Contents lists available at ScienceDirect

Journal of Informetrics

journal homepage: www.elsevier.com/locate/joi

Regular article

Megajournal mismanagement: Manuscript decision bias and anomalous editor activity at PLOS ONE

Alexander M. Petersen

Ernest and Julio Gallo Management Program, Management of Complex Systems Department, School of Engineering, University of California, Merced, CA 95343, United States

A R T I C L E I N F O

Article history: Received 12 July 2018 Received in revised form 24 August 2019 Accepted 29 August 2019 Available online 14 September 2019

Keywords: Science ethics Editorial board service Perverse incentives Bias Self-citation Human analytics

ABSTRACT

Since their emergence just a decade ago, nearly 2% of scientific research is now published by megajournals, representing a major industrial shift in the production of knowledge. Such high-throughput production stresses several aspects of the publication process, including the editorial oversight of peer-review. As the largest megajournal, PLOS ONE has relied on a single-tier editorial board comprised of ~7000 active academics, who thereby face conflicts of interest relating to their dual roles as both producers and gatekeepers of peerreviewed literature. While such conflicts of interest are also a factor for editorial boards of smaller journals, little is known about how the scalability of megajournals may introduce perverse incentives for editorial service. To address this issue, we analyzed the activity of PLOS ONE editors over the journal's inaugural decade (2006–2015) and find highly variable activity levels. We then leverage this variation to model how editorial bias in the manuscript decision process relates to two editor-specific factors: repeated editor-author interactions and shifts in the rates of citations directed at editors - a form of citation remuneration that is analogue to self-citation. Our results indicate significantly stronger manuscript bias among a relatively small number of extremely active editors, who also feature relatively high self-citation rates coincident in the manuscripts they handle. These anomalous activity patterns are consistent with the perverse incentives and the temptations they offer at scale, which is theoretically grounded in the "slippery-slope" evolution of apathy and misconduct in power-driven environments. By applying quantitative evaluation to the gatekeepers of scientific knowledge, we shed light on various ethics issues crucial to science policy - in particular, calling for more transparent and structured management of editor activity in megajournals that rely on active academics.

© 2019 The Author. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

The emergence and rapid growth of megajournals¹ in the last decade represents a drastic industrial paradigm shift in the production of scientific knowledge (Binfield, 2013; Bjork, 2015; Pan, Petersen, Pammolli, & Fortunato, 2018; Petersen, Pan, Pammolli, & Fortunato, 2019; Solomon, 2014; Solomon & Bjork, 2012; Wakeling et al., 2016). This transition places pressure on several fundamental aspects of the scientific endeavor. First, the personnel resources required to referee the

https://doi.org/10.1016/j.joi.2019.100974







¹ Megajournals are typically online-only e-journals, the largest of which publish upwards of 500 articles per month and serve a multi-disciplinary audience. Consequently, the production process is primed for growth, in particular international growth (Wakeling et al., 2016). The top 5 megajournals, ranked by the number of articles published in 2016 (in parenthesis) according to Scimago Journal & Country Rank, are: PLOS ONE (22,159), Scientific Reports (20,883), Royal Society of Chemistry Advances (13,025), Oncotarget (6391), and Physical Review B (5483).

^{1751-1577/© 2019} The Author. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/ 4.0/).

60,000+ megajournal articles each year is quite substantial (Binfield, 2013). Second, such publication volume also stresses the cognitive and technological capacity of individual scientists in their ability to search, retrieve, and organize the research literature. By way of example, over its first 6 years, PLOS ONE grew at an annual rate of 58%, roughly 18 times larger than the net growth rate of scientific publication over the last half-century (Petersen, Pan, et al., 2019). Consider for example 2012, in which the 23,468 articles published by PLOS ONE alone represented approximately 1 in every 1000 science publications indexed by the Web of Science (Pan et al., 2018). And third, megajournals rely on a highly scalable model for managing the scientific publication process. In particular, PLOS ONE relies on thousands of acting scientists comprising its editorial board, who simultaneously continue their role as research leaders. The dichotomy of being both a producer and gatekeeper of knowledge, which is also common to other journal editorial boards, brings forth the conditions for conflict of interest – as scientists must balance conflicting incentives arising from their duties as both authors and editors.

Oversight of the editorial board is a formidable challenge in megajournals, calling for strategic management policies to document editor activities and address unintended incentives. However, our fundamental understanding of this problem is hindered by the lack of transparency – both during the review process and also post-publication.² Ironically, while this lack of transparency – e.g. blinding of authors, reviewers, and editors – is justified as facilitating unbiased peer review, it makes external monitoring of operational misconduct difficult. This tradeoff is an important consideration for the management of the scientific practice, not least because research shows that misconduct can arise organically from the basic pursuit of internal (and external) power (Malhotra & Gino, 2011) and the innate difficulty of avoiding temptation in decision-heavy endeavors (Gino, Schweitzer, Mead, & Ariely, 2011) – conditions characteristic of science.

For these reasons, the oversight of editor activity is necessary in order to address social and cognitive biases that can enter into the manuscript decision process. For example, the multi-disciplinary journal *Proceedings of the National Academy of Sciences* (PNAS) has a two-tiered editorial board system in which National Academy of Science members serve as editors of individual submissions, and a smaller rotating Editorial Board provides an additional oversight layer for approving final decisions. Similarly, *Management Science* also employs a two-tier system comprised of a rotating Editor-in-Chief which serves above a second layer of "Department Editors" who oversee the review and manuscript decision process for individual submissions. There are plenty of additional examples of journals that also employ multiple-tiered editorial boards, e.g. comprised of principal editors, associate editors, and advisors. Compared with the volume of megajournals, the relatively small size of traditionally print journals limits the net activity of any given editor.

Contrariwise, megajournals have developed around principles of scalability. Consequently, megajournal editors also have the opportunity to scale their decision-making power beyond the levels that are available through editorial board service in smaller journals. Given the recency of this paradigm, little is known about the variation in megajournal editor activity, the upper limits of extreme activity, and the degree to which perverse self-citation incentives may explain such extreme activity. Here we address this knowledge gap by performing an in-depth analysis of the largest journal in the world, PLOS ONE, over its first 10-years of publishing (2006–2015). Not surprisingly, this journal also has the largest distributed editorial board of any megajournal. Another distinction, one that is common to the entire family of PLOS journals and commendable for its leadership in supporting open and transparent science, is the explicit reporting of the particular handling editor associated with each published article. Thus, by combining handling editor, manuscript, author and post-publication citation data for each article, we constructed a large multi-variable dataset centered around the 6934 PLOS ONE editors. The longitudinal nature of this dataset facilitates identifying the role of social factors underlying editor manuscript decisions, thereby providing insight into a domain of science that has traditionally been poorly documented, since most journals do not reveal editor-article associations. As such, we contribute to recent efforts aimed at measuring biases in the editor manuscript decision process (Bravo, Farjam, Moreno, Birukou, & Squazzoni, 2018; Card & DellaVigna, 2017; Colussi, 2018; Sarigöl, Garcia, Scholtes, & Schweitzer, 2017).

The results of our analysis indicate that a relatively small set of PLOS ONE editors are exceeding reasonable activity levels. By way of comparison, we observe 85 PLOS ONE editors (or 1% of all editors) with activity levels exceeding the most active PNAS editor (see Fig. 1). Moreover, we identify 10 extremely active editors (denoted collectively by XE) who exhibit significant differences when compared to the remaining PLOS ONE editors. To be specific, articles accepted by XE are accepted significantly faster, have higher rates of citations to the editors' research, and lower citation impact relative to the other PLOS ONE research articles.

As such, after we document the descriptive statistics on editor activity levels, we then shift the focus to the outliers in the editor activity distribution, and identify anomalous statistical differences in manuscript handling outcomes between XE and the remaining editors. Our results provide complementary lines of evidence supporting an underlying self-citation strategy. It is indeed challenging to determine with absolute certainty whether these differences point to either apathy or misconduct on the part of XE, which is beyond the scope of our analysis. Independent of the source cause of the identified

² We chose to study the megajournal PLOS ONE primarily because the full article text are readily available and in a stable format, because it has a single review and publication process, because it is not affiliated with a society (in which case we would likely be missing significant information capturing significant author-editor social relations), and because the sample size of publications and editors is large enough to be amenable to methods of statistical inference. Unlike most journals, PLOS ONE does provide the name of the editor overseeing each article, a crucial aspect which we leverage in this study. Two other multidisciplinary journals with a distributed editor management system that also provide the editorial board member name on each article are the weekly journal *Proceedings of the National Academy of Sciences of the United States of America* and the monthly journal *Management Science*.



Fig. 1. The distribution of editor activity in three multi-disciplinary journals. (A) Comparison of distributions of annual activity n_E (per editor per year) for three journals with distributed editorial boards comprised of acting academics. This complementary cumulative distribution, $CCDF(\ge n_E)$, plots the fraction of all editors that oversee n_E or more articles per year, showing that there is a small but significant subset of PLOS ONE editors (1%) exhibiting extremely high activity levels. (B) Comparison of $P(S_E)$, calculated as kernel density estimates, of normalized annual activity S_E (per editor per year); S_E measures the annual activity of a given editor relative to the median activity across all journal size. A small but significant subset of extremely active editors are distinguished by comparison of right tails of the distributions. Vertical dashed lines indicate the activity level corresponding to the top-5% of all editors for a specific journal.

editor manuscript handling bias, another issue we bring to light is the need for better editorial board oversight in order to address editorial apathy and/or misconduct, which may otherwise persist unchecked in a large single-tier editorial board system. Thus, in addition to standard editor policies (Editorial Policy Committee, 2012), our results lend support for a two-tiered editorial board system model, one that places increased emphasis on the accountability of manuscript-level editors. This is a relevant publication norm to address because it is quite possible that such anomalous editor activity is more widespread in science than currently appreciated, and would be evident for other journals if their editor data were also available.³ It is also a timely issue to address in light of the relatively rapid growth of the megajournal ecosystem (Petersen, Pan, et al., 2019).

Addressing underlying sources of the issue may be as simple as implementing editor activity quotas. Also, journals with distributed boards that do not specify the manuscript handling editor in the article metadata byline should strongly consider following the example of PLOS journals in making this editor handling information readily available. Taking this idea further, it would also benefit science to make all prior reviewer reports openly available after a certain number of years in order to support a more rigorous assessment of the peer-review system. Steps in this direction will increase the transparency of the publication process, facilitate trust in the peer review process, and foster the development of data sources, methods and protocols for rigorously assessing the integrity of the scientific peer review system (Editorial Policy Committee, 2012).

In what follows, we provide background and motivation for our study in Section 2, and specify our data sources and statistical measures in Section 3. We present descriptive results in Sections 4.1, 4.2, where we characterize the extremely skewed activity levels and manuscript decision times at PLOS ONE. By contrasting editor activity levels at PLOS ONE with *PNAS* and *Management Science*, two additional journals that also employ distributed editorial boards comprised of academic editors, we argue that the activity levels among the most prolific PLOS ONE editors is anomalous – thereby meriting further in-depth analysis.

As such, we continue in Sections 4.3, 4.4 to explore how and why an editor might take on extreme activity levels for strategic gain. In Section 4.3 we investigate the *how* by using longitudinal panel regression to show that when editors handle articles by the same set of returning *repeat authors*, they accept these particular manuscripts significantly faster than if the authors are first-timers; we then show that this pattern is even stronger among the 10 extremely active editors. Similarly, we also show that manuscript handling times are faster if the manuscript authors include references that cite the editor's research, an effect that is also larger in magnitude among XE.

In Section 4.4 we address the *why* by measuring the degree to which preferential treatment of repeat authors is reciprocated through citations directed back at the editor's research publications in the form of citation renumeration, providing substantial evidence for editor-author "backscratching". According to our analysis, we estimate a lower limit of 100s of citations that XE could reap by aggressively scaling up manuscript handling activity.

In Section 4.5 we pursue these lines of evidence further by focusing on three anomalous editors who are identified as simultaneous outliers in various categories, including their net activity and the frequency of citations to their work in

³ Our analysis does not specifically investigate or make any conclusions based on the fact that PLOS ONE is an open-access journal. Nevertheless, it is important to note that many journals entering the megajournal market use an APC ("pay-to-publish"), open-access, continuous publication model, e.g. Heliyon (Elsevier), Springer Plus (Springer), Scientific Reports (Nature), Royal Society Open Sciences, IEEE Access, Peerl, SAGE Open (Binfield, 2013; Bjork, 2015; Solomon, 2014; Solomon & Bjork, 2012). For example, only four of the top 10 journals by 2016 publication volume do not follow the APC open-access model: Physical Review B, ACS Applied Materials & Interfaces, Proceedings of the National Academy of Sciences of the United States of America (PNAS), and Physical Review D. Instead, here we address the need for transparent and centralized oversight of the editor team megajournals; The two-tiered editor Board and Manuscript Editor system used by PNAS is a good example, thereby providing a more robust three-tiered manuscript review process.

manuscripts they accept. Using the complete career publication records for these active researchers, we show that anomalous citation rates occur not just among the PLOS ONE articles they accept, but that self-citation rates by the editors within their own published research also exceeds their baseline citation rate in the literature.

