Author Accepted Manuscript

1 Citation counts and journal impact factors do not capture some indicators of research
2 quality in the behavioral and brain sciences

3 Michael R. Dougherty

4 Department of Psychology, University of Maryland

5 Zachary Horne

6 Department of Psychology, University of Edinburgh

"In press" at **Royal Society Open Science**

Abstract

Citation data and journal impact factors are important components of faculty dossiers and figure prominently in both promotion decisions and assessments of a researcher's broader societal impact. Although these metrics play a large role in high-stakes decisions, the evidence is mixed about whether they are strongly correlated with indicators of research quality. We use data from a large scale dataset comprising 45,144 journal articles with 667,208 statistical tests and data from 190 replication attempts to assess whether citation counts and impact factors predict three indicators of research quality: (1) the accuracy of statistical reporting, (2) the evidential value of the reported data, and (3) the replicability of a given experimental result. Both citation counts and impact factors were weak and inconsistent predictors of research quality, so defined, and sometimes negatively related to quality. Our findings raise the possibility that citation data and impact factors may be of limited utility in evaluating scientists and their research. We discuss the implications of these findings in light of current incentive structures and discuss alternative approaches to evaluating research.

Citation counts and journal impact factors do not capture some indicators of research quality in the behavioral and brain sciences

Introduction

Researchers and university administrators often assume that journal impact factors and citation counts are indicators of research quality (e.g., [McKiernan et al., 2019](#); [Sternberg, 2016](#)). This assumption seems plausible: High impact journals may seem to have a more selective and rigorous review process, thereby weeding out lower quality research as a consequence. One might also view citation counts as reflecting something akin to the wisdom of the crowd whereby high-quality research garners more citations than low-quality research. One need not look far to see these assumptions on display: Universities often promote bibliometric indices such as citation counts and journal impact factors as indices of “impact” or “quality”, academic departments use these metrics for important decisions such as hiring, tenure, and promotion, and science journalists promote research from high-impact journals. It is also common for authors to equate impact factors and citation counts with quality ([Ruscio, 2016](#); [Sternberg, 2016](#)) – an assumption that appears in university promotion and tenure policies ([McKiernan et al., 2019](#)). For instance, in their review of promotion policy documents, [McKiernan et al. \(2019\)](#) found that the vast majority of institutions that include impact factors describe them as positive indicators of quality, significance, importance, or impact. The inclusion of these metrics in high-stakes decisions starts with the assumption that there is a positive and meaningful relation between the quality of ones work on the one hand, and impact factors and citations on the other. This raises the question, are we justified in thinking that high-impact journals or highly-cited papers are of higher quality?

Before proceeding, it is important to note that citation counts and journal impact factors are often treated as variables to be maximized, under the assumption that high citation counts and publishing in high-impact journals demonstrate that ones work is of high quality. This view implicitly places these variables on the left-hand side of the

prediction equation, as if the goal of research evaluation is to predict (and promote) individuals who are more likely to garner high citation counts and publish in high-impact factor journals. This view, whether implicitly or explicitly endorsed is problematic for a variety of reasons. First, it neglects the fact that citation counts themselves are determined by a host of factors that may be unrelated to research quality, or for that matter even unrelated to the science being evaluated (Aksnes, Langfeldt, & Wouters, 2019; Bornmann & Daniel, 2008; Larivière & Gingras, 2010). For instance, citation counts covary with factors such as the length of the title and the presence of colons or hyphens in the title (Paiva, Lima, & Paiva, 2012; Zhou, Tse, & Witheridge, 2019) and are correlated with other non-scientific variables, such as the size of ones social network (Mählck & Persson, 2000), the use of social media, and the posting of preprints on searchable archives (Gargouri et al., 2010; Gentil-Beccot, Mele, & Brooks, 2009; S. Lawrence, 2001; Luarn & Chiu, 2016; Peoples, Midway, Sackett, Lynch, & Cooney, 2016). Citation counts also tend to be higher for papers on so-called “hot” topics and for researchers with greater citation-based reputations (Petersen et al., 2014). Not only are researchers incentivized to maximize these, but it is also easy to find advice on how to game them (Nature, 2010, 2019; Sage, 2022).

Second, treating these variables as values to be maximized invites problematic inferences such as inferring that mentorship quality varies by gender simply because students of male mentors tend to enjoy a citation advantage (see the now retracted paper by AlShebli, Makovi, & Rahwan, 2020), and even perpetrate systemic inequalities in career advancement and mobility due to biases in citation patterns that disfavor women and persons from underrepresented groups (Greenwald & Schuh, 1994; X. Wang et al., 2020b). Finally, treating citations and impact factors as the to-be-maximized variables may alter researchers’ behaviors in ways that can undermine science (Chapman et al., 2019). For example, incentivizing researchers to maximize citations may lead researchers to focus on topics that are in vogue regardless of whether doing so addresses key questions that will advance their field.

If we instead think of citation counts and impact factors as predictors of influential and quality science, we can ask whether they are valid proxies for assessing key aspects of the quality of a researcher. To be clear, these metrics are not and cannot be considered *direct* measures of research quality. But of course, addressing this question requires a way of measuring quality that is independent of citation counts. Past work on this topic has primarily addressed the issue by relying on subjective assessments of quality provided by experts or peer reviews. On balance, these studies have shown either weak or inconsistent relationships between quality (broadly defined by the authors) and citation counts (e.g., Nieminen, Carpenter, Rucker, & Schumacher, 2006; Patterson & Harris, 2009; West & McIlwaine, 2002). One challenge in relying on subjective assessments is that their use assumes that the judges can reliably assess quality—an assumption that has been challenged by Bornmann, Mutz, and Daniel (2010), who showed that the inter-rater reliability of peer review ratings is extremely poor. Indeed, controlled studies of consistency across reviewers also indicate a surprisingly high level of arbitrariness in the review process (see Francois, 2015; Langford & Guzdial, 2015; Peters & Ceci, 1982). Peters and Ceci (1982), for instance, resubmitted 12 articles (after changing the author names) that had previously been accepted for publication in psychology journals and found that the majority of the articles that had been accepted initially were rejected the second time around based on serious methodological errors. Similarly, N. Lawrence and Cortes (2014) reported a high degree of arbitrariness in accepted papers submitted to the *Neural Information Processing Systems* annual conference, widely viewed as one the premier peer reviewed outlets in Machine Learning. In this study, a random sample of submissions went through two independent review panels; the authors estimated that 60% of decisions appeared to arise from an arbitrary process. By comparison, a purely random process would have yielded an arbitrariness coefficient of 78%, whereas a process without an arbitrary component would have yielded a coefficient of 0%. If reviewers who are specifically tasked with judging quality cannot agree on the acceptability of research for

publication, then it is unlikely that impact factors or citation counts, which themselves are dependent on the peer review process, are reliable indicators of quality.

Several issues have been raised with the use of bibliometrics for faculty evaluation and the incentive structure for appointments, tenure, and promotion. First, there is growing concern about the *improper* use of bibliometrics when evaluating faculty. This concern has been expressed by several scholars ([Aksnes et al., 2019](#); [Hicks, Wouters, Waltman, de Rijcke, & Rafols, 2015](#); [Moher et al., 2018](#); [Seglen, 1997](#)) and has been codified in the San Francisco Declaration on Research Assessment (DORA) – a statement that has been endorsed by hundreds of organizations, including professional societies such as the Association for Psychological Sciences. Second, several recent analyses challenge traditional orthodoxy that bibliometrics are valid indicators of research quality. For instance, [Fraley and Vazire \(2014\)](#), and [Szucs and Ioannidis \(2017\)](#) report a moderate *negative* correlation (approximately -0.42) between impact factor and statistical power in the domains of social psychology and cognitive neuroscience, respectively (see also [Brembs, 2018](#); [Brembs, Button, & Munafò, 2013](#)). Despite these results and repeated calls for scaling back the use of bibliometrics in research evaluation, their use is still widespread in evaluating faculty job candidates and promoting them on the basis of their *research trajectory* ([McKiernan et al., 2019](#)).

The use of impact factors and citation counts as indices of a scholar's future performance also highlights that these metrics are not only used retrospectively, they are also intended to be diagnostic of how scholars will perform in the future. One of the principle functions of the hiring and tenure processes is prediction: The goal is to predict (and select) scientists with the most promise for advancing science. Realizing this goal requires both that we have access to valid predictors of future scientific accomplishments and that we use those indicators appropriately. To our knowledge, there is no evidence for either of these claims ([Chapman et al., 2019](#)). As discussed, there are several reasons to think that these metrics do not *predict* quality ([Aksnes et al., 2019](#); [Hicks et al., 2015](#);

130 [Moher et al., 2018](#); [Seglen, 1997](#)). In order for hiring and tenuring committees to establish
131 their usefulness would require formally modeling how these factors accounted for a scholars
132 future performance. Moreover, even if hiring and tenuring committees actually built such a
133 model, they would need to calibrate their decision-making in accordance with the model.
134 We suspect, again, that this is not done, nor would it be.

