A memory-theoretic account of citation counts.

Michael R. Dougherty

David A. Illingworth

Rosalind Nguyen

Department of Psychology, University of Maryland

Abstract

Despite the common assumption that citations are indicative of an article's scientific merit, increasing evidence indicates that citation counts are largely driven by variables unrelated to quality. In this paper, we treat people's decisions of what to cite as an instance of memory retrieval and show that observed citation patterns are well accounted for by a model of memory. The proposed exposure model anticipates that small alterations in factors that affect people's ability to retrieve to-be-cited articles from memory early in their life cycle are magnified overtime and can lead to the emergence of highly cited papers. This effect occurs even when there is no variation in the starting point exposure probabilities (i.e., when assuming a level playing field where all articles are treated equally and of equal 'quality'), and is exacerbated by natural variation in retrievability of articles due to encoding. We discuss the implications of the model within the context of research evaluation and hiring, promotion, and tenure decisions.

A memory-theoretic account of citation counts.

Introduction

As a scientist, it's nice to know (or at least believe) that one's work is being read and cited by others who find it useful. What makes an article highly cited, however, is of considerable debate. On the one hand, citation counts are commonly treated as proxies for quality. The assumption being that if a paper is of high quality, it is more likely to be cited. As intuitive as it may seem however, the inverse inference does not follow. We cannot infer quality from citation counts. Though it seems reasonable to assert that quality might indeed drive some citations, the bibliometric literature indicates that there are many other factors unrelated to quality that also relate to citation counts. For instance, posting an article to a preprint server (Gargouri et al., 2010), having a co-author whose name is well recognized (Petersen et al., 2014), being part of a large academic social network (Mählck & Persson, 2000), presenting at conferences (de Leon & McQuillin, 2018), and the use of social media can all positively boost citation counts (Eysenbach, 2011; Gentil-Beccot, Mele, & Brooks, 2009; Luarn & Chiu, 2016; Peoples, Midway, Sackett, Lynch, & Cooney, 2016), while the inclusion of mathematical equations decreases citation rates (Fawcett & Higginson, 2012). There is even evidence that citations covary with length of titles and the use of colons and hyphens in titles (Paiva, Lima, & Paiva, 2012; Zhou, Tse, & Witheridge, 2019). Lurking underneath these factors is another curious finding: A relatively large proportion of citations are perfunctory — made for reasons that are not clearly necessary for the authors scientific arguments — while others reflect 'negative' endorsements (Krampen, Becker, Wahner, & Montada, 2007; Stremersch, Camacho, Vanneste, & Verniers, 2015). Even more problematic is the possibility that authors cite articles without having read them or for

the purpose of appeasing reviewers or editors: Artino Jr., Driessen, and Maggio (2019), for example, found that approximately 49% of authors reported both of these behaviors.

The idea that citation counts are influenced by factors that are seemingly unrelated to quality is problematic – not for reasons related to the ethicality of scientific behavior, but for reasons related to how citation counts are used. It is no secret that counting the number of citations researchers garner figures prominently into people’s perceptions of scientific impact. These perceptions, founded on the belief that citations reflect quality or ‘impact’, influence important decisions. The most obvious decisions are those related to hiring, tenure, and promotion: University promotion committees from the department on up to the president’s office often want to see citation-based metrics. But it is also likely that citation counts have other knock-on effects, such as influencing speaker invitations, scientific awards, grant awards, and even affecting university rankings. Such knock-on effects can in principle further impact citation behaviors by amplifying the exposure of those who receive invitations and awards and increasing the opportunity for their research to be viewed, while restricting opportunities for less well-cited authors. As we discuss later, these effects can have real world practical implications that can adversely impact the career trajectories of some groups more than others.

Although it is tempting to maintain the belief that citation counts reflect some dimension of quality, the scientific literature examining this issue is equivocal. For instance, Dougherty and Horne (2019) examined the relationship between citation counts and objective measures of article quality such as the presence of inconsistencies in the reporting of statistics, the degree of evidential support provided by the data, and replicability of scientific findings. The revealed relationships were weak at best and

definitely not meaningful, and in some cases inversely related to quality. In general, article-level citation counts were either unrelated to objective measures of research quality, or inversely related (e.g., articles containing at least one statistical decision error – in which the p-value indicated significance, but the test statistic did not, or vice versa – were cited more often than articles with no errors). Nieminen, Carpenter, Rucker, and Schumacher (2006) examined the relation between methodological soundness and citation counts and found no meaningful relationship. Similarly, analyses examining experts' subjective ratings of quality often show weak or non-existent relationships between rated quality and future citation counts (Patterson & Harris, 2009; Ruano et al., 2018; West & McIlwaine, 2002). Wang, Veugelers, and Stephan (2017) derived an objective measure of novelty, an often cited dimension of 'quality', and showed that for truly novel research (that which makes unique connections across different literature's) citations accrue more slowly early in their life cycle than less novel work; although, citations to highly novel work is more likely to be in the top 10% 15 years post publication, it is also more likely to be in the bottom 10%. Even one of the most thorough and extensive exercises in evaluating research quality, the Research Excellence Framework (REF) (HEFCE, 2015)) failed to show consistent correlations between peer review and bibliometric indicators.¹

On the flip side, analyses of duplicate publication of identical papers (i.e., papers published twice in different journals) reveals that citation counts for articles that are objectively equivalent in quality (because they are identical) do vary, primarily as a

¹ Although the correlation between citation counts and REF quality scores was reasonably high in some disciplines, including psychology (e.g., Spearman's $\rho=0.41$), all reviewers had access to journal names and some review panels, including psychology, even had access to citation counts. It is reasonable to assume that reviewers are knowledgeable enough to know the impact factors of many of the journals. Thus, the quality ratings cannot be considered as independent of either the journal impact factors or citation counts (Teplitskiy, Duede, Menietti, & Lakhani, 2020). Indeed, the average correlation between citation counts and REF quality scores was substantially higher for panels who had access to citation data ($\rho=.50$) compared with panels who did not ($\rho=.17$, Bayes factor=85313).

function of the journal impact factor (Larivière & Gingras, 2010). Moreover, evidence suggests that people's perceptions of quality is systematically influenced by citation counts: People show a bias to overestimate research quality when told that a paper is associated with a high citation count, but underestimate quality when the paper is associated with a low citation count (Teplitskiy et al., 2020). Altogether, citation counts do not meaningfully vary as a function of objective indices of quality or subjective judgments of quality, *do* vary systematically within articles that are identical in quality, and appear to bias people's perceptions of quality. Coupled with the factors reviewed above that predict citation counts, these findings suggest that the link between quality and citation counts is tenuous at best (see Aksnes, Langfeldt, & Wouters, 2019).

Setting aside the factors reviewed above that affect citation counts, might there be another compelling yet parsimonious explanation for why some articles are cited a lot and others are not? In this paper we address this question by treating citation distributions as a memory-theoretic phenomena that manifests as a function of authors' citing choices. We demonstrate that citation patterns can be largely accounted for by a memory process model in which small variations in the retrievability of scientific papers early on in their life cycle is magnified overtime resulting in the emergence of highly cited articles. This effect occurs even when there is no variation in the starting point citation probabilities due to say, quality or any other factor, and can easily produce so-called 'super citations' as a result of natural variation in the memorability of an article early in its life cycle. Our analysis casts doubt on the use of citation counts as a metric of research quality or impact and suggest that the number of citations garnered by papers and people may well be more reflective of the cognitive processes of the people doing the citing than of the author whose work is cited.