We summarize our results in Section 5, which concludes with straightforward and feasible policy recommendations pertaining to megajournal management and editorial board oversight.

2. Background and motivation

This work contributes to multidisciplinary research streams converging around the modeling of the scientific endeavor (Fortunato et al., 2018; Scharnhorst, Börner, & Besselaar, 2012) by way of data-driven computational social science methods (Lazer et al., 2009). Together these efforts aim to improve our understanding of incentives and inequality in science (Stephan, 2012), the growth of science and its implications for knowledge production and research evaluation (Pan et al., 2018; Petersen, Pan, et al., 2019), ethics issues related to the growth and densification of academic social networks (Petersen, Pavlidis, & Semendeferi, 2014), and the knowledge about knowledge itself (Evans & Foster, 2011).

Introspective research on the scientific endeavor can provide much-needed guidance on challenging science policy issues (Börner, Edmonds, Milojević, & Scharnhorst, 2016; Fealing, 2011; Stephan, 2012) – e.g. how to increase the efficiency of scientific discovery (Rzhetsky, Foster, Foster, & Evans, 2015) and funding allocation (Bollen, Crandall, Junk, Ding, & Börner, 2016), how to foster sustainable careers paths despite the dichotomy of competition and collaboration (Petersen, Riccaboni, Stanley, & Pammolli, 2012), and how to appropriately use citation metrics for research evaluation (Petersen, Pan, et al., 2019; Wilsdon et al., 2015). Among these challenges is the task of maintaining high quality standards despite the rapid growth of scientific production, especially in light of disinformation campaigns aimed at discrediting science itself (Editorial, 2017; Oreskes & Conway, 2011; Petersen, Vincent, & Westerling, 2019). As the peer-review system continues to evolve new norms and practices, such as publishing open peer review reports, there are new opportunities for improving its fidelity (Bravo, Grimaldo, López-Iñesta, Mehmani, & Squazzoni, 2019).

To this end, a research field focusing on improving our understanding of the peer-review process has emerged, one which aims to provide generalizable insights and practical recommendations for improvement (Batagelj, Ferligoj, & Squazzoni, 2017). Early efforts to measure the prevalence of editor bias in science were motivated by the common perception that editors favor close colleagues, in accordance with the social theory of in-group favouritism (Becker, 1957). While evidence from select economics journals supports this perception (Laband & Piette, 1994; Medoff, 2003), these studies also found that this bias correlated with editors identifying higher impact research, thereby arguing that the ends justifies the means. This correlation appears to be discipline and context specific, as analysis of editorial bias in prestigious university-based law review journals indicates that editorial bias in favor of local faculty produced less-cited work relative to faculty from other law schools (Yoon, 2013); another study of international relations journals produced similar evidence for in-group bias correlating with less-cited research (Reingewertz & Lutmar, 2018). Other types of editor and referee bias derive from the differential review of positive versus negative results (van Lent, Overbeke, & Out, 2014), and gender bias in the selection of reviewers by editors (Helmer, Schottdorf, Neef, & Battaglia, 2017). Considered together, the imperfections associated with the peer review process often derive from the frequent nature of conflicts of interest arising between academic peers (Fong & Wilhite, 2017; Lee, Sugimoto, Zhang, & Cronin, 2013; Wilhite & Fong, 2012; Zaggl, 2017).

Increasingly, yet still in limited cases and extent, researchers are able to peer into the black box of peer-review by obtaining select sets of referee reports from compliant journals. These efforts reveal the previously obscured levels and types of peer-review bias. By way of example, recent work studying a select sample of ~8000 neuroscience manuscripts submitted to PLOS ONE identified reviewer bias associated with the social distance in the collaboration network between reviewers and manuscript authors (Teplitskiy, Acuna, Elamrani-Raoult, Körding, & Evans, 2018). Another study used structured peer-reviews to assess the frequency of disagreement between reviewer and editor recommendations (Kravitz et al., 2010). Other related work used the variation across reviewer recommendations combined with the editor's reject-revise-accept decision to develop a theoretical model for the editor decision-making process, complemented by empirical analysis and survey evidence indicating that reviewer and editor bias actually supports less prolific authors by holding more prolific authors to a higher standard (Card & DellaVigna, 2017).

Against this background, we focus on addressing timely issues at the intersection of the megajournal paradigm, management of distributed academic editorial boards, and the fidelity of the manuscript decision process. In particular, we contribute to the literature stream quantifying the role of social factors on manuscript decisions and decision timescales (Colussi, 2018; Powell, 2016; Sarigöl et al., 2017; Sugimoto, Lariviere, Ni, & Cronin, 2013; Teplitskiy et al., 2018). Our results demonstrate the value of transparency in the editorial process, and provide motivation for implementing transparent oversight of large editorial boards comprised of acting academics.

3. Material and methods

3.1. PLOS ONE article data

We gathered citation data for all PLOS ONE articles, hereafter indexed by *A*, from the Web of Science (WOS) Core Collection. From this data we obtained a master list of the unique digital object identifier, *DOI*_A, as well as the article's page

length, P_A , the number of authors, k_A , a list of their surnames and first-middle name initials, and the number of citations received by the article, c_A , at the time of the data download (census) date on February 25, 2019. We then used each DOI_A to access the corresponding online XML version of each article at PLOS ONE by visiting the unique web address "http://journals.plos.org/plosone/article?id=" + " DOI_A ". Because the full-text XML files have common and relatively stable structure over the 10-year period of analysis, we were able to collect the same metadata for each PLOS ONE article. As such, we parsed each article's reference list, resulting in a dataset of more than 6.7 million outgoing citations. We used the list of coauthor names for each citation to estimate the rate of citations directed at handling editors. Based upon a string match between the full last name and first initial of the coauthor and handling editor, we estimate that 0.3% of citations are directed at the handling editor, and roughly 8% of PLOS ONE articles have at least one such handling-editor citation [see Appendix Fig. 5].

Given that most journals do not make the editor-publication association data available, while also considering the magnitude of the data collection and cleaning effort, it is beyond the scope of our analysis to perform a comprehensive comparison of editor patterns at all megajournals. Thus, our study is primarily a case study of PLOS ONE. Nevertheless, in order to provide an initial comparison of editorial board size and activity, we collected, processed and analyzed editor-publication data for two other journals featuring distributed editor management systems comprised of acting academics: *Proceedings of the National Academy of Sciences of the United States of America* (PNAS), a large multidisciplinary journal publishing 3,799 articles in 2016, and *Management Science* (MS), a medium-sized monthly journal publishing 195 articles in 2016.

For comparison, the volume and editorial board size of each journal for the years in which editor-publication data are available are as follows: (i) PLOS ONE published 141,986 research articles from 2006–2015, relying on the efforts of 6934 editors; (ii) PNAS published 20,816 research articles from 2005–2014 (not including articles "Contributed by" or "Communicated by" PNAS editorial board members, which is a distinct contribution mechanism of PNAS which we excluded from our analysis), relying on the efforts of 2981 editors; (iii) and MS published 780 research articles from 2011–2015, relying on the efforts of 51 editors. The corresponding average editor activity rates are 20.5, 7.0, and 15.3 articles per editor, respectively.

3.2. Article and editor measures

The principal unit of analysis in our study is a PLOS ONE editor, which we denote by the index *E*. For each *E* we collected the corresponding group of N_E articles over which he/she has served as editor. This editor-article (*A*, *E*) association is publicly available in both the published electronic article as well as on the article webpage, appearing in the article abstract and author information byline, which is also the case for PNAS and MS. Nevertheless, we foster privacy by anonymizing the full names of individual editors throughout this analysis, instead referring to particular editors by their activity rank *r*, as in Fig. 2(A). The association of variables is denoted by the following index system: quantities that are mostly article-specific are denoted by the index *A*, those that are mostly editor specific are denoted by the index *E*, and quantities that are properties of both are indexed as $x_{A,E}$.

Embedded in the XML file for each article are various editor, coauthor, and article metadata which we extracted from the webpage of each *A* and then aggregated for each *E*. All together, the entire dataset for the 10-year period 2006–2015 is comprised of 141,986 articles and 6934 editors. In both of our panel regression models we refine this dataset to the 3749 editors with $N_E \ge 10$ articles to reduce small sample noise at the editor level, resulting in 128,734 articles. From these articles and their editors we define the following quantities:

- 1. The net editor activity, N_E , is the number of articles overseen by editor *E* over the total editor service period, L_E , which is the number of days between an editor's first and last article through the end of 2015.
- 2. The article acceptance time, Δ_A , is the number of days between the submission and acceptance of article *A*. Note that this duration does not include the time interval between acceptance and publication, as factors external to the editor process could affect this process, its timeline, and the ultimate duration.
- 3. The annual activity, n_E , is the mean number of articles edited per year while serving as editor at a particular journal, i.e. $n_E = 365N_E/L_E$. The inverse measures the editor turnover time $d_E = L_E/N_E$, which is the mean number of days between two articles published in PLOS ONE, a proxy for the intensity of the time commitment required of a given editor.
- 4. The normalized annual activity, S_E , facilitates comparison of editor activity levels between journals of different size. For a given journal, we calculate the median activity Median $[n_{E,y}]$ for a given year. We then normalize each editor's activity in a given year by the median activity, defining the standardized ratio $S_E = n_{E,y}$ /Median $[n_{E,y}]$ where Median $[n_{E,y}]$ is calculated across all active editors in a given year y for a particular journal.
- 5. The mean acceptance time, Δ_E , is the mean Δ_A calculated for a given editor. Likewise, we measured the variability in Δ_A using the coefficient of variation, $cov_E = \sigma_E[\Delta_A]/\Delta_E$, calculated within each editor's article subset (where $\sigma[\ldots]$ denotes the standard deviation).
- 6. The citation count *C_A* is the total number of references in a given article that cite the handling editor's research. This number is calculated by going through the reference list of each article, and identifying publications that include the editor's last name and first-name initial among the author list. A limitation to this string-matching approach is the misattribution error associated with the ambiguity of abbreviated names, e.g. the editor J. Doe and coauthor J. Doe of a referenced article are in fact distinct individuals. Another issue that contributes measurement error derives from reference styles using the



Fig. 2. Distributions of PLOS ONE editor and article characteristics. (A) The top-100 most active editors ranked according to N_E ; individual names are anonymized and replaced by their rank. Circle color indicates the editor's mean time to acceptance, Δ_E (days); green-yellow values are above the population mean of 130 days (see panel D); magenta-blue values are significantly below the population mean, and are typical of the top-10 editors. (B) The distribution $P(N_E)$ is extremely right-skewed. (C) The distribution of the number of days between an article was received and accepted for publication (i.e. not including the time between acceptance and publication). (D) The distribution of the mean number of days to accept an article calculated for each editor; comparable with panels A and C. (E) The distribution of f_A , the fraction of the references in a given article that cite other papers that include the editor as a coauthor: 92% of papers have $f_A = 0$, but there is an extremely long tail. In panels (D,E) we only included data for the 3749 editors with $N_E \ge 10$ articles in order to reduce the fluctuations due to small sample size; vertical dashed lines indicate distribution mean. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

shorthand abbreviation "et al.", which may obscure the name of the editor. Likewise, the editor citation rate, f_A , is the fraction of the total references for a given article that cite the handling editor's work.

7. The editor's PLOS ONE service age, $\tau = \tau_{A,E}$, is the time difference between the acceptance date of the first accepted article of editor *E* and the acceptance date of a given article *A*, measured in units of years.

Fig. 2 shows the statistical distribution of several important quantities, with article-level statistics shown in blue, and editor-level statistics shown in red throughout the remainder of the analysis; Appendix Fig. 6 shows the distribution of additional editor activity and repeat-author measures.