135 According to the view that impact factors and citations are useful for assessing
136 research quality, we would expect reasonable, or even strong, positive statistical
137 relationships with objective indices of research quality. Evidence to the contrary or even
138 evidence of a lack of relationship would suggest that their utility has been overestimated.
139 Existing research already suggests that journal impact factors and citation counts are given
140 weight beyond how well they could plausibly predict future performance, potentially at the
141 expense of other diagnostic measures. Although quality is a complex construct, there are
142 several factors that we believe are indicative of research quality (though, without question,
143 these indicators are themselves only proxies). These include methodological features of an
144 article such as use of random assignment (when feasible) and the use of appropriate control
145 groups, the representativeness of the participant population, the appropriate reporting and
146 transparency of statistical analyses, the accuracy with which results are reported, the
147 statistical power of the study, the evidentiary value of the data (i.e., the degree to which
148 the data provide compelling support for a particular hypothesis, including the null
149 hypothesis), and the replicability of research findings. Many indicators of research quality
150 can be assessed only with close inspection by domain experts (e.g., [Nieminen et al., 2006](#);
151 [Patterson & Harris, 2009](#); [West & McIlwaine, 2002](#)), but some others can be identified
152 using automated data mining tools.

153 In this paper, we take the latter approach by examining how citation counts and
154 journal impact factors relate to three aspects of research quality: (1) accuracy of reporting
155 statistical conclusions as measured by the number of errors in the reporting of statistics in a
156 paper, (2) the strength of statistical evidence provided by the data, as defined by the Bayes

factor (Edwards, Lindman, & Savage, 1963), and (3) the replication of empirical findings. Inasmuch as impact factors and citation counts are measures of research quality, we would expect that papers in higher impact journals and highly cited papers should “pass” some key diagnostic checks; for instance, some indicators of problems or virtues in a paper may be the presence (or absence) of reporting errors, the strength of evidence provided by their data, and the independent replicability of the findings. For citation counts and impact factors to be useful for research evaluation, at least as they are currently used, we expect that the associations should be reasonably large. The advantages of our approach are that our chosen indices are quantifiable, verifiable, transparent, and not reliant on subjective evaluations of reviewers who may not agree about the quality of a paper.

The analyses reported herein bring together three data sources to address whether citation counts and journal impact factors have as much utility for assessing research quality as how they are used in practice. These data sources include a data set used in a large-scale assessment of statistical reporting inconsistencies across 50,845 articles in behavioral science and neuroscience journals published between 1985 and 2016 (Hartgerink, 2016) and 190 replication studies from a collection of repositories (see Reiner et al., 2020). Each data set was merged with the 2017 journal impact factors and/or article-level citation counts. The raw data and code used to generate our analyses and figures are available on the Open Science Framework: <https://osf.io/hngwx/>.

There are benefits and limitations that must be acknowledged when using this kind of “big data” approach. Large data sets afford the possibility of detecting even small effects. We expect that citation counts and impact factors are over-weighted, but if these metrics are even relatively weak predictors, that might provide some reason for optimism about their use. However, there is one key limitation that must be acknowledged at the outset. To some extent, we sacrifice some fidelity in distinguishing between statistical tests that reflect more “central hypotheses” tested by researchers, versus hypotheses that reflect more “peripheral hypotheses.”

It may be tempting to assume that this limitation of our research indicates that—at the outset—our results may be quite limited because they do not focus on a paper’s central hypotheses. However, a paper is not encapsulated by a single or even a few hypothesis tests. Rather, a paper is the conjunction of the tests authors decided *ought* to be reported to provide the necessary details for the reader to understand the phenomena of interest; authors do not report tests at random. So, even if a given test does not represent the most central prediction of a study, the test reported should nonetheless be *accurate* and have *evidential value*.¹

Analytic strategy

Throughout, we perform a combination of Bayesian multilevel modeling to estimate effects and Bayes factors to quantify support for the alternative hypothesis, while including an index of out-of-sample predictive accuracy to aid in assessing the utility of each predictor. Bayesian analyses afford some advantages for fitting models with complex random-effects structures. For all Bayesian models, we set *weakly regularizing* priors in `brms` (Bürkner, 2017), which are detailed in our analysis scripts on the Open Science Framework. Crudely put, weakly regularizing priors guide the MCMC estimation process but, particularly when a data set is large, allow the posterior to primarily reflect the data. Posterior predictive checks were conducted for each model to verify model fit. Posterior predictive checking involves simulating the data from the fitted model and graphically comparing the simulations to the observed distribution. This step is useful for diagnosing mis-specifications of the statistical model and ensuring that the model adequately reproduces the distribution. Code for conducting the posterior predictive checks is included in the analysis scripts provided on Open Science Framework.

We had specific questions we sought to answer about the relationship between

¹ A further point worth considering is that the distinction between central and peripheral tests may not be so clear. For instance, an exploratory analysis indicating a manipulation check behaved as predicted might initially regarded as ‘peripheral,’ but if the check did not behave as predicted, this test could become more central because it would become an observation the author would need to account for.

citation counts, impact factors, and measures of research quality, but our analyses were not
 preregistered and are nonetheless exploratory. Consequently, we checked the robustness of
 our inferences by changing model parameterizations and prior choices, including or
 excluding different covariates and group-level effects, and making alternative distributional
 assumptions (e.g., comparing model fits with Gaussian vs. Beta distributions).
 Throughout the paper, we report the mean posterior parameter estimates and the 95%
 credible intervals, but our focus is on characterizing the magnitude of the effects, rather
 than on binary decisions regarding statistical significance (or not). To assess the usefulness
 of the various covariates for prediction, we also computed the Bayesian analogue of the R^2
 statistic based on leave-one-out (LOO) cross validation (Gelman, Goodrich, Gabry, &
 Vehtari, 2019). The LOO adjusted R^2 provides a better estimate of out-of-sample
 prediction accuracy compared to the Bayesian R^2 . Finally, we also estimated the Bayes
 factors (BF) for all single-predictor mixed effects models and ordinal analyses. The BF
 expresses the odds in favor of one model compared to an alternative. Here, we use BF_{10} as
 an index of the alternative hypothesis relative to an appropriately formed null hypothesis.
 Unless otherwise noted, we use a point-null hypothesis of no effect, with the prior centered
 on zero. BF's for mixed effects models used Bayes factor functions provided in the `brms`
 and `bayesrestR` (Makowski, Ben-Shachar, & Lüdtke, 2019) packages, whereas the BF's
 for Kendall's tau and Mann-Whitney U statistics used methods developed in van Doorn,
 Ly, Marsman, and Wagenmakers (2018) and the Bayesian t-tests used the `BayesFactor`
 package in R (Morey & Rouder, 2018). All Bayesian t-tests used the default Cauchy prior
 with a width parameter of 0.707. This prior places roughly 50% of the mass on effects
 between -.707 and .707.² A BF_{10} greater than one reflect the degree to which the data
 support the alternative hypothesis, while values less than 1 reflect cases in which the data

² Although impact factors and citations counts have traditionally been assumed to reflect quality, we
 anticipated that any observed relationships were likely to be modest at best. The default prior captures
 this prior belief by weighing small to modest effects more heavily than large effect. Aside from this general
 justification, we had no other information to guide the setting of the prior.

support the null hypothesis. As a rule of thumb, BF_{10} within the range of 1/3 and 3/1 is interpreted as uninformative, with more extreme values interpreted as providing greater evidence for one hypothesis relative to the other. Although prior authors provide categorical distinctions for BF's of different magnitudes (Jeffreys, 1961), we prefer to interpret BFs as a continuous index of degree of support. We note here that many of the BF's we report are based on large samples, which can result in extremely large values even for relatively small effects. Thus, readers should interpret the magnitude of the BF's within the context of the reported R^2 's, along with the graphical analyses of the data.

Question 1: Do article citation counts and journal impact factors predict more accurate statistical reporting?

Hartgerink (2016) used `statcheck` (Nuijten, Hartgerink, van Assen, Epskamp, & Wicherts, 2016), an automated data mining tool, to identify statistical reporting inconsistencies among 688,112 statistical tests appearing in 50,845 journal articles in content areas related to the behavioral and brain sciences. Articles included in the data set were published between 1985 and 2016. A full description of the data and initial analyses are provided by Nuijten and colleagues (2016), but we detail critical aspects of the `statcheck` data set which suggest it is a valid measure of our research question. `statcheck` is an R package that uses regular expressions to extract American Psychological Association (APA) formatted results in a Null Hypothesis Significance Testing framework (Epskamp & Nuijten, 2014). Among other things, `statcheck` recalculates the p -values within a paper based on the test statistics and degrees of freedom reported with each test. `statcheck` has been shown to be a reliable indicator of inconsistencies in APA formatted papers. For instance, Nuijten and colleagues 2016 calculated the interrater reliability between `statcheck` and hand coding finding that `statcheck` and handcoders had an interrater reliability of .76 for inconsistencies between p -values and reported test statistics and .89 for the “gross inconsistencies.” We should note that although quite reliable,

naturalistic datasets—including the data set used here—will often entail minor limitations that can affect data quality. These data quality issues can arise for several reasons including errors in converting PDF files to text files, or more systematic errors due to some subfields reporting their analyses in ways which deviate from the American Psychological Association Style guide. To consider one example, in reviewing the `statcheck` data, we observed that many χ^2 statistics appeared to be miscoded. In particular, we observed that `statcheck` sometimes mis-identified sample size as the degrees of freedom for the χ^2 statistic. This leads to an incorrect computation of the p -value. To address this issue, we ran analyses both including and excluding all χ^2 statistics. The reported conclusions are based on analyses that included the χ^2 statistics, but we note here that analyses with the χ^2 statistics converged with those without. The only analyses for which the mis-coded χ^2 are included are those reporting decision error rates.