The Exposure Model of Citation Counts

Fundamentally, the exposure model treats citation behavior as a memory-theoretic phenomena. We assume that authors choose who to cite by retrieving candidate 'references' from memory and selecting those that are most relevant for the citing context. The citing context is assumed to act as a memory retrieval prompt for the author, bringing to mind the set of potential articles to reference in the text. In this sense, our assumptions about how and why particular papers are chosen by authors to cite is no different than any other act of memory retrieval, including, for example the choice amongst words when producing written documents. Thus, we make two core assumptions: (a) people choose what to cite based on what comes to mind in a particular context and (b) what comes to mind in a particular context is largely a function of factors that influence retrieval (e.g., encoding, frequency of exposure, recency of exposure, etc).² These two assumptions are motivated by a wealth of prior research in cognitive science, including work on the availability heuristic (Tversky & Kahneman, 1973), the recognition heuristic (Goldstein & Gigerenzer, 2002), and the validity and mere exposure effects (Hasher, Goldstein, & Toppino, 1977). The availability heuristic implies that the choices one makes (e.g., what to cite) in a particular context is largely driven by what is immediately retrievable from memory (Tversky & Kahneman, 1973). The mere exposure and validity effects reflect the fact that frequently or recently experienced items (i.e., citations) are preferred over, and

² Although the assumption that citation behavior is a memory-theoretic phenomena forms the basis of the model, it is not necessary to assume that all referencing decisions are made solely on the basis of retrieval of the reference per se, as we assume that the memory processes detailed herein broadly apply to how people generate search terms (from memory) when searching electronic databases, and how they decide who to cite when evaluating the output of a search engine (i.e., by relying on familiarity or recognition of an authors name). Nevertheless, we do assume that the majority of articles one decides to cite when writing an academic article are ones the author generates from memory or are chosen based on a recognition-based heuristic (Goldstein & Gigerenzer, 2002)

viewed as more valid than less-frequently or less-recently experienced items.

Availability, the mere-exposure effect, and the validity effect all derive from a common underlying memory mechanism accommodated by models of memory (Dougherty, Gettys, & Ogden, 1999). The common theme running through these phenomena is that retrievability functions as a cue for inferring other quantities (Goldstein & Gigerenzer, 2002). For instance, prior exposure to causes of death in newsprint media or through one's social circle predicts the perceived frequency of those causes (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1978; Pachur, Hertwig, & Steinmann, 2012); perceived truthfulness of assertions increases as a function of repetition or recency (Hasher et al., 1977); and the perceived quality of products increases with brand familiarity (Griffith & Gray, 2002). Even the perceived quality of academic journals has been shown to correlate with self-rated familiarity of journals (Serenko & Bontis, 2011). As we show below, the two assumptions stated above are sufficient to account for the Matthew Effect in citation behavior, in which articles that receive more citations are more likely to be cited again in the future (Merton, 1968), the finding that articles cited earlier in their life cycle are more likely to be cited more often, and the common belief that more highly cited articles are more important or of higher quality.

Before continuing, it is useful to point out that our proposed models are broadly consistent with a class of steady state skew probability distribution functions outlined by Newman (2009), Price (1976), and Simon (1955). These authors provide elegant mathematical frameworks for understanding the emergent properties of distributions that well describe citation frequencies, but these frameworks do not provide a cognitive-theoretical account linking the distributions to plausible psychological mechanisms. Our approach is to model citation counts as reflecting fundamental

properties of memory. Our conceptualization and theoretical framework derive from well known memory mechanisms, reflecting the intuition that what authors decide to cite in many (if not most) contexts is, fundamentally, an act of memory retrieval. In this respect, our approach provides novel insights into the plausible behavioral and psychological processes that guide citation behavior. Broadly, we assume that the probability that an article is cited is proportional to its exposure to the scientific community, where exposure is modeled as a function of how well an article has been encoded in memory, and a count of the number of times in the past that an article has appeared in print, is cited in other articles, or is promoted in other ways (e.g., social media, blogs, presented at a conference, etc). The model assumes that each article begins its publication life cycle with some units (u) of exposure due to its publication. The probability that an article is selected for citation is simply the article's accrued exposure units (u) divided by the total number of exposure units across the set of hypothetical journal articles. Here we implement the exposure model within the context of a simple model of memory. We illustrate that the natural stochastic properties of our model, coupled with well-justified assumptions of memory processes, capture many of the properties observed in citation distributions, including the emergence of highly cited articles in the literature, without assuming that authors make citing decisions based on research quality or scientific merit.

We develop two versions of the exposure model: A stochastic-only version that merely assumes that the probability of citing an article is proportional to the number prior exposures an article has received and a memory-theoretic version that represents cited articles as vectors of features that are encoded in memory where forgetting can be modeled via a single parameter (Hintzman, 1988). The stochastic model allows us to

illustrate the properties of the model without imposing memory-theoretic assumptions.

It makes minimal assumptions and has no free parameters. We assume that all articles in reference class R start on a level playing field with $u=1$ and with an a priori equal probability of being cited of $p(cited) = u / \sum u$. R defines the set of articles to be modeled. Articles in R will accrue some number of citations, c , within some predefined time frame, or epochs. Each epoch corresponds to a single opportunity for one article to be cited. Thus, the number of epochs is set to the total number of citations in R .

Article citations are represented as frequency counts; thus, the stochastic model is essentially a frequency counter that tabulates how often each paper is cited over time. Articles selected for citation accrue additional exposure units randomly across epochs according to the iteratively updated distribution of citation probabilities. The value of u is incremented by 1 across epochs when an article is selected to be cited or receives other exposure. Articles that receive more citations therefore have a higher probability of being cited again, which in turn reduces the probability that previously uncited articles will receive future citations. Note that the model assumes that citations are a finite resource.³ Thus, one consequence of the model is that the presence of highly cited articles within a particular distribution necessarily results in stagnation of citations for other articles. To provide a concrete example of the model, imagine a set of 100 articles collectively cited 1000 times. Therefore, $R = 100$ and the number of epochs equals $c = 1000$. For each epoch of the simulation, one article from the field of 100 is selected at random to be cited. The probability of each article being cited in the first epoch is 0.01 (1/100). For the first epoch, the value of u for the cited article is increased

³ This assumption follows from both practical constraints, i.e., there is no need to cite all possible articles to substantiate a claim, and from actual constraints many journals place on the number of allowable references in a paper.

by 1. The exposure probability of that article is then increased to 0.0198 ($2/101$), while the exposure probability for the non-cited articles all decrease to 0.0099 ($1/101$).

During the 2nd epoch of the life cycle, an article is selected randomly according to the updated probability distribution and the process is repeated. The entire process is repeated for c epochs to simulate the evolution of citation patterns across the life cycle of all 100 articles. Individuals interested in exploring the predictions of the model can access R code to run the model at <https://osf.io/hmbup/>.

Note that the stochastic model has no free parameters and simulates the expected distribution of citations without the need for optimizing parameters to maximize fit. The assumption that all articles start with an equal value of u is obviously an oversimplification, as it is easy to imagine how various factors could lead some articles or authors to naturally have greater exposure, or receive greater attention than others. Nevertheless, we start with the assumption of the level playing field for purposes of demonstrating the stochastic properties model.