3.3. Repeated editor-author associations

We investigate whether editor-author associations correlate with publication outcomes by tabulating the set of N_k authors appearing within the article set of a given editor. That is, for each article we recorded the last name and first initial of each of the k_A coauthors. Then, for each editor we tallied the number of articles ($A_{E,k}$) he/she edited for a given author (proxied by the individual surname + first-initial combinations, e.g. "J"+"Smith"). Because of the author name ambiguity problem, it is difficult to distinguish authors with the same name, especially for authors with extremely common surnames. Thus, we removed from our analysis those authors with common surnames (e.g. Xie, Yang, Adams, Johnson), using the PLOS ONE editor name list to determine which surnames appear with significant frequency that might significantly contribute to false-positive union of coauthor counts. We describe this procedure in the Appendix text, where we provide the full list of surnames which we ignored in our editor-author analysis. After tallying the author names associated with each E, we then ranked the list of those coauthors within each editor's profile.⁴

Three recent studies have also estimated the strength of the "social tie" between editor and author using two similar methods, leveraging either prior department, coauthorship, and mentorship histories (Colussi, 2018) or by measuring the distance within the collaboration network (Sarigöl et al., 2017; Teplitskiy et al., 2018). Ideally, one would be able to capture the union of all three methods for estimating the strength of the social relation between editor and author. Due to data limitations this is not feasible, and so we use a simple heuristic: independent of the precise combination of $A_{E,k}$ attributed to the k_A coauthors of a given article, we estimate the presence of social ties between the authors and editor based on whether or not any of the authors have $A_{E,k} \ge 2$ articles with a given editor – i.e. "repeat authors". This definition also accounts for social ties that exist outside of the collaboration network. One clear limitation of this approach is that we do not account for the time dependence of $A_{E,k}$ at the time that each article was submitted to PLOS ONE. Instead, we choose a direct method to implement and document, in which we calculate $A_{E,k}$ using the article data pooled across all years; nevertheless, the temporal uncertainty associated with this approximation only affects the first article in the sequence of articles for each editor-author pair.

Because of its large publication volume, more senior coauthors are more likely to have prior experience publishing with PLOS ONE. As such, it is not unreasonable that the corresponding authors of new submissions would elect editor(s) that previously oversaw their accepted publications. Put another way, it is also unlikely that submitting authors would be assigned the same editor as in previously accepted articles just by chance. In what follows we test for spurious relations using a shuffling method to show that this repeat author criteria identifies a subset of articles that are statistically distinct, thereby supporting the accuracy of this method for identifying significant editor-author associations.

Applying this method, for each editor we tagged the articles featuring at least one repeat author using the indicator variable $R_{A,E} = 1$; the remaining articles are denoted by $R_{A,E} = 0$. In total, the articles with $R_{A,E} = 1$ represent 13.9% of all articles. We denote the number of repeat authors per editor by K_{2E} . Appendix Fig. 6(C) shows that editors have on average 5.2 repeat authors. Likewise, on a per-article basis, Appendix Fig. 6(D) shows that on average 11% of an editor's articles have $R_{A,E} = 1$. However the distribution of the fraction ρ_E is skewed with 10% of editors having 26% or more of their articles with $R_{A,E} = 1$; the median ρ_E value is 0.1. As such, this "repeat author" method identifies a sufficiently large number of articles with $R_{A,E} = 1$ that we can use this binary classification to estimate the impact that social factors have on editor decisions by comparing an editor's article set with $R_{A,E} = 1$ to the counterfactual set with $R_{A,E} = 0$.

4. Results

4.1. Skewed distribution of editor activity

Compared to the more traditional print journals MS and PNAS, the online-only megajournal PLOS ONE exhibits an extremely wide range of annual editor activity, n_E , which measures the mean number of articles edited per year by a given editor. The distribution averages calculated across editors from each journal are: 1.5 (PNAS), 4.4 (MS) and 4.9 (PLOS ONE) accepted articles per year per editor. Notably, we do not include "Contributed track" articles by National Academy of Science members in our analysis of PNAS, focusing only on the "Direct Submission track" articles.

To illustrate the broad range of activity, Fig. 1(A) shows the complementary cumulative distribution $CCDF(\ge n_E)$ conveying the fraction of editors with activity larger than a given n_E level. The activity distributions are generically right skewed – most editors oversee a few articles a year, whereas the most active editors are significantly more active than the average editor within each journal. However, the upper limits of editor activity allude to the distinct productivity differences between these journals. Comparing the larger journals PLOS ONE and PNAS directly, we count 85 editors at PLOS ONE that are more active than the most active PNAS editor who averages 22 articles per year. For a comparison at the distribution level, we calculate the Gini index, which is a useful sample-size independent measure of dispersion or "inequality" across the units of analysis. This standardized measure quantifies the mean difference between all pairs in the population, with larger values indicating higher levels of inequality. The resulting Gini index values are 0.47 (PLOS ONE), 0.40 (MS), 0.36 (PNAS).

To further highlight the anomalous activity of extremely active PLOS ONE editors, we calculated the normalized annual activity, S_E , which controls for journal-specific volume, acceptance criteria and acceptance rates. Fig. 1(B) shows the kernel density estimate of the probability distribution $P(S_E)$, which demonstrates data collapse across the three journals up until the value corresponding to the 95th percentile (indicated by the vertical dashed lines around $S_E \approx 4$). Thus, there are indications for a common distribution of editor activity, net of extrinsic factors, in the bulk of the distribution. Nevertheless, there are

⁴ We investigated the distribution of the rank-coauthor profile within an editor's article set, and found that the distribution $P(A_{E,k})$ decays like a binomial distribution, but with deviations in the tail. The maximum value Max[$A_{E,k}$] depends to a large degree on N_E . Thus, unlike the rank-coauthor distribution within a given researcher's publication profile, which is well-fit by a discrete exponential distribution and characterized by a subset of strong career partners (Petersen, 2015), the editor-author distribution is not characterized by very strong social ties. Thus, we settle on a simple heuristic based upon repeat authors, since it is the repeated interactions facilitated the large volume of articles published that is of interest.

significant differences between the three distributions when comparing the behavior in the right tails corresponding to the extreme values. By way of example, the truncation at the upper tails of $P(S_E)$, corresponding to the most extreme editor activity, is roughly 4 times the median activity for MS; roughly 10 times the median activity level for PNAS; and nearly 20 times the median activity level for PLOS ONE editors.

To demonstrate the extent to which editorial activity scales up over time, we also analyzed the total articles per editor, N_E . Fig. 2(A) shows the 100 most-prolific editors, who collectively oversaw 17,000 (12.2%) of the total 141,986 articles; and Fig. 2(B) shows the extremely right-skewed distribution, $P(N_E)$. Most editors have overseen a reasonable number of articles in their tenure, e.g. the median value of N_E is 11 articles. However, the most prolific editor (Editor r = 1) has served on roughly 27 times (557/20.5) as many articles as the average editor. In terms of the cumulative fraction of all articles edited by a given percentile: the bottom 25% of editors oversaw just 3% of the total 141,986 articles; the middle 65% of editors oversaw 55%; the top 10% of editors (693 editors) oversaw 42%; and the top 10 editors together oversaw 3408 articles, corresponding to 2.4% of all articles analyzed.

Together, these numbers illustrate the remarkable upper limits of editor activity and decision-making power facilitated by the scalability of high-throughput megajournals.

4.2. Skewed distribution of article acceptance time

The number of days between the editor receiving an article and eventually accepting the article, Δ_A , is highly variable [see Fig. 2(C)]. By way of example, close inspection of the 10 most active editors shows that their extreme editor activity is largely explained by their rapid acceptance times, with several editors averaging just around 2 months per article from submission to acceptance. Thus, the variation in Δ_A can partly explain the variation in net editor activity at PLOS ONE – faster editors churn through more submissions.

To further assess the variation at the editor level, we also calculated the mean acceptance time for articles handled by a given editor, Δ_E . Fig. 2(D) shows significant variation at the editor level as well, with the mean values at the editor and article levels in correspondence: $\langle \Delta_E \rangle \approx \langle \Delta_A \rangle$ =126 days or roughly 4 months. As shown above, some editors are significantly faster (slower) than others, as demonstrated by the standard deviation in Δ_E across editors, which is roughly one month (33 days).

Yet even among the 100 most active editors there is significant variation in Δ_E , as demonstrated by the color coding of individual editors in Fig. 2(A). Moreover, variation in Δ_A and n_E also implies that the number of articles being handled at any given time can also vary widely across editors. On average, the time between articles accepted by a given editor, d_E , is about half as long as the time for an article to be accepted ($2\langle d_E \rangle \approx \langle \Delta_A \rangle = 126$ days), meaning that even the average editor is handling roughly two articles at a time. This estimate does not include the additional effort associated with articles that are not ultimately accepted. Because keeping track of multiple tasks requires marginally more effort per task, the effort required to maintain activity at extreme levels would likely be daunting assuming that one maintains reasonable manuscript evaluation standards.

4.3. Modeling variation in article acceptance time within Editor profiles

The average PLOS ONE article takes 126 days from being officially received and processed by the editor, reviewed (possibly over several rounds), and finally accepted. That is, we do not include in our analysis the time between the manuscript being accepted and being published online, which could depend on spurious post-production factors. This characteristic 126-day timescale is higher than the global average across journals which was recently estimated to be roughly 100 days, with only slight variation observed when disaggregating journals by their impact factors (Powell, 2016).

However, we observe extremely wide variation in the acceptance time of individual articles within PLOS ONE – both across and within editor profiles. This variation is highlighted by two remarkable extremes – we observed one publication with $\Delta_A = 0$ days (DOI: https://doi.org/10.1371/journal.pone.0031292) and one publication (DOI: https://doi.org/10.1371/journal.pone.0028904) with $\Delta_A = 1927$ days, or more than 5 years to finally be accepted. Moreover, we find that 0.43% of articles are received and accepted within 7 days, possibly following the rapid transfer of a submission between PLOS journals.⁵

This variation in article acceptance time, which does not purely correlate with editor activity, suggests that editor-author social factors play a significant role – in particular in-group bias (Becker, 1957) relating to a specific subset of researchers. Sarigöl et al. (2017) first pursued this direction by analyzing the degree to which PLOS ONE editors handle papers differently if the manuscript authors are close versus distant within their local collaboration network. Indeed, they find that PLOS ONE editors handle manuscripts by former coauthors more often than what one would expect by random chance, and further show that submissions by prior co-authors tend to be accepted faster, with a reduction in Δ_A of 19 days on average. An

⁵ According to PLOS policy, in a effort shared by other journals to make the peer review system more efficient, manuscripts rejected by one PLOS journal can be nearly instantly transferred for submission to another PLOS journal, including the referee reports, which could in principle be used by an editor to make an immediate decision to accept the article.

important distinction from this work is that here we focus more on identifying editor-specific trends, whereas this initial study uses a pooled model specification that does not include editor fixed effects.

Continuing in this direction, we model the same outcome variable, Δ_A , using a editor fixed-effect model to account for time-independent unobserved variables that distinguish editor profiles. One basic motivation for this model refinement is the extremely broad distribution of editor activity (n_E), which gives rise to the fundamental question – why would an editor take on such extreme levels of service, especially considering the increasing marginal effort associated with handling each additional manuscript?

To a certain degree, variations in auxiliary teaching, institutional service, research load and efficiency across editors can explain variation in N_E , n_E and Δ_E . However, we also observe significant variability within editor profiles [Fig. 6(A)], and so even if the variation is explained by some editors being more efficient than others on average, the is still substantial variation in efficiency among the most proficient editors. To address this within-editor profile variation, we incorporate editor-specific fixed effects into our model specification, which distinguishes the model and the interpretation of the model estimates in comparison with Sarigöl et al. (2017). Thus, in addition to qualitatively reproducing their results, we use two different variables to measure the degree to which social factors explain the broad distribution of Δ_A – both across and within editor profiles. The first social factor is captured by the repeat author indicator variable, $R_{A,E}$, and the second is captured by the fraction of references directed at the editor's publications, $f_A \in [0, 1]$.

The first variable $R_{A,E}$ incorporates formal and informal social network information associating editors with manuscript authors. We assume that manuscript handling bias is more likely to manifest, and likely manifest more strongly, among ingroup members (Becker, 1957) due to trust established in prior interactions. In regards to the frequency of repeat authorship, in Section 3.3 we found that roughly 14% of PLOS ONE articles are authored by individuals who have also previously published via the same editor. We use this information to classify articles using the binary repeat-author indicator variable $R_{A,E}$. Likewise, since on average each editor has 11% of her articles with $R_{A,E} = 1$ and the remaining with $R_{A,E} = 0$ [Fig. 6(D)], we can further test for differences at the editor level according to a first hypothesis:

H1. Articles by "repeat authors" (denoted by the indicator variable $R_{A,E} = 1$) are accepted faster, and this correlation persist within editor profiles.