These issues notwithstanding, because of the established reliability of the `statcheck` data set, we merged this open-source data set with article-level citation counts and the impact factors from 2017 to understand the relationship between reporting errors and citation counts and journal impact factors. Citation counts were obtained using the `rcrossref` (Chamberlain, Zhu, Jahn, Boettiger, & Ram, 2019) package in R and the impact factors were obtained using the `scholar` (Keirstead, 2016) package in R. Although our dataset includes articles from 1985 to 2016, we used only the 2017 journal impact factors. We chose to use only impact factors from a single year for two reasons. First, review committees evaluating faculty for promotion, tenure, and awards generally are not provided JIFs for the year of publication, but rather are provided with the JIF in the most recent year. Given that our ultimate goal is to identify statistical relationships as they might manifest in professional contexts, this approach seems most defensible. However, we recognize that others addressing this problem may make other analytic decisions. Second, and more practically, obtaining JIFs for the full set of articles was not possible, given that existing databases only provided JIFs dating back to 1997. Limiting our analyses to papers

published since 1997 would have reduced our dataset by roughly 20%.

Data from each article were then summarized to reflect the number of statistical reporting errors per article and the total number of statistical tests. Articles with fewer than two statistical tests were excluded from our analyses. The motivation for this decisions was again twofold: First, it is rare for an empirical paper to report only a single statistical test. We reasoned that these papers likely reflected commentaries, notes, or editorials and were therefore not empirical contributions. Second, articles with fewer than two statistical tests do not allow for computing variability across tests. The final data set used for our analysis included 45,144 articles and 667,208 statistical tests.³

To address the question of whether article citation counts and journal impact factors predict more accurate statistical reporting, we examined the degree to which citation counts and impact factors predict diagnostics of quality in three ways. First, we examined whether articles containing at least one error were cited more or less than articles with no errors.⁴ Of the 45,144 articles included in the analysis, roughly 12.6% percent ($N = 5,710$) included at least one statistical reporting inconsistency that affected the decision of whether an effect was statistically significant or not. The majority of the decision errors ($N = 8,515$) included in these articles involved authors reporting a result as significant ($p \leq .05$), when the recalculated p -value based on the reported test-statistic (e.g., $t = 1.90$) and df (e.g., $df = 178$) was greater than .05. Only $N = 1,589$ consisted of the reverse error, in which a p -value was reported as non-significant or $p > .05$ when the test statistic indicated otherwise. This asymmetry suggests that these reporting errors are unlikely to be random. Indeed, of the reported tests in which the recomputed p -value was *greater* than .05 ($N = 178,978$), 4.76% ($N = 8,515$) were incorrectly reported as significant, whereas only 0.32%

³ Although this dataset is quite comprehensive, we nonetheless report inferential statistics throughout the manuscript. This is because even though we have the population of articles for specific years and journals, the statcheck dataset does not reflect the entire population of articles in, for example, cognition, social psychology, or clinical psychology. Rather, this dataset is specifically from large publishers like the American Psychological Association, Springer, and Taylor & Francis.

⁴ Analysis of non-decision errors are provided in the supplemental material. Conclusions based on these analyses are generally consistent with those based on the decision errors

($N = 1,589$) of the recomputed p-values that were significant ($N = 488,154$) were incorrectly reported as *non-significant* – a difference that was statistically reliable with strong evidence for the alternative hypothesis ($BF_{10} > 1.0e+16$ with a proportion test, $\text{beta}(1,1)$). This result directly replicates prior work by Nuijten and colleagues (2016). More interesting, articles containing at least one reporting error ($M = 52.1$, $SD=111.7$) garnered more citations than articles that did not contain any errors ($M = 46.8$, $SD=98.3$). Although the magnitude of the effect is small—indicating that papers with at least one decision error have on average 5 *more* citations than those without errors—the Bayes factor provides strong evidence for the alternative hypothesis, $BF_{10}=18.1$ (JZS two-sample Bayes factor, $r_{scale}=.707$, (Morey & Rouder, 2018)). This conclusion was robust to transformation: Correcting for skew in citation counts with the log transformation yielded even more convincing evidence that papers with errors are cited more frequently ($BF_{10}=2.4e+8$). Thus, at least by this criterion, and in contrast to the conventional assumption that citations are positively related to quality, this analysis indicates that citations are actually *inversely* related to quality.⁵ Although the above analysis of citation counts neglects the multilevel structure and complexity of the data, we share it because it is likely that people use citation counts in a similar way when inferring research quality, following the general principle of “more is better” without adjusting appropriately for differences across sub-disciplines or accounting for other relevant factors.

Next, we examined the relationship between frequency of errors and impact factor and citations counts using a Bayesian multilevel zero-inflated Poisson model. This model predicted the frequency of errors per paper on the basis of the log of a journal’s impact factor and the log of an article’s citations (+1), controlling for the year of publication and

⁵ One question is whether researchers conceive of citation counts as an index of research quality in the first place rather than as an index of research impact. Replacing ‘quality’ with ‘impact’, raises the question, what is an appropriate index of impact independent of citations? If what we mean by ‘impact’ is societal impact, then it is clear that citation counts are not the right metric *prima facie*, because citations counts primarily reflect academic, not societal consideration. We suggest that regardless of how else we might define impact, it is bound up with an initial assumption that the research in question is high-quality (i.e., replicable, robust, and so on).

number of authors, with the total number of reported statistical tests per paper included as the offset in the model. Including the offset allows us to appropriately model the error *rate* per paper, while adjusting for the total number of reported statistical tests. Year of publication was included because there have been large changes in APA formatting of statistical analyses since 1985, as well as changes in statistical software over this period of time. We also considered it plausible that the number of authors on a paper could affect error rates in a paper (e.g., more authors may check over the work being submitted and so could reduce the errors in a paper) and so we conducted our analyses with and without the number of authors as a covariate in this model. The inclusion of this covariate did not materially change the conclusion of this set of analyses. Finally, we also classified each journal into one of 14 content areas within psychology (clinical, cognitive, counseling, developmental, education, experimental, general, health, industrial/organizational, law, methodology, neuroscience, social, and “other”). The classification of the journals into content areas allowed us to control for potential variation in publication practices across areas by including content area as a grouping-factor. However, journal was not treated as a grouping-factor in the model because it is perfectly co-linear with journal impact factor (i.e., the population and grouping-level effects are identical).

Figure 1 plots the change in error rate (per 100 statistical tests for legibility) as a function of the predictors based on the posterior distribution of the Poisson regression model. As shown, the number of decision errors per 100 statistical tests drops from roughly 4 errors per 100 for the lowest impact factor journals to a little under 2 errors per 100 for the highest impact factor journals, and most of this drop occurs for journals with impact factors between 0 and 2; there is virtually no difference between journals with impact factors greater than 2. Statistically, only journal impact factor $b = -0.210$ ($-0.280, -0.139$), $BF_{10}=597.57$ and year of publication $b = -0.108$ ($-0.134, -0.082$), $BF_{10}=2.05e+06$ predicted decision errors, with both BF’s strongly supporting the alternative hypothesis, indicating 0.23 fewer errors for a one unit increase in the log of journal impact factor, and

0.11 fewer errors for papers published one year more recently, respectively.⁶ Analysis of both citation counts, $b = -0.008$ ($-0.030, 0.015$), $BF_{10}=1/1000$ and number of authors $b = 0.023$ ($-0.001, 0.046$), $BF_{10}=1/142$ strongly supported the null hypothesis. This later finding is consistent with the work of Veldkamp and colleagues (2014).

Table 1 provides *leave-one-out* R^2 's and the Bayes factor, along with the parameter estimates for each predictor when modeled separately. We include these analyses for two reasons. First, we wanted to verify that the specific findings obtained for the full model were not dependent on the inclusion of the covariates. Thus, the analyses provide a 'robustness' test. Second, given that impact factor and citation counts are likely treated independently in faculty evaluations, we wanted to construct models that more closely approximated this type of decision context. For these models, each predictor was entered into the model separately as a fixed effect, including content area as a random effect with total number of reported statistical tests per article controlled by including it as the offset. The Bayes factors were estimated from the posterior distributions using the bayestestR package (Makowski et al., 2019). All of the analyses included in Table 1 support the conclusions from the full model: The evidence for citations and number of authors strongly supports the null, while evidence for JIF strongly supports the presence of a positive relation.