It is also possible to implement the general processes detailed above within a more complex theory of memory. For this, we adopt an instance-based model of memory that represents experienced events as vectors of features (Hintzman, 1988). The model starts with the assumption that each article in reference class R is represented by a single trace in memory. On epoch 1, a probe vector representing the citing context is created and matched against all traces in memory. This match yields an activation value for each trace, which is simply the cube of the similarity between the probe vector and trace. The normalized activation's of all traces exceeding zero is then computed and a trace is sampled and retrieved probabilistically using Luce's choice rule. Because traces of references that have a greater similarity to the citing context are more highly active,

they are more likely to be sampled and retrieved from memory and cited by an author. We assume that these new citations represent new exposure opportunities. Thus, for each newly cited article, a vector representing the newly cited article is added to memory and the process is repeated. Thus, episodic memory grows by one trace for each epoch of the model, with there being $C+R$ total traces in memory at the completion of the simulation. The memory-theoretic version of the model is implemented in Python. Supplemental material providing a more detailed description of the model, along with Python code for running the model, is provided on <https://osf.io/hmbup/>. Computer code used to distill the model output and produce the plots included in this paper was written in R and is also provided at <https://osf.io/hmbup/>.

The memory-theoretic model has one parameter, L , reflecting the degree to which the traces in memory retain details of the original event. The L parameter can be thought of as capturing information loss that occurs either through lack of attention at encoding or via forgetting. Although encoding and forgetting are conceptually distinct, the model does not differentiate between these two types of information loss. Nevertheless, the model does allow us to explore the effect of information loss on citation patterns. The memory representations used in the exposure model consist of vectors of features, which can be thought of as representing conceptual and contextual attributes of scientific articles. In principle, the similarity structure across populations of scientific articles can be modeled by representing articles that share conceptual and contextual attributes with vectors that share a high degree of feature overlap.

Demonstration Simulations

To illustrate the properties of the memory-theoretic model, we conducted a series of simulations varying the L parameter across four levels and similarity across three

levels. The manipulation of L allows us to examine the impact of information loss on citation patterns. The manipulation of similarity allows us to examine whether citation patterns might be expected to differ between domains of research in which there is a high degree of conceptual overlap (e.g., domains in which the research community is narrowly focused) versus domains in which there is a lower degree of conceptual overlap (e.g., domains characterized by broad focus). It is well known that citation distributions show strong positive skew, with very few highly cited articles and relatively more poorly cited articles. The question addressed by these simulations is whether the degree of skew, and hence the emergence of highly cited articles, might be affected by natural variability in information loss and domain similarity.

We chose a research ecology consisting of $R=739$ articles that were collectively cited 57310 times, which corresponds to the number of articles (R) and citations (C) to papers published in 1989 and included the in statcheck database (Hartgerink, 2016; Nuijten, Hartgerink, van Assen, Epskamp, & Wicherts, 2016). The encoding parameter was varied across four levels: $L=.35$, $.55$, $.75$, and 0.95 , while sim was held constant at 0.1 . For reference, an encoding parameter of 0.35 corresponds to a case where 35% of the features corresponding to the original event are retained in the memory trace, whereas the 65% are lost. A value of $L=0.95$, corresponds to the case where 95% of an events original features are preserved in memory and only 5% are lost. Note that the operation of L is stochastic, such that the actual degree of encoding for any individual trace will vary around its expected value. To model similarity, we constructed memory traces that share Sim percent of features, where Sim varied across $\text{Sim}=0.1$, 0.25 , 0.50 , 0.75 . For these simulations, L was held constant at 0.50 . Only a single simulation run was used for each value of L and Sim to more clearly illustrate the variation in

predicted citation counts (averaging across runs can mask this variability).

Figure 1. Density plots of the memory-theoretic model's predicted distribution of citation counts as a function of the encoding parameter. The tick marks at the bottom correspond to individual articles.

Figures 1 and 2 provide the results of these demonstration simulations, along with the observed distribution of citation counts for the year 1989. A few observations are worth mentioning. First, decreased encoding increases the skew of the distribution. The model predicts that highly cited articles are more likely to emerge from an environment

Figure 2. Density plots of the memory-theoretic model's predicted distribution of citation counts as a function of the similarity of articles. The tick marks at the bottom correspond to individual articles.

in which the expected information loss is greater. The effect manifests in the model based on natural variation in how well articles are encoded in memory. When most articles are poorly encoded (low L), few that receive higher attention by random chance (i.e., higher encoding) will have a greater retrieval advantage compared to cases in which most articles are well represented in memory. For instance, when the expected

value of encoding is 0.95, there is little room for articles to be encoded at a higher rate (only 5%). In contrast, when encoding is 0.35, there is a much greater opportunity for an article to be encoded at a much higher rate simply by chance. The disparity between the activation's of the best and least-well encoded vectors should be greater when L is low.

The model also predicts increasing positive skew as a function of similarity: Highly cited articles are more likely to emerge when the reference class consists of a set of highly similar articles. But note also that high-similarity also increases the number of *uncited* articles (see Figure 2). These simulations suggest that research domains in which a lot of researchers are addressing the same (or highly similar) phenomena will produce highly skewed citation distributions with a small number of super citations or so-called 'citation classics'. Given the inherent challenges with operationalizing similarity, we have not examined these predictions in detail.⁴ In the simulations below, we hold similarity constant at a fairly low level ($\text{sim}=0.1$) and only manipulate L when fitting the memory-theoretic model.

Application of Models to Data

Data description

The models were fit to 46,690 articles (cited 2,161,838 times) from 32 consecutive years (1985-2016) of articles included in the statcheck database (Hartgerink, 2016; Nuijten et al., 2016). The citation data for each of these articles was obtained using the `scholar` (Keirstead, 2016) package in R (R Core Team, 2019) and is available at

⁴ Although similarity amongst articles is a difficult construct to define, it may be possible to capture some dimensions of similarity with high-dimensional models. Another way to define similarity is to examine the degree to which articles cite common articles (Wang et al., 2017).

<https://osf.io/hmbup/>, along with analyses for all 32 years of data. Citation data were gathered in January 2019 using the scholar API tool; citation counts for 1000 articles were checked manually by the first author to verify the accuracy of the automated tool. The goal of the modeling exercise was to evaluate the degree to which the model predicts the distribution of actual citations. The model does not enable us to predict the citation counts for individual articles, nor should it given that the underlying thesis is that citations derive from a stochastic memory-theoretic process. However, it does allow us to examine the degree to which observed distributions of citations follow the proposed exposure process. To this end, we are primarily concerned with the degree to which the characteristics of citation distributions are consistent with the exposure model.

For the stochastic version of the model, we ran 100 simulations to simulate the distribution of citations for each year. For the memory-theoretic version of the model, we varied the model's encoding parameter, L , and conducted 10 repetitions at each value of L for each each year that was simulated. L was manipulated across a minimum of 4 levels in steps of 0.1 for each year to find an approximate best fit. A total of 144 simulations were run across the 32 years of data, with 10 repetitions of each simulation. Note that the model produces highly consistent predictions across simulation runs, with pairwise correlations between runs ranging from 0.97 to 0.99 for a given value of L . This is largely due to the fact that the citation probabilities converge on a steady state. Thus, there is little to gain by running additional simulations.⁵ The value of Sim was held constant for all simulations at 0.1. We did not attempt to fit distributions by

⁵ The memory-theoretic version of the model is highly computationally intensive, as its memory computations are repeated for each epoch, with one new trace added to memory for each epoch. For our application of the model, the median number of epochs for a single simulation run was **66,769**. Thus a single simulation run entails a median of 66,769 iterations of the model, each of which was repeated 10 times for each value of L .

varying Sim.