The second variable f_A incorporates citation network information connecting editor and author. The statistical distribution $P(f_A)$ indicates that 92% of articles do not have any references that cite the handling editor [Fig. 2(G)]. Calculated across all articles, the mean value $\langle f_A \rangle = 0.003$ [Fig. 5]. However, among the remaining 8% of articles with $f_A > 0$, there is a wide range, with the average value $\langle f_A | f_A > 0 \rangle = 0.036$ corresponding to roughly one in every 28 references citing the editor's work. We generalize f_A with the indicator variable I_{SCA} , which takes the value 1 if an article has editor self-citations ($f_A > 0$) and 0 otherwise. This leads to our second hypothesis:

H2. Articles with higher f_A are accepted faster, and this correlation persists within editor profiles.

Based on the results in support of H1 and H2, we test a third hypothesis relating to extreme editor activity. To operationalize this distinction between 10 extremely active editors (denoted by XE), we define an indicator variable which takes the value $I_{XE} = 1$ if the editor is among the 10-most active editors (according to N_E), and takes the value $I_{XE} = 0$ otherwise.⁶ Based on this classification, we begin with preliminary evidence based upon the application of the T-test to measure differences in mean values (indicated below by $\langle ... \rangle$) between the two groups of editors:

- 1. XE accept papers significantly faster, on average by nearly 6 weeks: $\langle \Delta_A | I_{XE} = 1 \rangle \langle \Delta_A | I_{XE} = 0 \rangle = 43.7$ days. T-test results: T=30.116, N=109,377, $p \approx 0$, 95% confidence interval = [40.87, 46.56].
- 2. Papers accepted by XE have significantly higher likelihood of featuring repeat authors: $\langle R_{A,E} | I_{XE} = 1 \rangle \langle R_{A,E} | I_{XE} = 0 \rangle = 0.25$. T-test results: T=32.6, N=109,377, $p \approx 0,95\%$ confidence interval = [0.24, 0.27].
- 3. Papers accepted by XE have significantly higher citation rates directed at their own work: $\langle f_A | I_{XE} = 1 \rangle \langle f_A | I_{XE} = 0 \rangle = 0.0030$. T-test results: T=13.06, N=109,377, $p \approx 0$, 95% confidence interval = [0.0026,0.0035].

Based on these differences, we formulate two additional hypotheses:

H3. The bias associated with repeat authors tested in H1 is stronger in magnitude among the extremely active editors (XE).

H4. The bias associated with self-citations to the Editor's research tested in H2 is stronger in magnitude among XE.

We test H1–H4 by modeling $\ln \Delta_A$, an outcome variable that explicitly depends on the handling editor's action. We leverage the longitudinal nature of the data by including editor-specific time variables to account for trends within each editor profile. More specifically, we employ a ordinary least-squares (OLS) regression framework that leverages three particular control variables in the model specification: (i) editor fixed effects (FE) in order to account for time-invariant unobserved

⁶ Since Editor r = 6 has a rather common last name, we exclude this editor from our analysis (see Appendix), and replaced this individual within the XE set with Editor r = 41, who is subsequently identified in this analysis as an anomalous outlier. Consequently, there remain 10 editors with indicator value $I_{XE} = 1$.

Table 1

OLS modeling of Article acceptance time, Δ_A . The dependent variable is the logarithm of the acceptance time for an individual article (ln Δ_A). Only the editors with $N_E \ge 10$ are analyzed, corresponding to 3144 distinct editors profiles. Only Models B & C include editor fixed effects (FE), whereas all models are estimated using robust standard errors. Test variable coefficients highlighted in color are shown together in Fig. 3.

Model	(A) No Editor FE	(B) Editor FE	(C) Editor FE [w/ Interaction]	(D) No Editor FE [w/ Interaction]
Test variables Editor self-citations, β_f	-0.834* (0.347)	-0.459** (0.160)	-0.455** (0.162)	[III Interaction]
Editor self-citation, β_{SC} (0/1 indicator variable)				-0.0190** (0.00700)
Repeat author, β_R (0/1 indicator variable)	-0.0885*** (0.00841)	-0.0370*** (0.00435)	-0.0339*** (0.00426)	-0.0753***
Extreme Editor, β_{XE} (0/1 indicator variable)				-0.419*** (0.0139)
Extreme Editor \times Repeat author, $\beta_{XE \times R}$			-0.0713* (0.0305)	
Extreme Editor × Editor self-citation, $\beta_{XE \times SC}$				-0.131** (0.0417)
Control variables				
Article page-length, β_p	0.139*** (0.00772)	0.162*** (0.00604)	0.163*** (0.00605)	0.140*** (0.00573)
Article # coauthors, β_k	0.0292*** (0.00621)	0.0495*** (0.00343)	0.0496*** (0.00342)	0.0296*** (0.00323)
Article citation impact, β_z	-0.0401*** (0.00229)	-0.0485*** (0.00178)	-0.0484*** (0.00179)	-0.0400*** (0.00180)
Editor career year, β_τ	0.0112** (0.00397)	-0.00830 (0.00634)	-0.00831 (0.00634)	0.0194*** (0.00124)
Editor Fixed Effects	N	Y	Y	N
Year Dummy	Y	Y	Y	Y
Subject Area Dummy	Y	Y	Y	Y
N	109377	109377	109377	109377
adj. R^2	0.072	0.059	0.059	0.086
F	86.24	170.5	162.9	345.0
DF_{model}	20	19	20	22
Standard errors in parentheses below	point estimate.			
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.01$	0.001			

characteristics relating to each individual editor (captured by β_E); (ii) year fixed effects to account for shifts in the PLOS ONE manuscript submission and handling system among other time-dependent shocks (captured by D_t); (iii) and the amount of time in years since each editor began with PLOS ONE (captured by $\tau_{A,E}$), to control for temporal career trends (e.g. increasing business and prominence) across the editor's career. We focus our panel regression model on the editors with $N_E \ge 10$ so that on average each editor has at least one article with $R_{A,E}$, and those editors with sufficiently unique surnames that we can attribute citations to their work, which reduces the dataset from 141,986 to 109,377 articles (observations) and from 6934 to 3144 editors (observation clusters).

The specification of our linear fixed-effects model is thus given in compact form by

$$\ln \Delta_A = \beta_f f_A + \beta_R R_{A,E} + \beta \cdot \vec{x}_{\text{controls}} + \beta_E + \epsilon_{A,E}, \tag{1}$$

where $\vec{\beta} \cdot \vec{x}_{controls} = \beta_p \ln P_A + \beta_k \ln k_A + \beta_z \ln z_A + \beta_\tau \tau_{A,E} + D_s + D_t$ represents 6 article and editor-level controls to account for additional sources of variation in the dependent variable. The article-level variable P_A , the page length of the article, controls for the increasing effort required to review longer articles. The article-level variable k_A , the coauthor number, controls for team-size effects and is incorporated in logarithm since the distribution of authors per publication is rightskewed and approximately log-normal in various team-oriented disciplines (Petersen, Pavlidis, et al., 2014). The articlelevel variable z_A , the detrended citation impact which is standardized across time and discipline [see Appendix], accounts for variation in relative research quality and author reputation. And the editor-level variable $\tau_{A,E}$ controls for variation in the duration of each editor's service at PLOS ONE and learning associated with the journal's submission system and other tasks associated with serving as editor. We include subject area and publication year dummies variables to further control for disciplinary and temporal variation, respectively.⁷ Following from the properties of logarithms, the coefficients for independent variables that enter with a logarithmic transform (e.g. ln k_A) have the following interpretation: a one percent shift in the independent variable corresponds to a β % change in Δ_A (again without the logarithm).

⁷ We also ran the models without including subject area (SA) fixed effects, and did not observe a significant difference in the estimates of either model. This surprising result can be explained as a combination of the citation rates between SA being rather uniform (see Appendix Fig. 8) in addition to the fact that the editor fixed-effects approximately control for the variation in subject areas across individual articles if one assumes that editors do not expand their scope over time by handling articles from outside their principal research area.



Fig. 3. Regression model results. (A) Reduction in article acceptance time converted to % effects associated with the model variables capturing social factors between editor and author. See Table 1 for full set of model parameter estimates. Asterisks indicate significance level of point estimate: *p < 0.05; **p < 0.01; ***p < 0.001. (B,C) In order to rule out the possibility that spurious correlations are responsible for the significant estimates in (A), we shuffled the repeat author variable across all observations in the dataset without replacement (i.e. conserving the total number of observations of a given value), and then recorded the relevant coefficient estimate for each randomization. We ran this placebo model using 10,000 different randomized $R_{A,E}$ configurations, and plot the distributions of "placebo" estimates, $P(\beta_R)$ and $P(\beta_{XE\times R})$. The colored horizontal lines display the real estimates (dashed) along with their corresponding 95% confidence intervals (solid). (B) Out of 10,000 different randomization trials of Model (C), only 0.4% of the placebo models yield an estimate that is smaller than the real estimate $\beta_{XE\times R} = -0.071$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1 reports all model parameter estimates for four incrementally variant models, which we estimated using robust standard errors. We ran variations of the same model to demonstrate the robustness of our specification and to explore the additional interactions relating to H3 and H4. Model (A) does not include editor fixed-effects, and as such is a standard pooled regression similar to the model specification in (Sarigöl et al., 2017), which models the variation across editor profiles. Model (B) introduces editor fixed-effects, which alters the interpretation of the coefficients as they represent effects that are net of editor-specific time-independent average values (i.e. within-profile estimates as opposed to the across-profile estimates produced by Model A). Model (C) modifies Model (B) by introducing an additional interaction term $\beta_{XE\times R}(I_{XE} \times R_{A,E})$ into Eq. (1) in order to further test H3. And finally, to test H4, Model (D) modifies Model (A) by introducing an additional interaction term $\beta_{XE\times SE}(I_{XE} \times I_{SC})$ into Eq. (1), and replacing the covariate f_A with the indicator variable I_{SC} to simplify the interpretation of $\beta_{XE\times SE}$ in real terms. Importantly, Model (D) does not include editor fixed-effects because this model seeks to explain differences across profiles, in order to provide additional evidence that extreme editors are motivated by self-citations to a larger degree than their peers.

Importantly, the dependent variable enters the linear equation in logarithm in order to temper the skewed underlying distribution. As a result, the coefficients β_f , β_τ , β_R , β_{XE} and β_{SC} which do not enter in logarithm have the following interpretation: a one unit shift in the independent variable corresponds to a $100 \times \beta^{\%}$ percent change in Δ_A (i.e. without log). With this in mind, we focus our estimation summary on Model (B) which includes editor fixed effects, and highlight the main results in Fig. 3. First, we estimate the coefficient $\beta_R = -0.037$ (p < 0.001, 95% CI = [-0.046, -0.028]), which is in support of our H1. In real terms, this means that articles authored by repeat authors (with $R_{A,E} = 1$) have a $100|\beta_R| = 3.7\%$ reduction in acceptance time relative to the other articles (with $R_{A,E} = 0$) overseen by the same editor. Using the results of Model (A) to estimate this decrease for the average article, this effect corresponds to roughly a $\langle \Delta_A \rangle (1 - \exp[\beta_R]) \approx 126 * 0.0885 \approx 11$ -day decrease in acceptance time related to the shift in $R_{A,E}$ from 0 to 1.

Second, we estimate the coefficient $\beta_f = -0.459$ (p = 0.004, 95% CI = [-.77, -.14]). In real terms, this means that articles with 10% of their references citing the editor's research have a $0.1 \times 100 \times |\beta_f| = 4.6\%$ reduction in acceptance time relative to the other articles (with $R_{A,E} = 0$ and $f_A = 0$) overseen by the same editor.

In Models (C) and (D) we further test for differences between the 10 extremely active editors and the rest of PLOS ONE editors. The specification in Model (C) tests H3, yielding the baseline coefficient $\beta_R = -0.034$ (p < 0.001, 95% CI = [-.042, -.026]) and the interaction coefficient $\beta_{XE\times R} = -0.071$ (p = 0.019, 95% CI = [-.13, -.01]). In real terms, this indicates that among the majority of PLOS ONE editors, there is only a $100 \times |\beta_R| = 3.3\%$ reduction in article acceptance times associated with repeat

authors ($R_{A,E} = 1$); However, this manuscript bias is larger among the XE, for which we observe a $100 \times |\beta_R + \beta_{XE \times R}| = 10\%$ reduction in Δ_A , corresponding to roughly 12 days for the average article.