The above findings notwithstanding, it is clear that none of the predictors are of "practical" importance, as evidenced by the R^2 's. Articles in higher impact journals are associated with fewer statistical decision errors; however, the fact that the change in error rate only decreases from 4% for journals with impact factors close to zero to a little under 2% for journals with impact factors greater than 10 implies that the relation may not be particularly diagnostic for judging the quality of individual research papers.⁷

⁶ When only considering journals that have impact factors that are two or greater ($N=36,499$), the relation between journal impact factor and decision errors is diminished, $b = -0.009$ ($-0.116, 0.134$), $BF_{10}=1/166.67$. This suggests that the negative relationship between impact and quality should be interpreted with a considerable amount of caution.

⁷ To put this in perspective, the Kendall's tau rank-order correlation between impact factor and percent

382 The negative relation between the number of errors and the journal impact factor
383 may reflect the quality of the review process (maybe reviewers or editors of higher-impact
384 journals are more likely to catch errors), the strength of the journals editorial team, or
385 (perhaps) the possibility that authors are more careful when submitting to higher impact
386 factor journals. To address the question of quality more directly, we next evaluate whether
387 the strength of evidence presented in papers varies with impact factors or citation counts.

decision errors per paper, an analysis that doesn't control for the other factors, is inconsequentially small and favors the null ($\tau = -.002$, $BF_{10}=1/142$), while the rank order correlation between decision errors and citation counts is also small, the evidence supports a *positive* (Kendall's $\tau = .024$, $BF_{10}=9.1e+7$). Adjusting for the necessary grouping and population effects, however, decreases this small bivariate correlation.

Figure 1. Number of statistical decision errors per 100 tests as function (a) Number of times the article was cited, (b) journal impact factor, (c) number of authors on the paper, (d) year of publication. Shaded regions indicate 95% Credible Intervals. Analyses based on $N=45,144$ published articles.

Question 2: Do highly cited papers and papers in high impact factor journals report stronger statistical evidence to support their claims?

Carl Sagan noted that *extraordinary claims require extraordinary evidence*. From a Bayesian perspective, this can be interpreted to mean that surprising or novel results, or any result deemed a priori unlikely, require greater evidence. Many papers published in high-impact journals are considered a priori unlikely according to both online betting markets and survey data (Camerer et al., 2018). Moreover, editorial policies for some higher impact journals emphasize novelty as a criterion for publication—authors understand that research reports should be “new” and of “broad significance.” Assuming high-impact

Figure 2. Bayesian estimated posterior probability for H_1 for all t -test and 1 degree of freedom F -tests that were statistical significant at the $p < .05$ level ($N = 299,316$) as a function of (a) Number of times the article was cited, (b) 2017 journal impact factor, (c) year of publication, (d) number of authors on the paper, and (e) reported degrees of freedom. Shaded regions indicate 95% Credible Intervals.

journals are more likely to publish novel and surprising research, then one might also expect these journals to require a higher standard of evidence. One way to quantify the strength of the evidence is with the Bayes factor (Morey, Romeijn, & Rouder, 2016). The Bayes factor provides a quantitative estimate of the degree to which the data support the alternative hypothesis versus the null. For instance, a BF_{10} of seven indicates that the data are seven times more likely under the alternative hypothesis than under the null *as those hypotheses are specified*. All else being equal, studies that provide stronger statistical evidence are more likely to produce replicable effects (Etz & Vandekerckhove, 2016), and hold promise for greater scientific impact because they are more likely to be observable across a wide variety of contexts and settings (Ebersole et al., 2016; Klein et al., 2018).

To conduct these analyses, we computed the Bayes factors for all t -tests, one-degree of freedom F -tests, and Pearson r 's ($N = 299,504$ statistical tests from 34,252 journal articles) contained within the Hartgerink (2016) dataset in which the recomputed p -value based on the test-statistic and degrees of freedom was $p \leq .05$. This procedure allows us to evaluate the general strength of evidence presented by authors in support of their claims, but does not allow us to evaluate cases in which authors' theoretical conclusions are based on a single statistic amongst many that might be presented.

The BF was computed using the `meta.ttestBF` function in the `Bayesfactor` package in R using the default Cauchy prior ($r_{scale}=0.707$). This function computes the Bayes factor based on a t -statistic and sample size. The default prior was used in part because it weights the range of effect sizes that are plausible in behavioral science (i.e., small to medium) more heavily than larger effects and in part because we had no other prior beliefs to justify an alternative to the default. We did not evaluate any alternative priors. We first converted the F and r statistics to t statistics before computing the Bayes factor for a given test. Because our automated approach cannot differentiate between within and between-subjects designs, we computed two separate BF's for each statistic, one assuming within subjects and one assuming between-subjects. Sample size for the

meta.ttestBF function was estimated from the degrees of freedom, with $N = df + 1$ for the one-sample case and $N1 = N2 = (df + 2)/2$ for the two-sample case, rounded to the nearest integer. Because the two estimates are highly correlated ($r=0.99$) and yield similar findings, we report only the results of the two-sample case. It is important to note that while this procedure makes the intercept of our models difficult to interpret, our primary interest was in the *slope parameters* – we sought to understand the relationship between evidential strength and both journal impact factors and citation counts, *not* the overall or average magnitude of Bayes factors in the papers in our dataset.

Because the distribution of the Bayes factor was extremely skewed, we converted them to probabilities using $\frac{BF_{10}}{BF_{10}+1}$ and used Beta regression to analyze the data. This conversion is both natural and justified, and does not alter the interpretation of the data, as it merely re-expresses the Bayes factor from the relative probability of two competing models to the probability of a particular model (in our case, the model representing the alternative hypothesis). The Beta regression model is the appropriate approach for analyzing data bounded by 0 and 1. The conversion to the probability scale was not possible for a small number ($N = 188$) of the Bayes factors due to their extremity, resulting in a total of 299,316 total observations being used in the final analysis. The Bayes factors reported below for the model predictors were obtained using the bayestestR package (Makowski et al., 2019). As a robustness test, we also repeated all of the analysis after applying the rank normal transformation to the Bayes factor. These analyses are provided on OSF, and are consistent with the results of the beta regression models reported here.⁸

Under the common assumption that citation counts and impact factors reflect scientific impact or quality, one would expect these indices to be positively correlated with the evidentiary value of the data represented in the publications. However, this does not seem to be the case. In fact, as we demonstrate next, there is a tendency for the

⁸ The log transformation of the Bayes factor did not sufficiently correct for the skewness; models other than the beta model or a Gaussian model applied to the rank normal transformation generally failed posterior predictive checks.

evidentiary value of the data to be *weaker* in high-impact journals.

The “centrality” of a hypothesis can only be assessed indirectly in these analyses, but we work from the assumption that tests reported in a paper, particularly those that are statistically significant, will tend to be the tests most central to an author’s argument. Focusing on significant tests, we used citation counts and impact factors as predictors of a test’s Bayes factor, controlling for content area, year of publication, number of authors of the paper, and the reported degrees of freedom (as a surrogate for sample size).⁹ The magnitude of the Bayes factor was negatively related to the journal impact factor, $b = -0.033$ ($-0.046, -0.021$) and number of authors $b = -0.014$ ($-0.018, -0.010$). Both of these findings strongly support the alternative hypothesis (BF_{10} ’s = 60.52 and 3.91e+03 for impact factor and number of authors respectively). Citation counts also showed a small negative relationship, $b = -0.004$ ($-0.008, -0.0001$), though the evidence convincingly supported for the null hypothesis ($BF_{10}=1/500$) in this case. Both year of publication $b = 0.019$ (0.015, 0.023) and degrees of freedom, $b = 0.184$ (0.124, 0.248) were positively related and showed convincing support for the alternative hypothesis (BF_{10} ’s = 1.31e+07 and 1.30e+05 for year and degrees of freedom, respectively). Table 1 provides the regression coefficients, LOO R^2 ’s, and Bayes factor for each predictor modeled separately. Again, these models are included as both robustness tests and to examine impact factors and citations counts in particular, but also number of authors, independently as they might be used in tenure and promotion contexts. The BF_{10} ’s reported in Table 1 were computed using the bayestestR package (Makowski et al., 2019).

Although the BF’s for each predictor provide convincing support for each predictor, the magnitude of the effects are quite small and likely not of much practical relevance. For instance, as shown in Figure 2, the estimated probability of H_1 decreased from approximately 0.90 to 0.87 (corresponding to $BF_{10}=9$ and $BF_{10}=8$, respectively) across the

⁹ The inclusion of degrees of freedom also serves as a sanity check, as the Bayes factor should scale with sample size, *ceteris paribus*.

full range of impact factors plotted. Similarly, the probability of H_1 is only marginally higher in papers cited a mere 10 times relative to those cited over 1000. Thus, while the strength of evidence presented in papers is clearly related to JIFs, citations, and number of authors, this relation is small. Nevertheless, what should be clear from these analyses is that the use of impact factors and citations as positive indicators of quality is not justified. These indices are, at best, uninformative indicators of research quality (as defined by strength of evidence) and at worst, misleading indicators in which higher impact factors and greater citations reflect *poorer* quality.