To assess model fit, we use two indices to characterize the overlap between the predicted and actual citation distributions. One index is the A statistic, which is a non-parametric probabilistic measure of effect size. The value of A reflects the probability that a randomly drawn value from the observed distribution is greater than a randomly drawn value from the predicted distribution. A is a non-parametric estimator of the Common Language effect size, and has an expected value of 0.50 under the null hypothesis that two distributions are identical. The A statistic is closely related to both the Mann-Whitney U and Kendall's tau, with the U statistic providing a valid test of the null hypothesis. The second index is a measure of the degree to which the probability density functions (PDF's) between two distributions overlap (Pastore & Calcagni, 2019). We use the `gemm()` package (Chrabaszcz, Tidwell, & Dougherty, 2018) to compute A and the `overlapping` package (Pastore, 2018) in R to compute the similarity between two PDF's. Note that the expected overlap between any two PDF's generated from the same population is sample size dependent⁶. Thus, we provide both the raw and the adjusted value of overlap corrected for chance.⁷

Simulation results

Panels A, B, and C of Figure 3 plots the actual and predicted distributions of citations using the stochastic model. The three years that are plotted, 1989, 2003, and

⁶ The expected theoretical upper bound for overlap between two empirical probability density functions with the same population parameters is sample size dependent. This upper bound can be estimated by simulating two distributions with sample size N drawn from the same population and computing the overlap between the two densities. For the stochastic model, we estimate the upper bound by simulating 1000 pairs of log-normal distributions with $N = R$ for each year. The average overlap of these simulated distributions ranged from 88.6% to 94.5%. For the memory theoretic version of the model, the upper bound was estimated from all pairwise runs of the simulation, of which there are $10*(10-1)/2=45$ for each simulated year. The average upper bound from these simulations **ranged from 89.6% to 94.3%.**

⁷ Skewed distributions are oftentimes plotted on a log-log scale and fit with a linear model. These analyses are provided in online supplemental material.

2010, represent the years for which the stochastic model had the *poorest* fit (1989), *median* fit (2003), and the second *best* fit (2010) based on the overlapping index.⁸ Table 1 provides the fit statistics for these three years as well as the average fit across all 32 years. The average adjusted overlap for the stochastic model was 85.8%. The mean value of A was 0.548. Both of these metrics indicate that a large proportion of the variability in true-citation counts is consistent with a simple stochastic process, without assuming any specific psychological mechanisms, yet the model's fit is far from perfect.⁹

Panels D, E, and F of Figure 3 plots the actual and predicted citation distributions based on our memory-theoretic account for the three selected citation distributions used to demonstrate the stochastic model. The bottom half of Table 1 provides the summary fit statistics for these distributions along with the mean, minimum, and maximum values across all 32 years. Overall, the memory-theoretic account provides a better fit to the data than the stochastic model, as illustrated by the average values for A and the overlapping indices. Recall that the A statistic is an index of the degree to which two distributions differ, analogous to Cohen's D , but expressed as a probability. Under the point-null hypothesis of no difference, $A = .50$. By comparison, Cohen's $d = .10$ is roughly equivalent to $A = 0.528$ (Dunlap, 1999). The mean value of A across the 32 years was 0.508, indicating that the simulated and actual citation distributions are nearly identical. Indeed, a comparison of observed distribution with the memory-theoretic predicted distribution showed strong support the hypothesis of no

⁸ The year 2016 had a slightly better fit than 2010. We decided to plot data for 2010 instead of 2016 because 2016 included a much smaller number of articles and a much smaller number of citations than the several years prior. Data for 2016 was included in the calculation of the average across all 32 years and the average of 2016 fits are available on OSF for those who wish to examine that and all other years.

⁹ Simon (1955) used a version of this model to describe word frequency distributions in written text. In as much as word usage involves retrieval, it is plausible to assume that the data generating mechanism that produces word-frequency distributions is similar to the data generation underlying citation distributions.

difference for most years (median $BF_{01}=10.7$ in favor of the null).¹⁰ In only four of the 32 cases was evidence for the null inconclusive ($BF_{01}<3$), and in only one case (2010, $BF_{01}=0.65$) was the BF_{01} less than 1.0. Although the model provides an excellent fit to the data, two aspects of the distributions are not well accounted for: The model under-predicts a small number of papers with extreme citation counts and over predicts the number of papers that will be uncited. Although it's possible for the memory-theoretic model to better account for highly cited articles in the extreme tail of the distribution by decreasing L or varying the similarity of the vectors, doing so further increases the predicted proportion of uncited papers.

As shown in Panel A of Figure 4, the value of L for the best fit models generally increased across years, with the higher values of L needed for modeling citations to articles published more recently, and lower values of L for articles published in the more distant past. This pattern is consistent with the memory-theoretic interpretation of citation behavior, as we assume papers published further in the past are less well represented in memory due to forgetting.

An interesting question concerns why some articles accrue more citations than others. Adams (2005) showed that articles cited *earlier* in their life cycle tend to be cited more often *across* their life cycle compared to articles that receive their first citation later. This general trend is naturally captured by the exposure model, as shown in Figure 5, which shows the evolution of citation probabilities for eleven simulated articles drawn from the simulation of citation distribution for the year 1989. The numbers to the right correspond to quantiles of the final predicted citation count. For instance, the article with the highest final citation count is given by the upper most line

¹⁰ For this analysis we used a non-parametric Bayes factor based on the Mann-Whitney U statistic van Doorn, Ly, Marsman, and Wagenmakers (2020).

Figure 3. Density plots of actual citations (red) and model predictions (blue) for three years of publications. Model predictions are based on the median value of each quantile across 100 simulations for each year.

(1.0), and the article corresponding to the 99th percentile in terms of final citation counts is given by next highest line. The inset of this plots shows the evolution of citation probabilities predicted by the memory-theoretic model across the entire lifespan; the main figure shows the citation probabilities for the first 5000 Epochs of the model. A few patterns can be discerned from these plots. Consistent with Adams (2005), the most highly cited articles are those that are cited early in their life cycles. This pattern is important for at least three reasons: First, it indicates that highly cited articles are guaranteed to emerge naturally from a stochastic system, even without assuming that highly cited articles are any more meritorious than less-widely cited article. Second, the simulation suggests that interpretations of early citation counts as an indicator of long-term citation 'impact' is a self-fulfilling prophecy. As the model demonstrates, papers that are cited early will naturally accrue more citations simply due to the stochastic nature of model and the increased exposure that they enjoy. Third, to a large degree, what determines whether an article becomes highly cited can be plausibly accounted for by random variability. A second pattern that can be discerned is that although early citations do not guarantee that an article will become highly cited, the model anticipates that the lack of early citations can make it very difficult for meritorious ideas to emerge later. Put another way, with each passing epoch in which an article goes uncited, the odds that it will ever be cited decrease. Although the model anticipates only a small percentage of papers will go uncited, it clearly shows that these papers become increasingly less likely to be cited with time. The implication of this effect is that some articles, and by extension the scientific knowledge contained within those articles, are choked out by a small number of highly cited articles.