Similarly, the specification in Model (D) tests H4, yielding the baseline coefficient $\beta_{SC} = -0.019$ (p = 0.007, 95% CI = [-0.033, -0.005]) and the interaction coefficient $\beta_{XE\times SC} = -0.13$ (p = 0.002, 95% CI = [-0.21, -0.05]). In real terms, this indicates that among the majority of PLOS ONE editors, there is only a $100|\beta_{SC}| = 1.9\%$ reduction in article acceptance times associated with an article having at least one reference citing the handling editor's research ($I_{SCA} = 1$); However, this manuscript bias is significantly larger among the XE, for which we observe a $100 \times |\beta_{SC} + \beta_{XE\times SC}| = 15\%$ reduction in Δ_A , corresponding to roughly 18 days for the average article.

Among the other control variables, we observe consistent parameter estimates. Longer articles correlate with longer acceptance times ($\beta_p = 0.16$, p < 0.001, 95% CI=[0.15, 0.17]); in real terms this indicates that a 1% increase in article length correlates with a 0.16% increase in Δ_A . Higher impact articles tend to get accepted more quickly ($\beta_z = -.0484594$, p < 0.001, 95% CI=[-.052, -.045]), likely because higher quality research is more easy to identify, and so there is a faster consensus towards a decision to accept. Another factor possibly contributing to this estimate for β_z is recent work showing that higher author centrality in collaboration networks may confer a considerable citation advantage (Sarigöl, Pfitzner, Scholtes, Garas, & Schweitzer, 2014), an advantage which may extend to the speed of the peer-review process as well. And finally, more coauthors also correlates with acceptance time, in line with expected increasing coordination costs in assembling and submitting referee revisions in large team endeavors ($\beta_k = .050$, p < 0.001, 95% CI=[.043,.056]).

We conclude this section with a procedural robustness check, as it is possible that our reported estimates for the β_R and $\beta_{XE\times R}$ coefficients could arise by chance due to a spurious correlation associated with the (mis)classification of articles according to $R_{A,E}$. In order to demonstrate that our estimates do not arise as a result of chance configuration of $R_{A,E}$, we implemented a randomization scheme in which we shuffled the values of $R_{A,E}$ across the dataset, without replacement – thereby conserving the total number of observations with $R_{A,E} = 1$. We implemented this "placebo" regression 10,000 times for model (B) and 10,000 times for model (C), each time recording the placebo estimates β_R and $\beta_{XE\times R}$, respectively. Fig. 3(B) shows the distribution $P(\beta_R)$ calculated for 10,000 randomizations; indeed, we do not observe a single placebo estimate smaller than the real estimate, thereby ruling out the possibility that β_R is significant due to chance alone. Fig. 3(C) shows the distribution $P(\beta_{XE\times R})$, which also indicates that it is unlikely that the real estimate $\beta_{XE\times R} = -0.07$ arose by chance alone, as only 0.4% of the placebo models yield an estimate that is smaller than the real estimate.

4.4. Estimating the scalability of editor self-citations

In this section we estimate the total amount of citations directed at the handling editor's research.⁸ One limitation of our approach is that we do not assess the context or appropriateness of each citation directed at the handling editors' research. Thus our baseline is to assume that the majority of the citations directed at the handling editor's research follow the same intent purposes of any other reference (Bornmann & Daniel, 2008; Tahamtan, Afshar, & Ahamdzadeh, 2016; Vieira & Gomes, 2010). Nevertheless, it is well-documented that citation attribution is susceptible to factors that undermine the credit system in science, such as unjustified self-citation (Costas, van Leeuwen, & Bordons, 2010; Fowler & Aksnes, 2007; Ioannidis, Baas, Klavans, & Boyack, 2019; Seeber, Cattaneo, Meoli, & Malighetti, 2019), reciprocal-citation (Zaggl, 2017) and coerced citation (Fong & Wilhite, 2017; Wilhite & Fong, 2012). Indeed, in the previous section we provide evidence that when authors cite the handling editor's research – either via an unsolicited nudge or a coordinated remuneration – such behaviour could entice a faster and more positive decision; moreover, we demonstrated that the prevalence of this phenomena is strongest among the XE. We now shift to the following question that addresses the possibility of perverse incentives underlying the emergence of excessive editor activity in a megajournal lacking editorial board oversight – how many editor citations could such activity possibly produce at scale?

To address this question, we leverage the size of the PLOS ONE dataset to identify measurable differences in the citation rate to editors conditional on the article including or not including repeat authors (i.e. comparing articles classified according to $R_{A,E} = 1$ or 0, respectively). Thus, for each editor we collected the set of $N_{E,R=1}$ articles with $R_{A,E} = 1$ and counted the total number of references $C_{R=1}$ made by this set of articles, and also the number of those references citing the editor's work, $C_{E,R=1}$. Similarly, for the set of $N_{E,R=0} = N_E - N_{E,R=1}$ articles with $R_{A,E} = 0$, we also calculated $C_{R=0}$ and $C_{E,R=0}$. Thus, the total number of references from all articles overseen by an editor is simply $T_E = C_{R=0} + C_{R=1}$, and the total number of citation received by the editor, independent of R, is $C_E = C_{E,R=1} + C_{E,R=0} = f_E T_E$.

We then define the conditional editor citation rates $f_{E,1} = C_{E,R=1}/C_{R=1}$ and $f_{E,0} = C_{E,R=0}/C_{R=0}$ and plot their distributions $P(f_E|R_{A,E} = 0, 1)$ in Fig. 4(A). The mean value for repeat authors $\langle f_{E,R=1} \rangle = 0.0041$ is 46% larger than $\langle f_{E,R=0} \rangle = 0.0028$, and the probability distribution $P(f_E|R_{A,E} = 1)$ shows a prominent excess in the right tail, suggesting that a sufficiently large f may be an enticing nudge. Application of standard tests for difference in means (T-test), difference in median (Mann–Whitney test),

⁸ Citations directed at the editor's work in order to entice favorable results could be the result of coordinated or uncoordinated actions by the authors with the editor. With the PLOS ONE submission system, authors recommend specific editors to handle their manuscript, but do not necessarily know the editor's identity until the manuscript is accepted; however, it is not impossible that authors and the handling editor could communicate externally. Indeed, this is a generic possibility present with any journal, and not specific to any particular aspect of the review process at PLOS ONE.



Fig. 4. Identifying anomalous levels of editor self-citations. (A) Conditional citation rate distributions for the articles without any repeat authors $P(f_E|R=0)$ (magenta), and for the articles with repeat authors $P(f_E|R=1)$ (green); vertical dashed lines indicate the mean value. (B) Distribution of the excess editor citations associated with repeat-author bias, estimated using the empirical editor-specific difference $\Delta C_E \propto (f_{E,1} - f_{E,0})$ for each author (see Eq. (2)). For visual comparison, we show the normal distribution (gray curve) with mean and standard deviation equal to the empirical data; the outermost red dashed lines indicate the confidence intervals corresponding to $\langle \Delta C_E \rangle \pm 3\sigma_{\Delta C}$. The asymmetry in the tails of the distribution are evident when considering the 3σ outliers, with 34 observations in the right tail and only 2 in the left tail. (C) Scatter plot of editor activity (n_E) and net editor self-citations (C_E), with color indicating whether the mean citation impact $\langle z_E \rangle$ is significantly above (cyan) or below (orange) average; editors with $\langle z_E \rangle$ not differentiable from 0 at the p > 0.06 significance level are colored grey. Two of the three anomalous editors (AE) are simultaneously outliers in 3 categories. (D–F) Mean number of references citing each editor's publications in: (red) PLOS ONE articles overseen by the editor (editor self-citations); (yellow) articles authored by the editor (traditional self-citations); (gray) all other articles that cite the Editor's research at least once. Percent values indicate the percent increase in a given in this article.)

and difference in distribution (Kolmogorov–Smirnov test) all reject the null hypothesis that the data with $R_{A,E} = 1$ and $R_{A,E} = 0$ are statistically similar (at the $p < 10^{-9}$ level in each case).

Do these significant differences persist at the editor level? To address this question we calculated the expected number of citations that an editor might gain due to the differences in citing behavior of repeat versus non-repeat authors. We measure this difference as

$$\Delta C_E = (f_{E,1} - f_{E,0})T_E,\tag{2}$$

which should be equal to 0 for those editors who are completely unbiased with respect to *R*. Deviations from 0 naturally arise due to the finite sample size, N_E . Fig. 4(B) shows the probability distribution $P(\Delta C_E)$ with mean value $\langle \Delta C_E \rangle$ =3.1 and standard deviation $\sigma_{\Delta C}$ = 15.2. However, the distribution is leptokurtic and significantly right-skewed (skewness = 3.5) as compared to the normal distribution with the same mean and standard deviation. In particular, the skew points to an excess number of editors with relatively large and positive number of citations attributable to differences in the citation rates for R = 1 versus R = 0. Remarkably, we count 34 editors with $\Delta C_E > \langle \Delta C_E \rangle + 3\sigma_{\Delta C}$, representing 2% of the 1595 editors we analyzed with $N_E \ge 20$; however, we count only 2 editors with $\Delta C_E < \langle \Delta C_E \rangle - 3\sigma_{\Delta C}$. The positive outlier ΔC_E values are on the order of 100 citations. This number serves as a lower bound estimate for a hypothetical net gains, $C_{\text{remuneration}}$, associated with editor self-citations achieved at scale, i.e. $\Delta C_E \le C_{\text{remuneration}} \le C_{E,R=1}$.

While it is possible that statistical outliers could arise by chance, it is unlikely that an editor would simultaneously be an outlier in various categories. To this end, we collected three measures for each editor: editor activity, n_E , the average citation impact of articles overseen by each editor, $\langle z_E \rangle$, and the total citations directed at the handling editor's research, C_E . We restrict this analysis to the subset of 1595 editors with $N_E \ge 20$ articles. Fig. 4(C) shows the editor-level information across these three variables, indicating no significant relation between the C_E and n_E variables. In order to include information on the citation impact of the articles accepted by each editor, we calculated a T-statistic to assess the degree to which the mean normalized citation impact is significantly different from 0, a baseline value which corresponds to the average logcitations for articles from the same year. As such, the color of each data point indicates whether $\langle z_E \rangle$ is significantly above 0 (cyan), below 0 (orange), or not significantly different than the mean baseline 0 (grey); we use the false positive rate p = 0.06 as the threshold for T-test statistical significance. Accordingly, we identify 241 editors with $\langle z_E \rangle > 0.1$ and 215 editors with $\langle z_E \rangle < -0.1$.

In particular, we identify three editors who are anomalous outliers when all three variables are considered together. Two of these three editors, r=1 and r=7, are among the 10 most active editors and the third is the r=41 most-active editor.



Fig. 5. Growth of PLOS ONE: publications and references produced. (A) Number of publications and references produced by PLOS ONE by year. (B) Percent of references in PLOS ONE articles that cite the handling editor by year. (C) Percent of PLOS ONE articles that have at least one reference citing the handling editor by year. Horizontal dashed lines correspond to the mean value calculated across all years.

In terms of manuscript quality control, we also find that the articles these 3 editors accepted were cited significantly less than other PLOS ONE articles.⁹ This additional bias in quality assessment is reminiscent of the perverse outcomes identified by Yoon (2013) relating to editorial favoritism in prestigious law reviews. Yet despite this all, these editors each obtained several hundred citations from the articles they oversaw (C_E) – more than 10 times greater than the average across other editors, $\langle C_E \rangle$ =8.6 citations.

4.5. Additional evidence: Case study of self-citation patterns for the 3 anomalous editors

Up until now we have only considered articles published at PLOS ONE as sources of editor self-citations. In order to provide more comprehensive evidence for self-citation motivations, in this section we test for differences in citation rates for articles under the influence of each editor compared to articles not under their influence. To this end, we downloaded and analyzed the complete career publication records for each of these 3 anomalous editors (AE), available from the Web of Science (WOS); We also downloaded data for every article in WOS that cited their research at least once. Combining these two datasets for each editor, we then separated the citations they received into three groups – citations from:

- (i) the set of PLOS ONE articles each editor handled (indicator of self citations associated with editor service, red bar);
- (ii) the editor's own articles (indicator of traditional self-citations, yellow bar);
- (iii) the set of citing articles not belonging to group (i) or (ii), denoted as "other" (grey bar).