The small yet convincing negative relationship between evidentiary value and number of authors indicates that papers with greater numbers of authors were also associated with a given test providing less evidentiary value. To speculate, in the absence of preregistered analysis plans, the number of alternative ways the data could be analyzed may grow as a function of number of authors, a hypothesis consistent with the findings of [Silberzahn et al. \(2018\)](#). In turn, it's possible that weak evidence discovered during extended exploratory data analysis is over-interpreted. However, further research would be needed to confirm this hypothesis. As with analyses of the decision errors, the Bayesian R^2 for these analyses are small, indicating that despite the fact that the results are reliable, none of the predictors account for much variance in the magnitude of the Bayes factor (see Table 1).

Further exploratory analyses. Another way to look at the data is to model the distribution of p -values directly. We next examined the degree to which the magnitude of the computed p -value co-varied with impact factors, citation counts, and number of authors, controlling for year of publication, number of statistical tests reported for each article, and content area within psychology. We modeled only those p -values $\leq .05$ under the assumption that these are the most meaningful of authors' findings: Of the 667,208 total statistical tests included in the dataset, 488,151 of the re-computed p -values were less than .05. Bayesian beta regression again revealed that both the impact factor $b = .041$ (.032,.050) and the number of authors $b = .008$ (.006, .011) predicted the magnitude of the

reported p -values with Bayes factors supporting the alternative hypothesis (BF_{10} 's=1.31e+08 and 3.27 respectively): Journals with higher impact factors and papers with more authors were associated with reporting higher p -values (i.e., p -values closer to the critical threshold of .05), though the BF for number of authors indicates weak evidence. The Bayes factor for citation counts indicated support for the null ($BF_{10}=1/36$), though the majority of the posterior distribution was negative, $b = -.004$ ($-.007, -.002$). The Bayes factor for year of publication clearly supported a negative relation, $b = -.017$ ($-.020, -.014$), $BF_{10}=4.53e+09$, suggesting that more recent publications report smaller p -values to support their claims. Generally, the effect of the log of the total number of tests included in each paper supported the null hypothesis as well, $BF_{10}=1/16$, though again the majority of the posterior distribution was negative $b = -.007$ ($-.01, -.003$). As a robustness check, we fit a separate model without controlling for the total number of reported statistical tests. This analysis did not alter our conclusions.

Question 3: Do citation counts or journal impact factors predict replicability of results?

To address this third question, we used an openly available data set that included data from eight different replication projects, including the Reproducibility project [Open Science Collaboration \(2015\)](#), Many Labs 1 ([Klein et al., 2014](#)), Many Labs 2 ([Klein et al., 2018](#)), and Many Labs 3 ([Ebersole et al., 2016](#)), a special issues of Social Psychology ([Epstude, 2017](#)), the Association of Psychological Sciences Registered Reports Repository ([Simons, Holcombe, & Spellman, 2014](#)), the pre-publication independent replication repository ([Schweinsberg et al., 2016](#)), and Curate Science ([Aarts & LeBel, 2016](#)). [Reinero et al. \(2020\)](#) collated these data, which also includes the number of citations to each original article, the number of authors on the original article, year of publication, and the Altmetric score for the original article. The full dataset curated by [Reinero et al. \(2020\)](#) is available on the Open Science Framework (<https://osf.io/pc9xd/>). We merged these data

Figure 3. Bayesian estimated posterior probability of replication success as a function of (a) 2017 journal impact factor, (b) Number of times article was cited, and (c) number of authors on the paper. Shaded regions indicate 95% Credible Intervals. Ticks on the x-axis indicate distribution of individual cases for the predictor variable.

with the 2017 journal impact factors obtained using the scholar (Keirstead, 2016) package in R and coded each journal by subsdiscipline (clinical, cognitive, social, judgment and decision making, general, marketing, and ‘other’).¹⁰ Of the 196 studies included in Reiner

¹⁰ We also manually looked up each impact factor on the journal website to verify the accuracy of the automated tool. The manually obtained impact factors correlated $r = 0.99$ with impact factors obtained using the Scholar package. We therefore report analyses based on the scholar data. However, we note here that the conclusions based on the manual impact factors are essentially the same.

Figure 4. Violin plots of (a) number of times each paper is cited, (b) number of times the first author was cited, (c) rated institutional prestige, and (d) rated surprisingness of experimental findings split by studies that were and were not successfully replicated. Medians are given by boxed numbers; means are unboxed. Plots are based on the $N = 100$ replication attempts included in [Open Science Collaboration \(2015\)](#)

et al. (2020), we were unable to obtain journal impact factors for four of the publications; an additional two papers were not coded for replication success. Thus, the total number of studies included in our analysis was $N = 190$, of which 80 (42.1%) were successfully replicated.

A subset of the replications included in this dataset are from the [Open Science](#)

Collaboration (2015) replication project and include a variety of other variables, including ratings of the “surprisingness” of each original finding (the extent to which the original finding was surprising, as judged by the authors who participated in the replication project) as well as citations to the first author of the original paper and prestige of the authors institution. Our original analyses provided in early drafts of this manuscript (<https://psyarxiv.com/9g5wk/>) focused only on the 100 studies included in the Open Science Collaboration (2015) repository. We expanded our analysis in March 2021 to include all of the studies reported in Reinero et al. (2020). The Open Science Collaboration (2015) dataset includes only three journals (*Journal of Experimental Psychology: Learning, Memory and Cognition*; *Journal of Personality and Social Psychology*; and *Psychological Science*), and the impact factors for these three journals are completely confounded with content area within psychology. The expanded dataset includes 190 total publications in twenty-seven different journals, thereby enabling us to include impact factor as a predictor in our models.

To examine whether impact factors and citation counts predict replication success (as defined by the original authors who replicated the prior studies), we conducted a Bayesian multi-level logistic regression predicting replication success from the log(impact factor), log(citation count+1) and the log(number of authors) with year of publication and area within psychology included as a random or grouping factors, using weakly informative priors. Figure 3 plots the Bayesian estimated posteriors. The relation between impact factor $b = -0.822$ 95% CI $(-1.433, -0.255)$ and replication success was negative, indicating a 30% decrease in replication success for a one unit increase in the log of a journal impact factor. This is in the opposite direction than would be desired by proponents of their use, though the the Bayes factor was equivocal, indicating that the data do not differentiate between the null or alternative hypotheses ($BF_{10}=1.5$). Analyses of citation counts $b = 0.237$ 95% CI $(-0.010, 0.490)$ and number of authors $b = 0.179$ 95% CI $(-0.178, 0.536)$ indicated that the data strongly support the null hypothesis in both cases (BF_{10} 's=1/13

and 1/34 respectively). As a robustness analysis, we fit alternative models excluding ‘area’ as a grouping factor, modeling the predictors individually, and specifying alternative prior distributions. These models all yielded consistent findings, with the slope of the impact factor consistently negative ranging from -1.03 to -0.53 , and the evidence for the effect of citation counts and number of authors showing support for the null. Note that the bulk of the articles used in this analysis were from journals with impact factors less than 10, and that the relation between impact factors and replication success is strongest in this range. An exploratory analysis eliminating the 8 articles with impact factors greater than 10 indicated relatively stronger support for the negative relation amongst this subset of papers, $b = -1.227$ 95% CI $(-1.990, -0.492)$, $BF_{10}=6.6$.

As noted above, a subset of the data were drawn from the [Open Science Collaboration \(2015\)](#) repository. This dataset included a number of additional variables that might be viewed as reflective of quality, including institutional prestige, number of citations garnered by the author, as well as number of citations to the original paper. As an index of ‘novelty,’ the dataset also includes ratings of *surprisingness*. Figure 4 displays violin and boxplots for the subset of studies in [Open Science Collaboration \(2015\)](#), split by whether they were successfully replicated. We conducted a Bayesian logistic regression using number of citations to the original paper, number of citations to the first author, institutional prestige, and *surprisingness* as predictors of whether the original finding was replicated, with replication success defined based on whether the replication study was statistically significant. Of these indices, only surprisingness showed a relatively strong relation with successful replication $b = -0.663$, 95% CI $(-1.225, -0.121)$, though there was no clear support for either the alternative or null hypothesis ($BF_{10}=1/1.58$). Nevertheless, the observed negative relationship is consistent with other recent work ([Camerer et al., 2018](#); [Hoogeveen, Sarafoglou, & Wagenmakers, 2020](#)) showing that effects viewed as having low prior probability were less likely to replicate. Neither the number of citations to the original paper ($BF_{10}=1/1202$), $b = 0.005$ 95% CI $(-0.003, 0.014)$, the number of citations

garnered by the first author of the original paper ($BF_{10}=1/131062$), $b = -0.00005$ 95% CI
 ($-0.0002, 0.00005$), nor the institutional prestige of the first author ($BF_{10}=1/67$) were
 predictive of replication success $b = 0.014$ ($-0.317, 0.274$) among this subset of
 replications.¹¹

The above analyses are consistent with the results presented for Questions 1 and 2, in
 that neither journal impact factor nor citation counts reflect key aspects of research
 quality. Table 1 shows the summarizes the key findings for all three questions we sought to
 address for all of the single-predictor models. If anything, the evidence seems to suggest
 that higher impact factor journals publish work that is *less* replicable consistent with
 recent research using a smaller, less representative sample (Serra-Garcia & Gneezy, 2021).
 Thus, in regard to the question of whether citation counts or journal impact factors are
positive indicators of quality, as defined by replicability, the answer appears to be “no”, at
 least given the data used to address this question.