Panels B and C in Figure 4 provides a more detailed examination of the model's

predictions in regard to uncited articles. Panel A plots the best fitting value of the encoding parameter for each publication year. The positive slope of the best fitting line can be interpreted as reflecting greater forgetting over time. For instance, in order the best model citation patterns for articles published in the 1980's to the present compared to articles published in the 2000's, we must assume a greater degree of forgetting.

Figures B and C plot the actual and predicted proportion of uncited articles for both the stochastic model (panel B) and the memory-theoretical model (panel C). Overall, both models appear to over-predict the proportion of uncited articles. Although the stochastic model better accounts for the proportion of uncited articles in the 80's and 90's, it fails to accurately capture the non-monotonic function observed in the data.

Super citations

The simulations presented in the left column of Figure 3 and the top half of 1 assume that all articles start with an equal probability of being cited. As discussed above, however, there is plenty of evidence to suggest that the playing field is uneven and that some articles (or researchers) enjoy an exposure advantage that heightens the amount of attention their work receives. This advantage may be due to the topic (hot topics are cited more frequently because there are more people working on the problem), the size of ones academic or social network, the use of social media to promote ones work (pre-print servers, tweeting, etc), or merely the readership of the journal. Of course, it's also entirely possible that buried amongst the factors affecting exposure are attributes related to scientific merit — though it is important to note that (a) scientific merit is neither necessary nor sufficient to ensure exposure, and (b) there is little evidence to support the assertion that citations are related to merit or quality (Aksnes et al., 2019; Dougherty & Horne, 2019). This general pattern in which the

Figure 4. Plots of best fitting encoding parameter as function of publication year (A) and predicted number of papers with zero citations for both the stochastic model (B) and the memory-theoretic model (C).

occurrence of highly cited articles simultaneously increases the number of poorly cited articles is exactly what Michalska-Smith and Allesina (2017) observed in their analysis of citation patterns for National Academies members.

The above results indicate that small changes in the starting point exposure probabilities are sufficient to create highly cited articles. At the same time, the

Figure 5. Plot of the predicted evolution of citation probabilities over time for the memory-theoretic model for year 1989. Main panel shows first 5000 epochs. Inset shows evolution over all epochs. Each line represents a different quantile (1.0,.99,.95,.90,.80,.70,.60,.50,.40,.30,.20,.10,.05,.01,.00) based on the final probabilities. Thick black dashed line corresponds to the 50th percentile of the distribution. The red dashed line corresponds to the starting point probabilities.

simulations illustrate that articles that benefit from increased attention early on detrimentally impacts citations to other papers. Thus, not only do the rich get richer,

Table 1. *Average fit statistics for three representative years of article citations and the mean of all 32 years of citations. The years were selected to represent a year with poor fit (1989) a year with median fit (2003) and a year with one of the top fits (2010), based on the stochastic model and using the A statistic as the index of fit. R=number of articles; C=number of citations; A = non-parametric measure of effect size; OV = uncorrected overlap of real and simulated citation distribution; OVadj = overlap of real and simulated citations corrected for the chance.*

Stochastic model					
Year	R	C	A	OV	OV _{adj}
1989	739	57310	.575	71.12	79.65
2003	1131	92948	.550	77.41	85.21
2010	2128	67212	.514	85.25	92.08
Mean 32 years	1459	67557	.548	78.18	85.8
Memory-theoretic model					
1989 (L=0.55)	739	57310	.504	79.89	89.0
2003 (L=0.75)	1131	92948	.504	85.2	93.3
2010 (L=1.0)	2128	67212	.523	84.83	90.0
Mean 32 years	1459	67557	.508	83.46	90.54
Minimum 32 years	694	39182	.502	78.04	84.18
Maximum 32 years	3874	107733	.523	88.10	96.33

but the poor get poorer. This is a hallmark of the Matthew Effect. Although it is possible that at least some of these super citations do indeed reflect transformative or paradigm shifting contributions, it's worth questioning whether articles become

transformative because they have been widely disseminated, or if they become widely disseminated because they're transformative (for an evidence that citations counts change perceptions of quality Teplitskiy et al., 2020, see). It's also worth asking how many papers that aren't cited contain transformative ideas. In as much as scientists hold the belief that highly cited articles are more important or transformative than poorly cited articles, the mere fact that a large proportion of articles are either uncited or poorly cited raises the real possibility that scientists have fallen prey to the pseudodiagnosticity effect (Doherty, Mynatt, Tweney, & Schiavo, 1979) — inferring that highly cited articles are of higher quality without assessing the quality of poorly cited articles.

Summary

The primary finding derived from our simulations is that highly cited articles can arise from noisy memory retrieval. Indeed, the more noise there is in the memory system, the more likely it is that highly cited articles will emerge. The second finding is that relatively small amount of variation in exposure of an article early on in its life-cycle can create relatively large variation in future citation counts. Our findings demonstrate that random variation in citation patterns early on has an out-sized effect on the future citation counts by altering the exposure probabilities, a finding that is consistent with past research on citation patterns (Adams, 2005), and easily explained as a function of memory theory. Moreover, the simulations offer an alternative account for citation patterns that challenge traditional orthodoxy that citations reflect the underlying scientific merit of the research. Specifically, our simulations illustrate that the pattern of citation distributions can be largely accounted for by recasting authors' decisions about what to cite as an emergent property of memory retrieval. This account

of citation behavior suggests that the critical variables underlying whether a paper is cited frequently or less frequently is largely determined by variables that affect its memorability. We suggest that many of the variables previously identified as related to citation patterns (e.g., length of title, familiarity of the authors, 'surprisingness' of findings, use of social media to promote ones work) are in fact variables that affect retrieval. Not only does this explanation make intuitive sense, it is also consistent with a variety of findings within psychological science, such as the mere exposure and validity effects, as well recent work showing that citation counts are sensitive to primacy effects (Stevens & Duque, 2019).

The idea that citation patterns follow a stochastic process is consistent with prior work by Michalska-Smith and Allesina (2017). Michalska-Smith and Allesina (2017) argued that citation patterns, and the apparent positive relationship between number of publications and the occurrence of highly cited articles is consistent with a stochastic process. This finding is also accounted for by the exposure model. Mechanistically, our model assumes that this stochastic process is rooted in memory processes, such that the more available a reference is in memory, the more likely it is to be retrieved and cited when cued by the citing context.

The exposure model clearly challenges the common belief that citation counts reflect the underlying scientific merit of the cited work. Unfortunately, the emerging literature examining the relationship between scientific quality and citation counts has largely failed to find consistent relationships between measures of research quality and citation counts (Dougherty & Horne, 2019; Nieminen et al., 2006; Paiva et al., 2012; Ruano et al., 2018; West & McIlwaine, 2002). Indeed, in their review of the literature on this issue, Aksnes et al. (2019) concluded that "citation indicators seem of little help

in the evaluation of the solidity/plausibility, originality, and societal value of research (pg. 12).” At the same time, there is evidence that citation counts are related to a variety of non-quality related factors, including some that directly reflect the exposure of an article to the academic community.

Note that the choice of what is retrieved is based on the similarity between the citing context and the traces in memory. Thus, the model will be biased (in a good way) to select references to cite that have a strong match (high similarity) to the citing context. One implication of this feature is that articles written on so-called ‘hot topics’ will necessarily have a heightened probability of being cited relative to articles on less-popular topics. This is a natural by-product of using the citing context as the probe: If a particular context (e.g., on a popular topic) occurs more frequently, traces of articles that match that context are more likely to be retrieved from memory.