Fig. 4(D–F) shows the mean citations per publication for articles in each group. By way of example, Anomalous Editor 3 (AE3, corresponding to activity rank r = 1) cites his own work on average 12.4 times in each of his publications, while "other" researchers typically cite his work 1.7 times per article. Using the citation rate by "other" researchers as the baseline, then AE3 self-cites 643% more than the average article that cites his work; this excessive self-citation rate is also observed AE1 and AE2. Also, within the set of PLOS ONE articles overseen by AE3 (corresponding to r = 1), we calculate an average rate of

⁹ We calculated the two-sample T-test comparing the mean *z* value (measuring detrended citation impact) between articles handled by these 3 anomalous editors and articles handled by the remaining editors: $\langle z_A^s | \text{AnomalousEditor} \rangle - \langle z_A^s | \text{Rest} \rangle = -0.22$. T-test results: T = 6.81, N=109,377, $p \approx 0,95\%$ confidence interval = [-0.28, -0.16]. Since z_A^s is a logarithmic transform of citation counts, this corresponds to roughly a -22% difference in citations between the two groups, on average. See Fig. 8 for the correspondence between z_A^s and nominal citation counts.



Fig. 6. Editor-level article characteristic distributions. (A) Probability distribution $P(cov_E)$ of the variability in Δ_E expressed as the coefficient of variation (the ratio of the standard deviation of Δ_E normalized by the mean Δ_E for a given editor). For most editors the mean Δ_E is rather characteristic, however some editors show a wide range of variability. (B) The distribution of the turnover time (or inverse activity) defined as the average number of days d_E between articles accepted by the same editor. (C) Probability distribution $P(K2_E)$ of the number $K2_E$ of repeat authors, i.e. the authors that have appeared on 2 or more of the N_E articles within a given editor set. (D) Probability distribution $P(\rho_E)$ of the fraction ρ of the total articles of a given editor sample size fluctuations; vertical dashed lines indicate distribution mean.

2.3 citations per article, which is 39% more than his baseline citation rate; this was the smallest percentage excess for PLOS ONE edited articles observed among the three anomalous editors.

For comparison, we juxtapose each anomalous editor with a second active PLOS ONE editor from the same research area. For this comparison editor, we also calculated the corresponding citation rates from their complete publication records. The comparison editor serves as a *counterfactual*, an editor from the same discipline and with relatively high total editor activity, but not excessive: Counterfactual E1 (N_{E1} = 93 articles), Counterfactual E2 (N_{E2} = 74), and Counterfactual E3 (N_{E3} = 73). For each comparison editor, the citation rate for "other" articles and PLOS ONE edited articles are quite similar, whereas this is not the case for the anomalous editors. Moreover, the *Editor remuneration citation rate* for anomalous editors is on par with the *self-citation rate* of their corresponding comparison editor. As an external validation of this self-citation assessment, we use a large database of scholarly citation totals compiled and shared openly by loannidis et al. (2019). This database tracks citations between articles indexed by Scopus, and includes an estimate for the percent of references in each individual's articles that are citing their own research, i.e. self-citations, and represented by the "self%" field in the supplementary dataset provided in (loannidis et al., 2019). Using this data, we also identified a large representative pool of researchers from the same broad research area as each AE, based upon overlap between their top two Science-Metrix categories (based upon the "sm-1" and "sm-2" fields). Within this pool of researchers we then computed the self-citation percentile of each AE – representing the percentage of researchers within the pool with self-citation rates lower than the particular AE.

This exercise serves as a consistency check, while also demonstrating supplementary evidence and methods for evaluating the self-citation motivations among extremely active editors. The self-citation percentage (and corresponding sample percentiles in parenthesis) for each AE are: 13.47% self-citation rate (54th percentile) for AE1; 35.36% (99th percentile) for AE2; and 25.11% (95th percentile) for AE3. In each case, the AE has a self-citation rate higher than the median value for researchers from his/her broad research area, and in the most extreme case of AE2, out of 11,796 researchers from the same research area, only 93 have a higher self-citation rate.

All together, we demonstrate the scalability of editor citation remuneration, which can yield significant returns when the excess citations for edited articles are compounded by excessive editorial activity. This advantage could compound into even



Detrended citation impact, zA

Fig. 7. Distribution of article characteristics: subject area and detrended citation impact. (A) Count distribution of the number of articles by principal subject area (no articles were observed with "Science policy" (SA = 10) as the principle SA). (B) Count distribution by SA after applying redistribution rule that if the principal SA = 1, then use the SA with the second-highest weight. (C) Count distribution by SA after merging into 6 refined subject areas, which are used throughout the analysis. (D) Empirical probability distribution P(z|s, t) for each SA and year combination (red bins) and baseline normal distribution N(0, 1) (blue curve) shown to demonstrate the time-independence of the normalized citation impact variable. Since all 2006 articles were published in December, we merged these publications with 2007. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

more citations due to the effects of reputation and cumulative advantage across the career (De Solla Price, 1976; Petersen, Fortunato, et al., 2014; Petersen & Penner, 2014). Moreover, this analysis of editor remuneration only compares articles handled at PLOS ONE and does not include editor activity and referee service at other journals. Indeed, based upon manual inspection, we confirm that several of the extremely active editors simultaneously served on the editor or academic advisory board for various other megajournals such as PeerJ, Palgrave Communications, Royal Society Open Science, Frontiers and Scientific Reports.



Fig. 8. Correspondence: citations - versus - Detrended citations. Logarithmic relation between $c_{A,t}^s$ and z_A^s defined in Eq. (1), shown for publications from 3 annual cohorts. Surprisingly, there is less substantial variation between the refined subject areas than expected. Nevertheless, for the sake of methodological completeness and robustness, we maintain the classification by refined subject areas throughout the analysis.

5. Summary and discussion

Little is known about how the ecosystem of megajournals have altered the production and consumption of peer-reviewed research. To address this gap, we analyzed comprehensive PLOS ONE article metadata over the inaugural 10-year period 2006–2015, leveraging this publisher's transparency-oriented policy of listing the handling editor on each individual article in order to operationalize a large-scale analysis of scientific gatekeepers. Our results reveal an extremely right-skewed distribution of editor activity, which is dominated by a relatively small number of 10 extremely active editors (denoted collectively by XE). Furthermore, by comparing the distribution of editor activity levels at PLOS ONE with two other journals featuring distributed academic editorial boards (*MS* and *PNAS*), we are able to objectively identify anomalous editor activity levels at PLOS ONE [Fig. 1].

Conflicts of interest affecting the neutrality of the peer-review process are not unique to PLOS ONE or any other megajournal for that matter. Indeed, manuscript handling bias likely occurs at all journals at hitherto imperceptible rates since it is not common to list the particular handling editor for each article. Highly concentrated editor activity is also implicit in small journals that feature a single editor or small team of editors. Yet the principle distinction we address here is a scenario in which a high-throughput megajournal lacks appropriate editorial board oversight – which may provide perverse incentives for editors to scale up manuscript handling activity to unreasonable levels.

In particular, the PLOS ONE editorial board is comprised of acting academics who serve a double role as gatekeepers and producers of knowledge. This raises cause for concern based upon research showing that when the acceptability of misconduct increases gradually in power-driven environments, that a "slippery-slope effect" (Gino & Bazerman, 2009; Gino et al., 2011; Malhotra & Gino, 2011) may facilitate the spread of misconduct, which may be inevitable even among individuals who initially had good intentions. Scientific actors in gatekeeper positions may be particularly susceptible to subtle forces of misconduct, because in addition to being a process mediated by negotiation, the information concerning the review process is tightly concealed. Ironically, instead of protecting the system, this lack of transparency may harbor the evolution of strategies for "gaming the scientific system" (Biagioli, Kenney, Martin, & Walsh, 2019).

Against this background, we developed various complementary methods for assessing activity levels across the entire editorial board comprised of nearly 7000 editors. We started with a descriptive approach, focusing on the 100 most active editors, ranked according to the total number of articles N_E . Even among this highly active set of editors we observed a wide range of mean acceptance times, ranging from $\Delta_E = 175$ days to as short as $\Delta_E = 56$ days [Fig. 2A]. Notably, several editors with the shortest Δ_E were among the 10 most active editors. For example, articles edited by the most active editor (corresponding to activity rank r = 1), appeared on average every 3.2 days in PLOS ONE over the 10 year analysis period. This extreme activity follows from this editor's relatively short acceptance time of $\Delta_E = 77$ days, as compared to the editor average of 130 days. The variability in acceptance time within each editor's profile was also high [Fig. 7]. One potential explanation for the prevalence of such short review times is the transferability of reviews from other PLOS journals, which can then be used in the PLOS ONE editor decision process (Bjork, 2015). Journals that have similar transferability policies should consider monitoring its use, in particular the frequency in which it results in extremely short acceptance times; by way of example, roughly 1 out of every 200 articles accepted by PLOS ONE were received and accepted within 7 days.

Analysis of the 10 most-active editors (representing just 0.14% of the PLOS ONE editor population size, who nevertheless together oversaw 2.4% of all articles analyzed) revealed significant differences when compared with the remaining set editors, including: relatively short acceptance times, relatively high rates of repeat authorship and relatively high rates of citations directed at the handling editor's research. It is unlikely that these observations all occur by chance. Thus, we pursued a panel regression modeling framework to measure the role of additional social factors associated with editor decision-making process. This approach leverages the longitudinal features of the data in order to explain variation within each editor profile. More specifically, we employ a regression model specification that includes editor fixed effects, which

controls for unobserved time-independent factors, thereby emphasizing variations net of the editor-specific baseline. Eq. (1) specifies our OLS regression model, in which the dependent variable is the log of the article acceptance time ($\ln \Delta_A$). We hypothesize that variations in acceptance time correlate with two particular factors: repeated submission of manuscripts by the same authors H1 and higher rates of citations to the handling editor's research H2.

There are two fundamental reasons why authors might cite the handling editor's work. First, PLOS ONE assigns articles to editors with relevant expertise by way of a subject area category system that facilitates identifying suitable editors, and so in principal the editor's published research is within the scope of the submitter's research topic. Second, editors are in principle appointed according to their prestige within the community, and so they are likely to have prominent research. Despite these considerations, we provide several lines of evidence showing how editor behavior is correlated with 'citation remuneration' – citations to the editor's own research occurring simultaneously in the article he/she is overseeing as handling editor – an analog to the traditional "self-citation" (Costas et al., 2010; Fowler & Aksnes, 2007; Ioannidis et al., 2019). There are various explanations for relatively high self-citation rates, such as signaling prestige in cross-disciplinary mobility (Hellsten, Lambiotte, Scharnhorst, & Ausloos, 2007), as well as bias towards citing one's past collaborators.¹⁰ However, unlike traditional forms of self, reciprocal, or coercive citation rigging (Fong & Wilhite, 2017; Wilhite & Fong, 2012), the incentives to sway an editor's decision are significantly larger.

The results of our panel regression provide support for H1 – articles by "repeat authors" are accepted faster – and H2 – articles with higher rates of citations to the editor's work are also accepted faster. Both of these results persist within editor profiles, which provides additional support for causal identification of these effects. As such, the difference in acceptance times for repeat authors provides further evidence that citations are indeed an effective form of remuneration (Fowler & Aksnes, 2007). We also tested wether H1 and H2 are amplified among the subset of extremely active editors, relative to the rest of the PLOS ONE editors. Indeed, we observe statistical support for H3 – the correlation between repeat authorship and faster acceptance time is even stronger among the XE; we also observe statistical support for H4 – the correlation between editor self-citations and faster acceptance time is even stronger among the 10 most-active editors.

From the alternative gatekeeper perspective: what does an editor stand to gain from such extreme activity? A formidable challenge is to determine whether differences among editors derives from either apathy or misconduct, which is beyond the scope of our analysis, however statistical support for H3 and H4 point to the latter. Because it is extremely difficult to measure and interpret the motivation and context associated with individual citations, we leverage the multi-dimensional information within our dataset to measure the interrelation between two social factors by measuring how citations directed at the handling editor's research (measured by f_E) relates to the authors having had repeated interactions with the editor. To this end, Fig. 4 summarizes multiple lines of evidence consistent with XE leveraging the high-throughput volume and insufficient editorial board oversight for their own benefit. For example, by comparing the citation rate for articles with and without repeat authors, we arrive at a lower bound for the impact of citation remuneration in the hundreds of citations. This quantity may seem like a relatively small amount to some, however note that this citation total represents revenue from just one source – not including editorial board service at other journals in addition to auxiliary referee service.