Altogether, the only reason for optimism for the use of journal impact factors as a
 metric for research quality is the slight reduction in decision errors in these journals
 (Question 1). Though, as indicated elsewhere in our analyses, this optimism should be
 tempered by the fact that higher-impact factor journals also tend to publish papers that
 present less evidentiary support (Question 2).

Discussion

Goodhart’s Law states that a measure ceases to be a good measure when it becomes
 a target (Goodhart, 1975). Even if citation counts and journal impact factors were once
 reasonable proxies for quality (though there is no evidence to suggest this), it is now clear
 that they have become targets and as a result lack validity as *positive* indicators of faculty

¹¹ We also conducted analyses using the total number of citations to the senior author of each paper and institutional prestige of the senior author’s institution. The conclusions are unchanged.

Table 1

Parameter estimates, Bayesian leave-one-out (LOO) adjusted R^2 for each one predictor modeled separately, and Bayes factors. Each model includes the predictor, with the random-effects factor (Area of psychology), and when appropriate the offset, included as control variables. Bayes factors express the evidence in favor of each model relative to a model with the control variables only. Question 3 includes the year as a random factor as well.

<i>DV</i>	<i>Predictor</i>	<i>Effect (95% CI)</i>	<i>LOO R^2</i>	<i>BF₁₀</i>
Question 1				
<i>Decision Errors</i>	2017 JIF	-.146 (-.214, -.079)	.021	inf
	Citations	.004 (-.016, .024)	.021	1/945.4
	Authors	.002 (.012, -.024)	.021	1/861.65
	Year	-.079 (-.11, -.06)	.022	1.48e+7
Question 2				
<i>Bayes factor</i>				
	2017 JIF	-.056 (-.067, -.044)	.0049	4.23e+15
	Citations	-.014 (-.017, -.011)	.0048	4.31e+12
	Authors	-.010 (-.014, -.006)	.0050	124
	Year	.021 (.018, .025)	.0045	1.75e+23
	Degrees of freedom	.190 (.131, .254)	.0044	1.47e+7
Question 3				
<i>Replication</i>				
	2017 JIF	-.640 (-1.202, -.115)	.022	1/2.31
	Citations	.068 (-.141, .281)	-.007	1/75.32
	Authors	.158 (-.187, .503)	-.005	1/38.51

performance (Fire & Guestrin, 2019). On balance, our findings are consistent with this conclusion: In only one case were either impact factor or citation counts a positive predictor of quality. In almost all other cases, there was either convincing evidence that these variables were associated with poorer quality or convincing evidence for the null. Regardless, in *all* cases, the magnitude of the observed relations suggest that neither citation counts nor impact factors were *meaningfully* related to research quality, as defined by errors in statistical reporting (Question 1), strength of evidence as determined by the Bayes factor (Question 2), or replicability (Question 3). The strongest relation was observed for replicability, though this finding should be interpreted cautiously due to the limited nature of the dataset and the uncertainty of the estimates. Though there is

evidence in Question 1 that impact factor is inversely related to statistical reporting errors (fewer errors in higher impact journals), this finding comes with numerous caveats. For example, the magnitude of this relationship was trivially small and completely absent for the subset of articles in journals with impact factors greater than 2.0 and for articles with fewer than 10% decision errors. Nevertheless, it is possible that some aspect of quality could underlie this relation, whether its due to the quality of the copy-editing team, better reviewers catching obvious statistical reporting errors, or diligence on the part of authors.

More problematic, however, is that some of our analyses indicate that impact factors and citation counts are associated with *poorer* quality. For instance, articles in journals with higher impact factors are associated with *lower* evidentiary value (Question 2) and appear to be less likely to replicate (Question 3) – again noting the considerable uncertainty in the estimates, the magnitude of the R_2 's, and the Bayes factors. Unlike the presence of statistical reporting errors, these later findings cannot be easily dismissed as typographical errors, as they speak directly to the strength of the statistical evidence and are generally consistent with prior results that show a negative relationship between impact factors and estimates of statistical power in social psychological [Fraley and Vazire \(2014\)](#) and cognitive neuroscience [Szucs and Ioannidis \(2017\)](#) journals. Regardless of the causal mechanisms underlying our findings, they converge on the emerging narrative that bibliometric data are weakly, if at all, related to fundamental aspects of research quality (e.g., [Aksnes et al., 2019](#); [Brembs, 2018](#); [Brembs et al., 2013](#); [Fraley & Vazire, 2014](#); [Nieminen et al., 2006](#); [Patterson & Harris, 2009](#); [Ruano et al., 2018](#); [Szucs & Ioannidis, 2017](#); [J. Wang, Veugelers, & Stephan, 2017](#); [West & McIlwaine, 2002](#)).

Practical relevance

Promotion and tenure is, fundamentally, a process for engaging in personnel selection, and it is important that it is recognize it as such. In personnel psychology, researchers often look for metrics that predict desired outcomes or behaviors without creating adverse

649 impact. Adverse impact occurs when a performance metric inherently favors some groups
650 over others. In this context, a good metric that does not create adverse impact might have
651 a modest correlation with work performance yet not produce differential preferences.
652 Conscientiousness, for example, has a validity coefficient of only about $r = .2$ (Barrick &
653 Mount, 1991), but is widely used because it produces minimal adverse impact (Hough &
654 Oswald, 2008). It is questionable that journal impact factors and citation counts are
655 predictively valid and have minimal potential for adverse impact. As suggested above,
656 citation counts and journal impact factors likely play an out-sized role in hiring and
657 promotion and for predicting future research productivity (see also McKiernan et al.,
658 2019). Yet, the research presented above and elsewhere suggests these metrics are
659 overweighted relative to their utility. There is also plenty of evidence to suggest that they
660 hold the potential to produce adverse impact. Citation counts are known to be lower for
661 women and underrepresented minorities (Dion, Sumner, & Mitchell, 2018; Dworkin et al.,
662 2020; X. Wang et al., 2020a), and there is some evidence for a negative relationship
663 between impact factor and women authorship (Bendels, Müller, Brueggmann, &
664 Groneberg, 2018) and so hiring, tenuring, or promoting on their basis may perpetuate
665 structural bias. Second, the present research highlights that the use of these indicators in
666 this context is predicated on the *assumption* that they reflect latent aspects of either
667 research quality or impact. But, inasmuch as the accuracy of statistical reporting (decision
668 errors), evidentiary value (Bayes factor), and replication are key components of research
669 quality, our results are inconsistent with this assumption. More problematic, however, is
670 the potential for the *misuse* of these indicators to ultimately select for bad science. Using
671 an evolutionary computational model, Smaldino and McElreath (2016) showed that the
672 incentive and reward structure in academia can propagate poor scientific methods. Insofar
673 as impact factors and citation counts are used as the basis for hiring and promotion, our
674 analyses, as well as several other recent findings (e.g. Fraley & Vazire, 2014; Szucs &
675 Ioannidis, 2017; J. Wang et al., 2017) are consistent with this model.

The observed inverse relationships between the key indicators of quality on the one hand, and citations and impact factors on the other, have been reported across numerous publications. That is to say, decisions that reward researchers for publishing in high impact journals and for having high citation counts may actually promote the evolution of poor science (Smaldino & McElreath, 2016) as well as select for only certain types of science or scientists (Chapman et al., 2019). Though our analyses indicate that the consequences might be rather small for any given decision, these selection effects could theoretically accumulate over time.

Limitations and Future directions

A rebuttal to our observations that impact factors are either unrelated or negatively related to quality is that our indices of quality (e.g., evidential strength) do not capture other dimensions such as “theoretical importance” or novelty. One can interpret “surprisingness” (see Question 3) as an indicator of novelty, and we are willing to cede that higher impact journals may well publish more novel research on balance (but see J. Wang et al., 2017, for data suggesting otherwise) — indeed, the ratings of ‘surprisingness’ of studies included in the OSF replication project were positively, though weakly associated with Impact factors, $b = 0.151$ 95% CI = (.011, .293). However, acceptance of this assumption only strengthens our argument. If papers in higher impact journals are *a priori* less likely to be true (i.e., because they are novel), then it would require more evidence, not less, to establish their validity. Our analyses suggest that the opposite is actually the case. This finding implies the presence of a perverse trade-off in research evaluation: That reviewers and editors for higher impact journals are more likely trade strength of evidence for perceived novelty. Regardless of the underlying reasons for this pattern, higher impact journals are more likely to be read, cited, and propagated in the literature, despite the fact that they may be based on less evidence.