Implications

An important question raised by our model concerns the factors that might affect the exposure probabilities of one’s work. Although there are a few seemingly obvious candidate hypotheses, we suspect that what drives exposure probabilities is complex and multiply determined. Obvious sources of exposure include the size of a journal’s readership, social media, press releases, news media reports, open-access, and use of pre-print servers. However, other factors likely enhance the visibility and exposure of a paper, including name recognition (Petersen et al., 2014).

While the results provide a recipe book for cooking up ones citation counts, they should also stimulate concerns about factors that might limit the exposure of one’s research to the scientific community. One such concern is parenting. Women are far more likely to be the primary caregiver for young children and as a result forego

conference attendance during critical stages in their career (Mason, n.d.). Because conferences provide an important opportunity for researchers to expose their latest research findings to the research community, missing these opportunities can negatively impact the visibility of these scientists (Calisi, 2018). The same could be said for scientists who forego conference attendance in order to reduce their carbon footprint. The exposure model anticipates that these early missed opportunities can be magnified over the course of one's career and affect citation rates. It's also useful to point out that the model suggests that whether or not a particular paper is cited is not just a function of its own exposure, but also a function of the exposure of other papers. That is, the probability that a paper is cited will decrease overtime as a function of the increased exposure of other papers. In other words, in the marketplace of scientific ideas, what may matter most is how much effort is put into promoting ones work, not the merit of the idea being promoted.

Our simulations assume that articles and researchers all start on a level playing field. But this is not the case in real life. It's easy to imagine variables that advantage some and disadvantage others: Do researchers from more prestigious institutions receive greater exposure of their work, or does it receive more attention, simply due to their affiliations? Does co-authoring with an already well-known researcher increase the attention the article receives? Might subtle racial, gender or other biases alter who researchers decide to cite (yes, see Greenwald & Schuh, 1994), and therefore affect exposure probabilities? Might the cultural climate of a scientific discipline in terms of gender, racial, or even political inclusivity affect researchers' willingness to attend conferences to present their work? Not only is it easy to imagine cases of each of these scenarios, evidence suggests that such effects are plausible. Decades or work on the

mere exposure effect shows that people have a preference for familiarity and even believe more strongly in claims deemed more familiar (Hasher et al., 1977). Work on racial prejudice has documented evidence of racial preferences in citation patterns. The immediate effects of any one of them may not be particularly large, but it doesn't have to be in order to have a major impact across the life cycle. Small changes in starting values at the outset of one's career or in the initial phases post-publication are easily magnified over time. There is clearly variation to be explained in citation counts, and aspects of the quality of the research may well be one small source of variation. However, what is clear from the simulations is that one need not assume that citation counts are based on quality to observe large disparities in citation counts.

Our simulations have a clear connection to prior work on the effect of perverse incentives on research quality. Smaldino and McElreath (2016) illustrated that incentives that focus on maximizing publication quantity can result in the selection of poor scientific methods. This selection process can happen in two ways. First, increasing the number of research products necessarily comes at a cost (Edwards & Roy, 2017). Smaldino and McElreath (2016) showed that those costs can manifest in poor research methods and ultimately lead to the selection (hiring, tenure, and promotion) for researchers who use substandard research practices. Second, as the number of publications one produces increases, the chances that one of those articles becomes highly cited also increases — merely as a function of chance. This later prediction is a natural consequence of the exposure process described above and is consistent with the finding observed by Michalska-Smith and Allesina (2017). Thus, not only will researchers who publish more look more productive (in terms of quantity), but they will also benefit from greater citation counts and a higher chance of at least one

paper becoming a super citation. By this account, one's citation "impact" has more to do with chance, than it does with the merits of one's work.

One potential criticism of our work is that perhaps the articles that are cited early on in the life cycle are chosen for their overall quality. If this were indeed the case then one would expect that an article's future citation counts would correlate with objective measures of quality. Unfortunately, there is little to no convincing evidence that this is the case: The evidence linking citations to objective measures of quality, and even subjective measures of quality, is weak at best

Our model offers a plausible competing account of citation patterns rooted in memory theory. Of course, even within this mechanistic account, one can still postulate that how well an article is represented in memory is due to some aspect of quality. Although we cannot fully reject this possibility, it is clear to us that research quality is neither a necessary nor sufficient factor in determining whether an author will retrieve a particular citation in a given context. One interpretation of our model's predicted citation distributions is that they represent the expected distribution of citations under the null hypothesis, where the null here is defined as a probabilistic data generating mechanism. Regardless of whether one maintains the belief that citations reflect quality or impact, the fact that the true distribution of citations is indistinguishable from a probabilistic process should give us pause and lead us to question the discriminate validity of citations counts. The standard in classical statistics is to infer the presence of an effect by comparing the observed data to a properly formulated null hypothesis. That same standard should apply to the use of citation counts to measure scientific merit, impact, or research quality.

A natural response to our conclusions could be that rather than measuring quality,

perhaps citations reflect influence or 'impact'. The problem with this this response is that it creates a potential tautology in which the very thing we're trying to measure is the definition of the measured variable. Even if one dismisses the tautology, redefining citations as a measure of impact doesn't actually solve the central problem. Whether one believes citations reflect quality, 'impact', or something else, its use as a selection variable is not justified if it is indistinguishable from a random process (as our model results imply) or largely determined by a host of non-merit based variables (as implied by the extant literature). Put more bluntly, we suspect that most people would object to the use of any variable to select who to hire or promote that is either random or is driven by extraneous factors.

A path forward

If citation counts and impact factors are to be abandoned, how might one evaluate the quality and impact of scholarly work? A number of proposals have been put forth that we believe provide a viable set of responsible indicators for assessing scientists (RAIS's, Moher et al. (2018)). Frank (2019) summarized a proposal that involves serious evaluation of a small set of a researcher's best research products (N-best); Moher et al. (2018) and others suggest the inclusion of criteria that capture ones use of research practices that promote reproducibility (making research tools, analysis code, and data publicly available); Fraley and Vazire (2014) proposed the N-pact factor — a metric for assessing journal quality based on statistical power or sample size. In addition to these, we suggest conducting reproducibility audits of scientists' research as a way of verifying and scoring ones use of open science practices, while asking for authors to more fully document on their CV's their practices for ensuring reproducible science (Dougherty, Slevc, & Grand, 2019). In contrast to research metrics that emphasize quantity

(quantity of publications, quantity of citations, high impact factors), which can promote the use of poor scientific methods (Edwards & Roy, 2017; Smaldino & McElreath, 2016), the alternative metrics proposed above are designed to incentivize what we as a scientific community value: quality science based on sound methodology whose conclusions can be trusted. More than that, because science is a public good often funded by taxpayer dollar, we should strive to maximize the benefit of our science to the public. Doing so requires sharing our research in the broadest way possible to both facilitate future science and make the research products available to the general public.

In sum, it is time to stop pretending that citation counts are a reasonable measure of the quality of a scientific paper or researcher. Although one cannot reject quality as one of the possible variables that affect the exposure probabilities of a paper, there is sufficient evidence that citation counts are multiply determined by a host of factors unrelated to scientific quality. These factors are likely to drown out any signal due to quality. Indeed, the fact that highly cited articles can emerge from processes early on in the life cycle of an article should give us pause when considering these numbers for tenure and promotion decisions. Maybe there is value in the use of citation counts for some other reason, such as an index of the degree to which one has promoted their work to the broader scientific community. Yet, even so, the citation count is probably not the best indicator of promotion for the same reasons as it's not a good measure of quality. What meaning we ascribe to citation counts in the context of quality or 'impact' undoubtedly rests on the validity of the assumptions we make regarding the reasons or intentions of authors to decide to cite a particular article.