We provide additional evidence by performing in-depth analysis of three editors with anomalous activity levels identified in Fig. 4(C). For each of these editors (denoted by AE), we combined the citation statistics for the set of PLOS ONE articles they handled with the citation statistics associated with the rest of their research portfolio. By combining these multiple reference points, we compare the rate of self-citations to their own research from two distinct sources: the citation rate from the PLOS ONE articles they handled, and the baseline citation rate from articles that they did not handle. We incorporated a second counterfactual baseline by comparing the differential self-citation rates against the same rates calculated for similar PLOS ONE editors matched by research area and activity. As a final robustness check, we used self-citation data constructed and shared openly by loannidis et al. (2019) indicating that two AE are in the 95th percentile or greater of self-citation rates within their respective research areas. The results of this analysis provide additional evidence that these AE are motivated to extreme activity levels by self-citation strategies, including leveraging their editorial power at PLOS ONE and possibly at other megajournals where they serve as editors.

We digress by discussing alternative explanations and limitations of our data and methods. First, we lack information concerning the quality of the referee reports, which could additionally explain variation in acceptance time scales. Second, we lack data on the rejected manuscripts handled by each editor. Variation in acceptance rates across editors could manifest as biased model estimates if the correlation between rejection rate and other model variables is strong or time-dependent. Third, we do not account for variations in research area distance between editor and the topic of each article, which could also explain variations in article acceptance times. Fourth, we established social ties between author and editor based upon statistical arguments, since it is unfeasible to account for all possible social relations and the variation in their strengths. For sake of simplicity, we estimated the existence of informal social ties using an indicator variable $R_{A,E} = 1$ that tags articles containing at least one author that published two or more times with a given editor. In small research areas, with lower citation rates and lower representation among the editorial board, it is possible that there would only be a single editor with expertise in the area, making repeat interactions more likely. However, with an editorial board size of nearly 7000 individuals, it is hard to imagine this scenario being the rule rather than the exception. While it is possible that spurious correlations

¹⁰ Analysis of self-citation report per-citation rates with ranges between 20% to 40% of references according to a 2010 meso- and individual-level study (Costas et al., 2010), and an interquartile range of 8.6% to 17.7% in a individual-level 2019 study (Loannidis et al., 2019).

alone can give rise to significant regression model estimates, we rule out this possibility by using a "placebo" randomization scheme showing that the "repeat author" heuristic we use does not produce an entirely spurious configuration [Fig. 3(B,C)]. To summarize, PLOS ONE has an enormous impact on the production of scientific literature and the connectivity of the science citation network (Pan et al., 2018). Given the implicit constraints in monitoring and managing such a large and distributed organization, it is not unlikely that a small set of individuals might take advantage of the system for personal gain. Thus, our study highlights why large megajournals should record, monitor, and embrace transparency. As science continues to grow, these conflicts-of-interest may become more difficult to avoid, e.g. in large teams or a large journal, due to the challenges in monitoring individual activity and managing incentives in distributed operations (Petersen, Pavlidis, et al., 2014). As such, it is important to develop methods for identifying anomalous behavior, and to sanction these cases after thorough internal review. While manuscript editors certainly deserve credit for their service as scientific gatekeepers, it is important to raise the possibility that despite their laudable achievements meriting gatekeeper status, that editors too may develop ulterior motives directed at the "gamification" of a research evaluation metrics (Biagioli et al., 2019).

We conclude with some policy recommendations. We recommend that journals follow the lead of PLOS journals by publicly recording the specific editor overseeing the review of accepted articles, e.g. printing this information on the article cover page. This will facilitate the transparent evaluation of editors activities and can readily be justified on account of transparency, sanctioning, quality management, and responsible science. As gatekeepers to our knowledge base, manuscript editors have a pronounced responsibility to remain unbiased, despite a growing literature indicating otherwise (Card & DellaVigna, 2017; Helmer et al., 2017; Kravitz et al., 2010; Laband & Piette, 1994; Medoff, 2003; Sarigöl et al., 2017; Teplitskiy et al., 2018; van Lent et al., 2014). We also recommend that electronic-only megajournals that do not have volume restrictions should nevertheless place restrictions on the number of articles an editor can oversee at a time and per year. In addition to discouraging editors from taking advantage of their gatekeeper power, it would also encourage higher quality standards for accepting articles. And finally, megajournals should consider implementing additional levels of oversight in an effort to reduce temptations associated with perverse incentives. A good starting point may be the two-tiered editorial board system implemented by the journals *Proceedings of the National Academy of Sciences* and *Management Science*, in which a rotating body of managing editors oversees the board of article editors.

Acknowledgments

We are grateful to Andrew Johnston and three anonymous referees for helpful comments. **Data accessibility:** The data used in this study were obtained from the openly accessible websites of PLOS ONE, Proceedings of the National Academy of Sciences, Management Science, and Web of Science; parsed and merged data, and supplementary code are available at https://doi.org/10.6071/M39W8V.

Appendix A

A.1 Citation normalization to account for temporal and research subject area variation

PLOS ONE is a multi-disciplinary journal, accepting submissions from all research domains. The disciplinary diversity among accepted papers introduces a significant measurement challenge, because we seek to explain meaningful variations in citation impact for articles, net of publication year (indexed by *t*) and discipline-specific subject areas (indexed by *s*). To address the latter, we use the classification system defined and maintained by PLOS ONE to assign each article to one of six primary subject areas: (i) Biology and life sciences, (ii) Medicine and health sciences, (iii) Physical sciences, (iv) Social sciences, People and places, (v) Engineering and technology, Computer and information sciences, (vi) Ecology and environmental sciences, Earth sciences.

We must also standardize the citation impact measure to address three principal statistical biases: variation in publication rates across discipline, censoring bias and citation inflation. The first refers to the fact that larger disciplines, e.g. "Biology and life sciences", produce more publications, and hence, more citations than other disciplines such as "Earth sciences". The second bias reflects the fact that older publications have had more time to accrue citations than newer ones. And the third bias refers to the fact that more citations are produced over time as a product of increasing publication rates and reference list lengths, leading to a significant inflation in the relative value of citations. By way of example, a recent study demonstrated that the total number of references produced by all scientific articles is growing by 5.6% annually, and hence doubling every 12.4 years (Pan et al., 2018).

To address these three measurement problems, we map the raw citation count $c_{A,t}^{s}$ of a given article¹¹ to a normalized or "detrended" value

$$z_A^{\rm s} = \frac{\ln(1+c_{A,t}^{\rm s}) - \langle \ln(1+c_t^{\rm s}) \rangle}{\sigma[\ln(1+c_t^{\rm s})]}.$$

⁽³⁾

¹¹ The census date for tabulating WOS citation counts is Y = 2/25/2019.

The mean, $\langle \ln(1 + c_t^s) \rangle$, and the standard deviation, $\sigma[\ln(1 + c_t^s)]$, are calculated only over publications from the same year *t* and refined subject area *s*. The constant 1 is added to each citation count in order to avoid the divergence (ln 0) associated with uncited articles, and does not affect the results.

By analyzing the logarithm of the citation count, this normalization leverages the universal log-normal statistics of citation distributions (Radicchi, Fortunato, & Castellano, 2008). Moreover, by rescaling the logarithm by the standard deviation, the underlying inflationary bias has been removed, and so the distribution of P(z) is stationary, thereby permitting cross-temporal comparison. As such, z is particularly well-suited for regression analysis, as recently demonstrated in longitudinal analyses of cumulative advantage (Petersen & Penner, 2014) and collaboration (Petersen, 2015) within researcher careers. Fig. 7 demonstrates that the probability distributions P(z|s, t) are all approximately normally distributed, and thus sufficiently time invariant for the purposes of our analysis, for each subject area and year. Fig. 8 provides a correspondence chart relating z_A^s and c_t^s values to provide an estimate of effect sizes in our regression analyses. For example, in 2008 a publication with s = 0 (Biology and life sciences) and baseline value z = 0 (meaning that $\ln c_A$ is equal to the logarithmic mean value for publications from that year, $\ln c_{A,t} = \langle \ln c_t \rangle$) corresponds to 33 citations; for 2010 z = 0 corresponds to 24 citations; and for 2012 z = 0 corresponds to 16 citations.

A.2 Article subject area classification

It is well known that citation rates are affected by discipline-dependent factors. Indeed PLOS ONE is comprised of articles from a range of disciplines, and is classified by WOS as a "Multidisciplinary" journal. Thus, we were careful not to blindly pool the citation impact measures from all articles together. Instead, we methodically separated the articles into subsets, so that the relative citation difference between two articles is less biased by disciplinary and even sub-disciplinary factors, such as research community size and innovation level. As a result, we are able to more accurately estimate differences in citation impact, used here as a proxy for scientific impact.

We grouped the articles by subject area (SA) based on the internal PLOS ONE classification subject area classification system derived from a controlled thesaurus of nearly 8000 keywords. To be specific, we started with the keywords appearing on the webpage of each article *A*. Nearly all articles have 8 keywords per article, with only a handful of articles containing less than 8. PLOS ONE also has an article-classification scheme which is used to group articles for comparing article visibility.

While these keywords are helpful for classifying articles, they are not fully sufficient. Instead, PLOS ONE implements a 2-level classification system which is evident using the "page-views" applet on each article's "Metrics" page. By way of example, the article with DOI https://doi.org/10.1371/journal.pone.0000112 is classified primarily as "Biology and life sciences", with 3 subclassifications (Evolutionary Biology, Genetics, and Population Biology). At the core of this classification system are 10 top-level groups: ranked according to their empirical frequency, they are: (i) Biology and life sciences, (ii) Medicine and health sciences, (iii) Physical sciences, (iv) People and places, (v) Social sciences, (vi) Engineering and technology, (vii) Computer and information sciences, (viii) Ecology and environmental sciences, (ix) Earth sciences, (x) Science policy.

Thus, a fundamental problem is the fact that articles have multiple subclassifications, and so there is no 1-1 correspondence between a given article and a single top-level classification. A second problem is that not all articles have the classification data, despite the fact that all articles do have keywords. Thus, we developed an algorithmic method to classify articles into a small set of refined subject areas using only their keywords as classifier inputs. To be specific, we calculated the weighted bipartite network associating keywords and top-level classifications by aggregating the statistics for all publications with top-level classifications and keywords. In this way, we calculated a vector of 10 weights for each keyword corresponding to the 10 top-level classifications. Then, we identified the principal SA of each article by combining the weight vectors for each of the individual keywords, and choosing the SA with the largest weight. Take again the article DOI:https://doi.org/10.1371/journal.pone.0000112 with the 8 keywords "Chromosome 4", "Genetic loci", "Centromeres", "X chromosomes", "Population genetics", "Chromosomes", "Sex chromosomes", "Alleles". As one might expect, these article keywords give the largest weight to SAs (i)"Biology and life sciences" and (ii)"Medicine and health sciences".

Applying this method to all articles, we found that the biomedical classifications (i) and (ii) are the most common first and second ranked classifications. Fig. 7(A) shows the SA count histogram for all PLOS ONE articles, with 123,750 (87.1% of all articles) having "Biology and life sciences" as the principal classification, and none having "Science policy" as the principal classification. Contrariwise, only 17 articles had "Earth sciences" as the principal classification. To account for the fact that the majority of the keywords in the PLOS ONE thesaurus are related to (i) and (ii), leading to the disparity in the principal classification, we created an exception rule in order to better account for the second-ranked classification. First, if the principal classification was (i), then we instead used the second-ranked classification as the principle classification. This rule helped to classify more publications for SA (iii)–(ix), as demonstrated by the second count histogram shown in Fig. 7(B). As one final step to condense the SA classifications, we joined the groups (iv) and (v), (vi) and (vii), and (vii) and (ix), since intuitively, there is considerable overlap between these SAs. Thus, Fig. 7(C) shows the final refined distribution of articles across the 6 *refined SA* used in our analysis: the smallest refined SA is 6/7 with 533 articles and the second-smallest is 4/5 with 1839 articles over 2006–2015; the remaining refined SA are comprised of 8000 or more articles over the 10-year period. In what follows, we define the index variable s=1...6 to denote the refined SA of an individual article. In the following section, *s* is used to indicate the SA-specific mean and standard deviations used to define the normalized citations. And in our regression models, we use *s* as an indicator variable to control for variable citation rates and acceptance times across SA.

A.3 Name disambiguation problem among editors and authors

Due to the name disambiguation problem – i.e. it is difficult to distinguish common last name and first name initial combinations in WOS data – there are certain abbreviated name combinations that we ignored in aspects of our analysis. First, in order to determine if an article was coauthored by a PLOS ONE editor, there were certain editor names which were too similar in their abbreviated forms, e.g. Shree Singh and Seema Singh, who both occur in WOS records as "Singh S". Thus, for those editor name abbreviations which have a multiplicity of 2 or greater, we do not count articles with these abbreviated names as being coauthored by an editor.