We note that each of the data sets we used has certain limitations, although they do

not undercut the importance of our results. For example, the program `statcheck` was used to curate the data for Questions 1 and 2. This program only extracts statistics reported in written text and in APA format (e.g., statistics reported in tables are excluded) and does not differentiate between statistics reported for ‘central’ hypotheses versus more ‘peripheral’ hypotheses, though as we have noted, there is some reason to be cautious in drawing this distinction. A second limitation is the studies included in the replication studies for Question 3 were not a representative sample of many psychology studies. Keeping these limitations in mind, all of our analyses converge on the same point: Impact factors and citation counts may be over-weighted relative to their predictive utility.

Several recent working groups have begun drafting recommendations for changes in how the quality of research is evaluated. For instance, [Moher et al. \(2018\)](#) provide six principles to guide research evaluation. Chief amongst these principles are the elimination of bibliometric indicators such as the journal impact factor and citation counting, reducing emphasis on the quantity of publications, and developing new responsible indicators for assessing scientists (RIAS’s) that place greater emphasis on good research practices such as transparency and openness. Moher and colleagues’ recommendations also include the need to incentivize intellectual risk taking, the use of open science practices such as data sharing, the sharing of research material and analysis code, and the need to reward honest and transparent publication of all research regardless of statistical significance. Some of these recommendations may require fundamental change in how researchers and administrators view publication.

To conclude, our analysis supports the growing call to reduce the role that bibliometrics play in the evaluation system, and line up with recommendations made in the San Francisco Declaration On Research Assessment (DORA, <https://sfdora.org/>), the Leiden Manifesto ([Hicks et al., 2015](#)), and by numerous scholars ([Aksnes et al., 2019](#); [Moher et al., 2018](#)). More concretely, several researchers have argued that evaluation should focus more on research process and less on outcome, to incentivize behaviors that

729 support open, transparent, and reproducible science ([Dougherty, Slevc, & Grand, 2019](#);
730 [Frankenhuis & Nettle, 2018](#); [Moher et al., 2018](#)). Changing the focus of the evaluation to
731 *how* faculty conduct their work from *what* is produced may shift the incentives to aspects
732 of the research enterprise that is both under the control of the researcher and arguably
733 promotes good scientific practices.

References

- Aarts, A., & LeBel, E. (2016). Curate science: A platform to gauge the replicability of psychological science.
- Aksnes, D. W., Langfeldt, L., & Wouters, P. (2019). Citations, citation indicators, and research quality: An overview of basic concepts and theories. *SAGE Open*, 9(1), 1–17. Retrieved from <https://doi.org/10.1177/2158244019829575>
- AlShebli, B., Makovi, K., & Rahwan, T. (2020). The association between early career informal mentorship in academic collaborations and junior author performance. *Nature Communications*, 11(1), 1–8.
- Barrick, M. R., & Mount, M. K. (1991). The big five personality dimensions and job performance: A meta-analysis. *Personnel Psychology*, 44(1), 1-26. Retrieved from <https://doi.org/10.1111/j.1744-6570.1991.tb00688.x> doi: 10.1111/j.1744-6570.1991.tb00688.x
- Bendels, M. H., Müller, R., Brueggmann, D., & Groneberg, D. A. (2018). Gender disparities in high-quality research revealed by nature index journals. *PloS One*, 13(1), e0189136.
- Bornmann, L., & Daniel, H.-D. (2008). What do citation counts measure? A review of studies on citing behavior. *Journal of Documentation*, 64(1), 45–80. Retrieved from <https://doi.org/10.1108/00220410810844150>
- Bornmann, L., Mutz, R., & Daniel, H.-D. (2010). A reliability-generalization study of journal peer reviews: A multilevel meta-analysis of inter-rater reliability and its determinants. *PLOS ONE*, 5(12), Article e14331. Retrieved from <https://doi.org/10.1371/journal.pone.0014331>
- Brembs, B. (2018). Prestigious science journals struggle to reach even average reliability. *Frontiers in Human Neuroscience*, 12, Article 37. Retrieved from <https://www.frontiersin.org/article/10.3389/fnhum.2018.00037>
- Brembs, B., Button, K., & Munafò, M. (2013). Deep impact: Unintended consequences of

761 journal rank. *Frontiers in Human Neuroscience*, 7, Article 291. Retrieved from

762 <https://www.frontiersin.org/article/10.3389/fnhum.2013.00291>

763 Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan.

764 *Journal of Statistical Software*, 80(1), 1–28. doi: 10.18637/jss.v080.i01

765 Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., . . .

766 Wu, H. (2018). Evaluating the replicability of social science experiments in Nature

767 and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637–644.

768 Retrieved from <https://doi.org/10.1038/s41562-018-0399-z>

769 Chamberlain, S., Zhu, H., Jahn, N., Boettiger, C., & Ram, K. (2019). rcrossref: Client for

770 various 'crossref' 'apis' [Computer software manual]. Retrieved from

771 <https://CRAN.R-project.org/package=rcrossref> (R package version 0.9.2)

772 Chapman, C. A., Bicca-Marques, J. C., Calvignac-Spencer, S., Fan, P., Fashing, P. J.,

773 Gogarten, J., . . . Chr. Stenseth, N. (2019). Games academics play and their

774 consequences: how authorship, <i>h</i>-index and journal impact factors are

775 shaping the future of academia. *Proceedings of the Royal Society B: Biological*

776 *Sciences*, 286(1916), 20192047. Retrieved from

777 <https://royalsocietypublishing.org/doi/abs/10.1098/rspb.2019.2047> doi:

778 10.1098/rspb.2019.2047

779 Dion, M. L., Sumner, J. L., & Mitchell, S. M. (2018). Gendered citation patterns across

780 political science and social science methodology fields. *Political Analysis*, 26(3),

781 312–327. doi: 10.1017/pan.2018.12

782 Dougherty, M. R., Slevc, L. R., & Grand, J. A. (2019). Making research evaluation more

783 transparent: Aligning research philosophy, institutional values, and reporting.

784 *Perspectives on Psychological Science*, 14(3), 361–375. Retrieved from

785 <https://doi.org/10.1177/1745691618810693>

786 Dworkin, J. D., Linn, K. A., Teich, E. G., Zurn, P., Shinohara, R. T., & Bassett, D. S.

787 (2020, Aug 01). The extent and drivers of gender imbalance in neuroscience reference

- lists. *Nature Neuroscience*, 23(8), 918-926. Retrieved from
<https://doi.org/10.1038/s41593-020-0658-y> doi: 10.1038/s41593-020-0658-y
- Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., ... others (2016). Many labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67, 68–82.
- Edwards, W., Lindman, H., & Savage, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, 70(3), 193–242. Retrieved from
<https://doi.org/10.1037/h0044139>
- Epskamp, S., & Nuijten, M. (2014). statcheck: Extract statistics from articles and recompute p values (r package version 1.0. 0.).
- Epstude, K. (2017). *Towards a replicable and relevant social psychology*. Hogrefe Publishing.
- Etz, A., & Vandekerckhove, J. (2016). A bayesian perspective on the reproducibility project: Psychology. *PloS one*, 11(2), e0149794.
- Fire, M., & Guestrin, C. (2019). Over-optimization of academic publishing metrics: observing Goodhart’s Law in action. *GigaScience*, 8(6). Retrieved from
<https://doi.org/10.1093/gigascience/giz053>
- Fraley, R. C., & Vazire, S. (2014, 10). The N-pact factor: Evaluating the quality of empirical journals with respect to sample size and statistical power. *PLOS ONE*, 9(10), 1–12. Retrieved from <https://doi.org/10.1371/journal.pone.0109019>
- Francois, O. (2015). *Arbitrariness of peer review: A bayesian analysis of the nips experiment*.
- Frankenhuis, W. E., & Nettle, D. (2018). Open science is liberating and can foster creativity. *Perspectives on Psychological Science*, 13(4), 439–447.
- Gargouri, Y., Hajjem, C., Lariviere, V., Gingras, Y., Carr, L., & Brody, T. (2010). Self-selected or mandated, open access increases citation impact for higher quality

research. *PLOS ONE*, 5(10), e13636.

Gelman, A., Goodrich, B., Gabry, J., & Vehtari, A. (2019). R-squared for bayesian regression models. *The American Statistician*, 73(3), 307-309. Retrieved from <https://doi.org/10.1080/00031305.2018.1549100> doi: 10.1080/00031305.2018.1549100

Gentil-Beccot, A., Mele, S., & Brooks, T. C. (2009). Citing and reading behaviours in high-energy physics. *Scientometrics*, 84(2), 345–355.

Goodhart, C. (1975). Problems of monetary management: The U.K. experience. In *Papers in monetary economics, 1975*. Reserve Bank of Australia.

Greenwald, A. G., & Schuh, E. S. (1994). An ethnic bias in scientific citations. *European Journal of Social Psychology*, 24(6), 623-639. doi: 10.1002/ejsp.2420240602

Hartgerink, C. (2016). 688,112 statistical results: Content mining psychology articles for statistical test results. *Data*, 1(3), 14–20.