References

- Adams, J. (2005, Jun 01). Early citation counts correlate with accumulated impact. *Scientometrics*, 63(3), 567–581. Retrieved from <https://doi.org/10.1007/s11192-005-0228-9> doi: 10.1007/s11192-005-0228-9
- Aksnes, D. W., Langfeldt, L., & Wouters, P. (2019). Citations, citation indicators, and research quality: An overview of basic concepts and theories. *SAGE Open*, 9(1), 2158244019829575. Retrieved from <https://doi.org/10.1177/2158244019829575>
- Artino Jr., A. R., Driessen, E. W., & Maggio, L. A. (2019). Ethical shades of gray: International frequency of scientific misconduct and questionable research practices in health professions education. *Academic Medicine*, 94(1). Retrieved from https://journals.lww.com/academicmedicine/Fulltext/2019/01000/Ethical_Shades_of_Gray__International_Frequency_of.27.aspx
- Calisi, R. M. (2018). Opinion: How to tackle the childcare–conference conundrum. *Proceedings of the National Academy of Sciences*, 115(12), 2845–2849. Retrieved from <https://www.pnas.org/content/115/12/2845> doi: 10.1073/pnas.1803153115
- Chrabaszcz, J., Tidwell, J., & Dougherty, M. (2018). gemmR: General monotone model [Computer software manual]. (R package version 1.3-4)
- de Leon, F. L. L., & McQuillin, B. (2018). The role of conferences on the pathway to academic impact: Evidence from a natural experiment. *The Journal of Human Resources*. doi: 10.3368/jhr.55.1.1116-8387R
- Doherty, M. E., Mynatt, C. R., Tweney, R. D., & Schiavo, M. D. (1979).

- Pseudodiagnosticity. *Acta Psychologica*, 43(2), 111 - 121. Retrieved from <http://www.sciencedirect.com/science/article/pii/0001691879900179>
doi: [https://doi.org/10.1016/0001-6918\(79\)90017-9](https://doi.org/10.1016/0001-6918(79)90017-9)
- Dougherty, M. R., Gettys, C. F., & Ogden, E. E. (1999). Minerva-dm: A memory processes model for judgments of likelihood. *Psychological Review*, 106(1), 180-209. doi: 10.1037/0033-295X.106.1.180
- Dougherty, M. R., & Horne, Z. (2019, Sep). *Citation counts and journal impact factors do not capture research quality in the behavioral and brain sciences*. PsyArXiv. Retrieved from psyarxiv.com/9g5wk doi: 10.31234/osf.io/9g5wk
- Dougherty, M. R., Slevc, L. R., & Grand, J. A. (2019). Making research evaluation more transparent: Aligning research philosophy, institutional values, and reporting. *Perspectives on Psychological Science*, 14(3), 361-375. Retrieved from <https://doi.org/10.1177/1745691618810693>
- Edwards, M. A., & Roy, S. (2017, Jan 01). Academic research in the 21st century: Maintaining scientific integrity in a climate of perverse incentives and hypercompetition. *Environmental engineering science*, 34(1), 51-61. Retrieved from <https://www.ncbi.nlm.nih.gov/pubmed/28115824> (28115824[pmid]) doi: 10.1089/ees.2016.0223
- Eysenbach, G. (2011, Dec 16). Can tweets predict citations? metrics of social impact based on twitter and correlation with traditional metrics of scientific impact. *J Med Internet Res*, 13(4), e123. Retrieved from <http://www.jmir.org/2011/4/e123/> doi: 10.2196/jmir.2012
- Fawcett, T. W., & Higginson, A. D. (2012). Heavy use of equations impedes communication among biologists. *Proceedings of the National Academy of*

- Sciences*, 109(29), 11735–11739. Retrieved from <https://www.pnas.org/content/109/29/11735> doi: 10.1073/pnas.1205259109
- Fraley, R. C., & Vazire, S. (2014, 10). The n-pact factor: Evaluating the quality of empirical journals with respect to sample size and statistical power. *PLOS ONE*, 9(10), 1-12. Retrieved from <https://doi.org/10.1371/journal.pone.0109019> doi: 10.1371/journal.pone.0109019
- Frank, M. C. (2019, Aug). *N-best evaluation for academic hiring and promotion*. PsyArXiv. Retrieved from psyarxiv.com/tzaqh doi: 10.31234/osf.io/tzaqh
- Gargouri, Y., Hajjem, C., Lariviere, V., Gingras, Y., Carr, L., & Brody, T. (2010). Self-selected or mandated, open access increases citation impact for higher quality research. *PloS ONE*, 5(10), e13636.
- Gentil-Beccot, A., Mele, S., & Brooks, T. C. (2009). Citing and reading behaviours in high-energy physics. *Scientometrics*, 84, 345-355.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75-90. doi: 10.1037/0033-295X.109.1.75
- Greenwald, A. G., & Schuh, E. S. (1994). An ethnic bias in scientific citations. *European Journal of Social Psychology*, 24(6), 623-639. doi: 10.1002/ejsp.2420240602
- Griffith, D. A., & Gray, C. C. (2002). The fallacy of the level playing field. *Journal of Marketing Channels*, 9(3-4), 87-102. Retrieved from https://doi.org/10.1300/J049v09n03_05 doi: 10.1300/J049v09n03_05
- Hartgerink, C. (2016). 688,112 statistical results: Content mining psychology articles for statistical test results. *Data*, 1(3), 14.
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of

- referential validity. *Journal of Verbal Learning & Verbal Behavior*, 16(1), 107-112.
doi: 10.1016/S0022-5371(77)80012-1
- HEFCE. (2015). *The metric tide: Correlation analysis of ref2014 scores and metrics (supplementary report ii to the independent review of the role of metrics in research assessment and management)*. doi: 10.13140/RG.2.1.3362.4162
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95(4), 528-551. doi: 10.1037/0033-295X.95.4.528
- Keirstead, J. (2016). scholar: analyse citation data from google scholar [Computer software manual]. Retrieved from <http://github.com/jkeirstead/scholar> (R package version 0.1.5)
- Krampen, G., Becker, R., Wahner, U., & Montada, L. (2007, May 01). On the validity of citation counting in science evaluation: Content analyses of references and citations in psychological publications. *Scientometrics*, 71(2), 191-202. Retrieved from <https://doi.org/10.1007/s11192-007-1659-2> doi: 10.1007/s11192-007-1659-2
- Larivière, V., & Gingras, Y. (2010). The impact factor's matthew effect: A natural experiment in bibliometrics. *Journal of the American Society for Information Science and Technology*, 61(2), 424-427. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.21232>
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory*, 4(6), 551-578. doi: 10.1037/0278-7393.4.6.551
- Luarn, P., & Chiu, Y.-P. (2016). Influence of network density on information diffusion

