting the neuronal modulation could last at least this long. Considering the possible effect of social influence on the subjective quality rating of an abstract (devoid social influence) reported in this study, it is not unreasonable to suspect neuronal modulation as a possible mechanism of action. Continued work on the neuroscience of social influence may shed more light on this potential pathway.

Conclusion. This study is the first to report on the relationship between digital academic literature and reading behaviors among academics in an experimental cultural market setting. It found that psychologists use social influence cues to guide their reading of the academic literature and that the presence of social influence cues makes it more difficult to predict which academic papers will become the most read or most popular. It also found the presence of social influence cues increased the subjective quality rating of abstracts by psychologists. Together, this study sheds light on the processes that govern how psychologists interface with the published literature, an important step in generating new scientific knowledge.

- Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. *Groups, Leadership and Men*, 119–124. <https://doi.org/10.1017/CBO9781107415324.004>
- Atkinson, A. B. (1970). On the Measurement of Inequality. *Journal of Economic Theory*, 2, 244–263.
- Chen, C. (1974). How Do Scientists Meet Their Information Needs? *Special Libraries*, 65(7), 272–280.
- Henriksen, D. (2016). The rise in co-authorship in the social sciences (1980–2013). *Scientometrics*, 107(2), 455–476. <https://doi.org/10.1007/s11192-016-1849-x>
- Krumme, C., Cebrian, M., Pickard, G., & Pentland, S. (2012). *Quantifying Social Influence in an Online Cultural Market*. 7(5), 0–5. <https://doi.org/10.1371/journal.pone.0033785>
- Larivière, V., Haustein, S., & Mongeon, P. (2015). The oligopoly of academic publishers in the digital era. *PLoS ONE*, 10(6), 1–15. <https://doi.org/10.1371/journal.pone.0127502>
- Muchnik, L., Aral, S., & Taylor, S. J. (2013). Social influence bias: A randomized experiment. *Science*, 341(6146), 647–651. <https://doi.org/10.1126/science.1240466>
- Pain, E. (2016). *How to Keep Up with the Scientific Literature*. Science Magazine.
<http://www.sciencemag.org/careers/2016/11/how-keep-scientific-literature>
- Perneger, T. V. (2004). Relation between online “hit counts” and subsequent citations: Prospective study of research papers in the BMJ. *Bmj*, 329(7465), 546–547.
<https://doi.org/10.1109/MMIT.2008.206>
- Ripley, B. (2021). *Package “boot.”* 117.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006a). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854–856.
<https://doi.org/10.1126/science.1121066>
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006b). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311(5762), 854–856.
<https://doi.org/10.1126/science.1121066>

Sen, A. (1973). *On Economic Inequality*. Clarendon Press.

Subramanyam, R. (2013). Art of reading a journal article: Methodically and effectively. *J Oral Maxillofac Pathol*, 17(1), 65–70. <https://doi.org/10.4103/0973-029X.110733>

Value all research products. (2013). 2013.

Van Noorden, R. (2014). *Scientists may be reaching peak in reading habits*. Nature.

<https://www.nature.com/news/scientists-may-be-reaching-a-peak-in-reading-habits-1.14658>

Zaki, J., Schirmer, J., & Mitchell, J. P. (2011). Social influence modulates the neural computation of value.

Psychological Science, 22(7), 894–900. <https://doi.org/10.1177/0956797611411057>

Zeileis, A., & Kleiber, C. (2015). R Package “ineq.” *CRAN Package Library*, 15.