Second, this name disambiguation problem occurs in the identification of the top authors within the article set of each editor. Thus, using the editor name set as our baseline, we also ignored all surnames – independent of first name initial – for those PLOS ONE editors with common last names. As such, the list of 160 common surnames ignored in the coauthor analysis are: Singh, Isalan, Hoheisel, Lo, Castresana, Liu, Zheng, Yang, Deb, Qiu, Chang, Zhou, Bhattacharya, Tang, Lee, Xu, Li, Cheng, Wang, Scott, Yu, Tan, Miao, Williams, Klymkowsky, Kaltenboeck, Zhang, Chen, He, Song, Brown, Lin, Brody, Wei, Kumar, Yan, Shi, Carvalho, Rogers, Ng, Ray, Phillips, Soriano-Mas, Paul, Fox, Butler, Ma, Wu, Carter, Xie, Hector, Wright, Caldwell, Fang, Sorensen, Lam, Chan, Stewart, Huang, Gravenor, Pan, Gupta, Smith, Lu, Cao, Xia, Ho, Moore, Liang, Franco, Parida, Zhao, Wilson, Gilbert, Nigou, Redfield, Paci, Park, Sun, Zhu, Chalmers, Clark, Colombo, Zuo, Das, Tian, Moreno, Meng, Gray, Schweisguth, Lopez-Garcia, Yue, Johnson, Wong, Medina, Fung, Kato, Roberts, Hwang, Hsieh, Wen, Knight, Csernoch, Anderson, Grant, Clarke, Jiang, Jones, Rao, Feng, Nguyen, Choi, Thomas, Chiu, Samuel, Gordon, Heutink, Evans, Martin, Ren, Berger, Kim, Han, Mao, White, McCutcheon, Temussi, Taylor, Schmitt, Kerby, Miller, Roy, Pereira, Shankar, Aoki, Jackson, Adams, Russell, Thompson, Abe, Duan, Hong, Borras, Costa, Yam, Porollo, Stumbles, Agarwal, Beier, Xiao, Beaudoin, Nosten, Shen, Feldman, Hall, Raible, Yin, Kelly, Simos, Knudsen.

References

- Batagelj, V., Ferligoj, A., & Squazzoni, F. (2017). The emergence of a field: A network analysis of research on peer review. *Scientometrics*, 113, 503–532. Becker, G. S. (1957). *The economics of discrimination*. Chicago: University of Chicago Press.
- Biagioli, M., Kenney, M., Martin, B. R., & Walsh, J. P. (2019). Academic misconduct, misrepresentation and gaming: A reassessment. Research Policy, 48, 401-413.
- Binfield, P. (2013). Open access megajournals Have they changed everything?
- Bjork, B.-C. (2015). Have the 'mega-journals' reached the limits to growth? PeerJ, 3, e981.
- Bollen, J., Crandall, D., Junk, D., Ding, Y., & Börner, K. (2016). An efficient system to fund science: From proposal review to peer-to-peer distributions. Scientometrics, 1–8.
- Börner, K., Edmonds, B., Milojević, S., & Scharnhorst, A. (2016). Editorial. Scientometrics, 1-4.
- Bornmann, L., & Daniel, H.-D. (2008). What do citation counts measure? A review of studies on citing behavior. Journal of Documentation, 64, 45–80.
 Bravo, G., Farjam, M., Moreno, F. G., Birukou, A., & Squazzoni, F. (2018). Hidden connections: Network effects on editorial decisions in four computer science journals. Journal of Informetrics, 12, 101–112.
- Bravo, G., Grimaldo, F., López-Iñesta, E., Mehmani, B., & Squazzoni, F. (2019). The effect of publishing peer review reports on referee behavior in five scholarly journals. *Nature Communications*, 10, 322.
- Card, D., & DellaVigna, S. (2017). What do editors maximize? Evidence from four leading economics journals. National Bureau of Economic Research.

Technical Report.

Colussi, T. (2018). Social ties in academia: A friend is a treasure. Review of Economics and Statistics, 100, 45-50.

- Costas, R., van Leeuwen, T. N., & Bordons, M. (2010). Self-citations at the meso and individual levels: Effects of different calculation methods. Scientometrics, 82, 517–537.
- De Solla Price, D. (1976). A general theory of bibliometric and other cumulative advantage processes. Journal of the American Society for Information Science, 27, 292–306.
- Editorial. (2017). Fake news threatens a climate literate world. Nature Communications, 8, 15460.
- Editorial Policy Committee. (2012). CSE's white paper on promoting integrity in scientific journal publications, 2012 update.
- Evans, J. A., & Foster, J. G. (2011). Metaknowledge. Science, 331, 721-725.
- Fealing, K. H. (Ed.). (2011). The science of science policy: A handbook. Stanford, CA, USA: Stanford Business Books.
- Fong, E. A., & Wilhite, A. W. (2017). Authorship and citation manipulation in academic research. PLOS ONE, 12, e0187394.
- Fortunato, S., Bergsrom, C. T., Borner, K., Evans, J. A., Helbing, D., Milojevic, S., et al. (2018). Science of science. Science, 359, eaao0185.

Fowler, J. H., & Aksnes, D. W. (2007). Does self-citation pay? Scientometrics, 72, 427-437.

- Gino, F., & Bazerman, M. H. (2009). When misconduct goes unnoticed: The acceptability of gradual erosion in others' unethical behavior. Journal of Experimental Social Psychology, 45, 708–719.
- Gino, F., Schweitzer, M. E., Mead, N. L., & Ariely, D. (2011). Unable to resist temptation: How self-control depletion promotes unethical behavior. Organizational Behavior and Human Decision Processes, 115, 191–203.
- Hellsten, I., Lambiotte, R., Scharnhorst, A., & Ausloos, M. (2007). Self-citations, co-authorships and keywords: A new approach to scientists' field mobility? Scientometrics, 72, 469–486.
- Helmer, M., Schottdorf, M., Neef, A., & Battaglia, D. (2017). Gender bias in scholarly peer review. eLife, 6, e21718.
- Ioannidis, J. P. A., Baas, J., Klavans, R., & Boyack, K. W. (2019). A standardized citation metrics author database annotated for scientific field. *PLOS Biology*, 17, 1–6.
- Kravitz, R. L., Franks, P., Feldman, M. D., Gerrity, M., Byrne, C., & Tierney, W. M. (2010). Editorial peer reviewers' recommendations at a general medical journal: Are they reliable and do editors care? PLOS ONE, 5, 1–5.
- Laband, D. N., & Piette, M. J. (1994). Favoritism versus search for good papers: Empirical evidence regarding the behavior of journal editors. *Journal of Political Economy*, 102, 194–203.
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabasi, A.-L., Brewer, D., et al. (2009). Computational social science. Science, 323, 721–723.

Lee, C. J., Sugimoto, C. R., Zhang, G., & Cronin, B. (2013). Bias in peer review. Journal of the Association for Information Science and Technology, 64, 2–17.

van Lent, M., Overbeke, J., & Out, H. J. (2014). Role of editorial and peer review processes in publication bias: Analysis of drug trials submitted to eight medical journals. *PLOS ONE*, 9, 1–8.

Malhotra, D., & Gino, F. (2011). The pursuit of power corrupts. Administrative Science Quarterly, 56, 559-592.

- Medoff, M. H. (2003). Editorial favoritism in economics? Southern Economic Journal, 425–434.
- Oreskes, N., & Conway, E. M. (2011). Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming. USA: Bloomsbury Publishing.
- Pan, R. K., Petersen, A. M., Pammolli, F., & Fortunato, S. (2018). The memory of science: Inflation, myopia, and the knowledge network. *Journal of Informetrics*, 12, 656–678.
- Petersen, A. M. (2015). Quantifying the impact of weak, strong, and super ties in scientific careers. Proceedings of the National Academy of Sciences of the United States of America, 112, E4671–E4680.
- Petersen, A. M., Fortunato, S., Pan, R. K., Kaski, K., Penner, O., Rungi, A., et al. (2014). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences of the United States of America*, 111, 15316–15321.
- Petersen, A. M., Pan, R. K., Pammolli, F., & Fortunato, S. (2019). Methods to account for citation inflation in research evaluation. Research Policy, 48, 1855–1865.
- Petersen, A. M., Pavlidis, I., & Semendeferi, I. (2014). A quantitative perspective on ethics in large team science. Science and Engineering Ethics, 20, 923–945.
- Petersen, A. M., & Penner, O. (2014). Inequality and cumulative advantage in science careers: A case study of high-impact journals. *EPJ Data Science*, 3, 24. Petersen, A. M., Riccaboni, M., Stanley, H. E., & Pammolli, F. (2012). Persistence and uncertainty in the academic career. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 5213–5218.
- Petersen, A. M., Vincent, E. M., & Westerling, A. L. (2019). Discrepancy in scientific authority and media visibility of climate change scientists and contrarians. *Nature Communications*, 10, 3502.
- Powell, K. (2016). The waiting game: Does it take too long to publish research? Nature, 530, 148-151.
- Radicchi, F., Fortunato, S., & Castellano, C. (2008). Universality of citation distributions: Toward an objective measure of scientific impact. Proceedings of the National Academy of Sciences of the United States of America, 105, 17268–17272.
- Reingewertz, Y., & Lutmar, C. (2018). Academic in-group bias: An empirical examination of the link between author and journal affiliation. Journal of Informetrics, 12, 74–86.
- Rzhetsky, A., Foster, J. G., Foster, I. T., & Evans, J. A. (2015). Choosing experiments to accelerate collective discovery. Proceedings of the National Academy of Sciences of the United States of America, 112, 14569–14574.
- Sarigöl, E., Garcia, D., Scholtes, I., & Schweitzer, F. (2017). Quantifying the effect of editor-author relations on manuscript handling times. *Scientometrics*, 1–23.
- Sarigöl, E., Pfitzner, R., Scholtes, I., Garas, A., & Schweitzer, F. (2014). Predicting scientific success based on coauthorship networks. *EPJ Data Science*, 3, 9. Scharnhorst, A., Börner, K., & Besselaar, P. v. d. (Eds.). (2012). *Models of science dynamics: Encounters between complexity theory and information sciences.*
- Heidelberg: Springer.
- Seeber, M., Cattaneo, M., Meoli, M., & Malighetti, P. (2019). Self-citations as strategic response to the use of metrics for career decisions. *Research Policy*, 48, 478–491.
- Solomon, D. J. (2014). A survey of authors publishing in four megajournals. PeerJ, 2, e365.
- Solomon, D. J., & Bjork, B.-C. (2012). A study of open access journals using article processing charges. Journal of the American Society for Information Science and Technology, 63, 1485–1495.
- Stephan, P. E. (2012). How economics shapes science. Cambridge, MA: Harvard University Press.
- Sugimoto, C. R., Lariviere, V., Ni, C., & Cronin, B. (2013). Journal acceptance rates: A cross-disciplinary analysis of variability and relationships with journal measures. *Journal of Informetrics*, 7, 897–906.
- Tahamtan, I., Afshar, A. S., & Ahamdzadeh, K. (2016). Factors affecting number of citations: A comprehensive review of the literature. *Scientometrics*, 107, 1195–1225.
- Teplitskiy, M., Acuna, D., Elamrani-Raoult, A., Körding, K., & Evans, J. (2018). The sociology of scientific validity: How professional networks shape judgement in peer review. *Research Policy*, 47, 1825–1841.
- Vieira, E. S., & Gomes, J. A. (2010). Citations to scientific articles: Its distribution and dependence on the article features. *Journal of Informetrics*, 4, 1–13. Wakeling, S., Willett, P., Creaser, C., Fry, J., Pinfield, S., & Spezi, V. (2016). Open-access mega-journals: A bibliometric profile. *PLOS ONE*, 11, e0165359. Wilhite, A. W., & Fong, E. A. (2012). Coercive citation in academic publishing. *Science*, 335, 542–543.
- Wilsdon, J., Allen, L., Belfiore, E., Campbell, P., Curry, S., Hill, S., et al. (2015). The metric tide: Report of the independent review of the role of metrics in research assessment and management. Higher Education Funding Council for England (HEFCE). Technical Report.
- Yoon, A. H. (2013). Editorial bias in legal academia. *Journal of Legal Analysis*, 5, 309–338.
- Zaggl, M. A. (2017). Manipulation of explicit reputation in innovation and knowledge exchange communities: The example of referencing in science. *Research Policy*, 46, 970–983.