- on social network sites: The mediating effects of transmitter activity. *Information Development*, 32(3), 389-397. Retrieved from <https://doi.org/10.1177/0266666914551072>
- Mählck, P., & Persson, O. (2000). Socio-bibliometric mapping of intra-departmental networks. *Scientometrics*, 49(1), 81-91. Retrieved from <https://doi.org/10.1023/A:1005661208810>
- Mason, M. A. (n.d.). The baby penalty. the chronicle of higher education. [Computer software manual]. Retrieved from <https://www.chronicle.com/article/The-Baby-Penalty/140813>
- Merton, R. K. (1968). The matthew effect in science. *Science*, 159(3810), 56-63. Retrieved from <https://science.sciencemag.org/content/159/3810/56> doi: 10.1126/science.159.3810.56
- Michalska-Smith, M. J., & Allesina, S. (2017, 06). And, not or: Quality, quantity in scientific publishing. *PLOS ONE*, 12(6), 1-12. Retrieved from <https://doi.org/10.1371/journal.pone.0178074> doi: 10.1371/journal.pone.0178074
- Moher, D., Naudet, F., Cristea, I. A., Miedema, F., Ioannidis, J. P. A., & Goodman, S. N. (2018). Assessing scientists for hiring, promotion, and tenure. *PLOS Biology*, 16(3), 1-20. Retrieved from <https://doi.org/10.1371/journal.pbio.2004089>
- Newman, M. E. J. (2009, jun). The first-mover advantage in scientific publication. *EPL (Europhysics Letters)*, 86(6), 68001. Retrieved from <https://doi.org/10.1209/0295-5075/86/68001> doi: 10.1209/0295-5075/86/68001

- Nieminen, P., Carpenter, J., Rucker, G., & Schumacher, M. (2006). The relationship between quality of research and citation frequency. *BMC Medical Research Methodology*, 6(1), 42. Retrieved from <https://doi.org/10.1186/1471-2288-6-42>
- Nuijten, M. B., Hartgerink, C. H. J., van Assen, M. A. L. M., Epskamp, S., & Wicherts, J. M. (2016). The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior Research Methods*, 48(4), 1205–1226. Retrieved from <https://doi.org/10.3758/s13428-015-0664-2>
- Pachur, T., Hertwig, R., & Steinmann, F. (2012). How do people judge risks: Availability heuristic, affect heuristic, or both? *Journal of Experimental Psychology: Applied*, 18(3), 314–330. doi: 10.1037/a0028279
- Paiva, C. E., Lima, J. A. P. d. S. N., & Paiva, B. S. R. (2012). Articles with short titles describing the results are cited more often. *Clinics*, 67, 509 - 513. Retrieved from http://www.scielo.br/scielo.php?script=sci_arttext&pid=S1807-59322012000500017&nrm=iso
- Pastore, M. (2018). overlapping: Estimation of overlapping in empirical distributions [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=overlapping> (R package version 1.5.2)
- Pastore, M., & Calcagni, A. (2019). Measuring distribution similarities between samples: A distribution-free overlapping index. *Frontiers in Psychology*, 10, 1089. Retrieved from <https://www.frontiersin.org/article/10.3389/fpsyg.2019.01089> doi: 10.3389/fpsyg.2019.01089

- Patterson, M., & Harris, S. (2009). The relationship between reviewers' quality-scores and number of citations for papers published in the journal physics in medicine and biology from 2003–2005. *Scientometrics*, 80, 343–349. Retrieved from <https://doi.org/doi.org/10.1007/s11192-008-2064-1>
- Peoples, B. K., Midway, S. R., Sackett, D., Lynch, A., & Cooney, P. B. (2016). Twitter predicts citation rates of ecological research. *PLOS ONE*, 11(11), 1-11. Retrieved from <https://doi.org/10.1371/journal.pone.0166570>
- Petersen, A. M., Fortunato, S., Pan, R. K., Kaski, K., Penner, O., Rungi, A., ... Pammolli, F. (2014). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences*. Retrieved from <https://doi.org/10.1073/pnas.1323111111>
- Price, D. D. S. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27(5), 292-306. Retrieved from <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/asi.4630270505> doi: 10.1002/asi.4630270505
- R Core Team. (2019). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/> (Version 3.6.0)
- Ruano, J., Aguilar-Luque, M., Isla-Tejera, B., Alcalde-Mellado, P., Gay-Mimbrera, J., Hernández-Romero, J. L., ... García-Nieto, A. V. (2018). Relationships between abstract features and methodological quality explained variations of social media activity derived from systematic reviews about psoriasis interventions. *Journal of Clinical Epidemiology*, 101, 35-43. Retrieved from <https://doi.org/10.1016/j.jclinepi.2018.05.015>

Serenko, A., & Bontis, N. (2011). What's familiar is excellent: The impact of exposure effect on perceived journal quality. *Journal of Informetrics*, 5(1), 219 - 223.

Retrieved from

<http://www.sciencedirect.com/science/article/pii/S1751157710000696>

doi: <https://doi.org/10.1016/j.joi.2010.07.005>

Simon, H. A. (1955, 12). On a class of skew distribution functions. *Biometrika*, 42(3-4), 425-440. Retrieved from <https://doi.org/10.1093/biomet/42.3-4.425> doi: 10.1093/biomet/42.3-4.425

Smaldino, P. E., & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, 3(9), 160384. Retrieved from <https://royalsocietypublishing.org/doi/abs/10.1098/rsos.160384> doi: 10.1098/rsos.160384

Stevens, J. R., & Duque, J. F. (2019). Order matters: Alphabetizing in-text citations biases citation rates. *Psychonomic Bulletin & Review*, 26(3), 1020–1026. Retrieved from <https://doi.org/10.3758/s13423-018-1532-8> doi: 10.3758/s13423-018-1532-8

Stremersch, S., Camacho, N., Vanneste, S., & Verniers, I. (2015). Unraveling scientific impact: Citation types in marketing journals. *International Journal of Research in Marketing*, 32(1), 64 - 77. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0167811614000792> doi: <https://doi.org/10.1016/j.ijresmar.2014.09.004>

Teplitskiy, M., Duede, E., Menietti, M., & Lakhani, K. R. (2020). *Citations systematically misrepresent the quality and impact of research articles: Survey and experimental evidence from thousands of citers.*

- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207 - 232. Retrieved from <http://www.sciencedirect.com/science/article/pii/0010028573900339>
doi: [https://doi.org/10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9)
- van Doorn, J., Ly, A., Marsman, M., & Wagenmakers, E.-J. (2020). Bayesian rank-based hypothesis testing for the rank sum test, the signed rank test, and spearman's . *Journal of Applied Statistics*, 0(0), 1-23. Retrieved from <https://doi.org/10.1080/02664763.2019.1709053> doi: 10.1080/02664763.2019.1709053
- Wang, J., Veugelers, R., & Stephan, P. (2017). Bias against novelty in science: A cautionary tale for users of bibliometric indicators. *Research Policy*, 46(8), 1416 - 1436. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0048733317301038>
doi: <https://doi.org/10.1016/j.respol.2017.06.006>
- West, R., & McIlwaine, A. (2002). What do citation counts count for in the field of addiction? an empirical evaluation of citation counts and their link with peer ratings of quality. *Addiction*, 97(5), 501-504. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1360-0443.2002.00104.x>
- Zhou, Z. Q., Tse, T. H., & Witheridge, M. (2019). Metamorphic robustness testing: Exposing hidden defects in citation statistics and journal impact factors. *IEEE Transactions on Software Engineering*, 1-1. Retrieved from <https://doi.org/10.1109/TSE.2019.2915065>