Hicks, D., Wouters, P., Waltman, L., de Rijcke, S., & Rafols, I. (2015). Bibliometrics: The Leiden Manifesto for research metrics. *Nature*, 520, 429–431.

Hoogeveen, S., Sarafoglou, A., & Wagenmakers, E.-J. (2020). Laypeople can predict which social-science studies will be replicated successfully. *Advances in Methods and Practices in Psychological Science*, 3(3), 267–285.

Hough, L. M., & Oswald, F. L. (2008). Personality testing and industrial–organizational psychology: Reflections, progress, and prospects. *Industrial and Organizational Psychology: Perspectives on Science and Practice*, 1(3), 272-290. Retrieved from <https://doi.org/10.1111/j.1754-9434.2008.00048.x> doi: 10.1111/j.1754-9434.2008.00048.x

Jeffreys, H. (1961). *Theory of probability*. Oxford: Oxford University Press.

Keirstead, J. (2016). scholar: analyse citation data from google scholar [Computer software manual]. Retrieved from <http://github.com/jkeirstead/scholar> (R package version 0.1.5)

- Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr, R. B., Bahník, Š., Bernstein, M. J., . . . others (2014). Investigating variation in replicability. *Social psychology*.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams Jr, R. B., Alper, S., . . . others (2018). Many labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443–490.
- Langford, J., & Guzdial, M. (2015, March). The arbitrariness of reviews, and advice for school administrators. *Commun. ACM*, 58(4), 12–13. Retrieved from <https://doi.org/10.1145/2732417> doi: 10.1145/2732417
- Larivière, V., & Gingras, Y. (2010). The impact factor’s Matthew Effect: A natural experiment in bibliometrics. *Journal of the American Society for Information Science and Technology*, 61(2), 424–427. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.21232>
- Lawrence, N., & Cortes, C. (2014). *The nips experiment*. <http://inverseprobability.com/2014/12/16/the-nips-experiment>. (Accessed: 2021-03-03)
- Lawrence, S. (2001). Free online availability substantially increases a paper’s impact. *Nature*, 411(6837), 521–521. Retrieved from <https://doi.org/10.1038/35079151>
- Luarn, P., & Chiu, Y.-P. (2016). Influence of network density on information diffusion on social network sites: The mediating effects of transmitter activity. *Information Development*, 32(3), 389–397. Retrieved from <https://doi.org/10.1177/0266666914551072>
- Mählck, P., & Persson, O. (2000). Socio-bibliometric mapping of intra-departmental networks. *Scientometrics*, 49(1), 81–91. Retrieved from <https://doi.org/10.1023/A:1005661208810>
- Makowski, D., Ben-Shachar, M. S., & Lüdtke, D. (2019). bayestestr: Describing effects and their uncertainty, existence and significance within the bayesian framework. *Journal of Open Source Software*, 4(40), 1541. Retrieved from

<https://doi.org/10.21105/joss.01541> doi: 10.21105/joss.01541

McKiernan, E. C., Schimanski, L. A., Muñoz Nieves, C., Matthias, L., Niles, M. T., & Alperin, J. P. (2019). Meta-research: Use of the journal impact factor in academic review, promotion, and tenure evaluations. *eLife*, 8, e47338. Retrieved from <https://doi.org/10.7554/eLife.47338>

Moher, D., Naudet, F., Cristea, I. A., Miedema, F., Ioannidis, J. P. A., & Goodman, S. N. (2018). Assessing scientists for hiring, promotion, and tenure. *PLOS Biology*, 16(3), 1–20. Retrieved from <https://doi.org/10.1371/journal.pbio.2004089>

Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of Bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, 72, 6–18. Retrieved from <https://doi.org/10.1016/j.jmp.2015.11.001> (Bayes Factors for Testing Hypotheses in Psychological Research: Practical Relevance and New Developments)

Morey, R. D., & Rouder, J. N. (2018). Bayesfactor: Computation of bayes factors for common designs [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=BayesFactor> (R package version 0.9.12-4.2)

Nature. (2010). *An easy way to boost a paper's citations*. Retrieved from <https://www.nature.com/news/2010/100813/full/news.2010.406.html>

Nature. (2019). *Study suggests 5 ways to increase citation count*. Retrieved from <https://www.natureindex.com/news-blog/studies-research-five-ways-increase-citation-counts>

Nieminen, P., Carpenter, J., Rucker, G., & Schumacher, M. (2006). The relationship between quality of research and citation frequency. *BMC Medical Research Methodology*, 6(1), 42. Retrieved from <https://doi.org/10.1186/1471-2288-6-42>

Nuijten, M. B., Hartgerink, C. H. J., van Assen, M. A. L. M., Epskamp, S., & Wicherts,

- J. M. (2016). The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior Research Methods*, 48(4), 1205–1226. Retrieved from <https://doi.org/10.3758/s13428-015-0664-2>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251). Retrieved from <https://doi.org/10.1126/science.aac4716>
- Paiva, C. E., Lima, J. A. P. d. S. N., & Paiva, B. S. R. (2012). Articles with short titles describing the results are cited more often. *Clinics*, 67, 509–513.
- Patterson, M., & Harris, S. (2009). The relationship between reviewers' quality-scores and number of citations for papers published in the journal physics in medicine and biology from 2003–2005. *Scientometrics*, 80, 343–349. Retrieved from <https://doi.org/doi.org/10.1007/s11192-008-2064-1>
- Peoples, B. K., Midway, S. R., Sackett, D., Lynch, A., & Cooney, P. B. (2016). Twitter predicts citation rates of ecological research. *PLOS ONE*, 11(11), 1–11. Retrieved from <https://doi.org/10.1371/journal.pone.0166570>
- Peters, D. P., & Ceci, S. J. (1982). Peer-review practices of psychological journals: The fate of published articles, submitted again. *Behavioral and Brain Sciences*, 5(2), 187–255. doi: 10.1017/S0140525X00011183
- Petersen, A. M., Fortunato, S., Pan, R. K., Kaski, K., Penner, O., Rungi, A., . . . Pammolli, F. (2014). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences*, 111(43), 15316–15321.
- Reinero, D. A., Wills, J. A., Brady, W. J., Mende-Siedlecki, P., Crawford, J. T., & Bavel, J. J. V. (2020). Is the political slant of psychology research related to scientific replicability? *Perspectives on Psychological Science*, 15(6), 1310–1328. Retrieved from <https://doi.org/10.1177/1745691620924463> (PMID: 32812848) doi: 10.1177/1745691620924463
- Ruano, J., Aguilar-Luque, M., Isla-Tejera, B., Alcalde-Mellado, P., Gay-Mimbrera, J.,

Hernández-Romero, J. L., . . . García-Nieto, A. V. (2018). Relationships between abstract features and methodological quality explained variations of social media activity derived from systematic reviews about psoriasis interventions. *Journal of Clinical Epidemiology*, 101, 35–43. Retrieved from

<https://doi.org/10.1016/j.jclinepi.2018.05.015>

Ruscio, J. (2016). Taking advantage of citation measures of scholarly impact: Hip hip h index! *Perspectives on Psychological Science*, 11(6), 905–908.

Sage. (2022). *Increasing citations and improving your impact factor*. Retrieved from

<https://uk.sagepub.com/en-gb/eur/increasing-citations-and-improving-your-impact-factor>

Schweinsberg, M., Madan, N., Vianello, M., Sommer, S. A., Jordan, J., Tierney, W., . . . others (2016). The pipeline project: Pre-publication independent replications of a single laboratory's research pipeline. *Journal of Experimental Social Psychology*, 66, 55–67.

Seglen, P. (1997). Why the impact factor of journals should not be used for evaluating research. *BMJ (Clinical research ed.)*, 314(7079), 498—502. Retrieved from

<https://europepmc.org/articles/PMC2126010>

Serra-Garcia, M., & Gneezy, U. (2021). Nonreplicable publications are cited more than replicable ones. *Science Advances*, 7(21), eabd1705. Retrieved from

<https://www.science.org/doi/abs/10.1126/sciadv.abd1705> doi:
10.1126/sciadv.abd1705

Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., . . . others (2018). Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science*, 1(3), 337–356.

Simons, D. J., Holcombe, A. O., & Spellman, B. A. (2014). An introduction to registered replication reports at perspectives on psychological science. *Perspectives on*

- Psychological Science*, 9(5), 552-555. Retrieved from
<https://doi.org/10.1177/1745691614543974> (PMID: 26186757) doi:
10.1177/1745691614543974
- Smaldino, P. E., & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, 3(9), 160384. Retrieved from
<https://royalsocietypublishing.org/doi/abs/10.1098/rsos.160384>
- Sternberg, R. J. (2016). "Am I Famous Yet?" Judging scholarly merit in psychological science: An introduction. *Perspectives on Psychological Science*, 11(6), 877-881.
- Szucs, D., & Ioannidis, J. P. A. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLOS Biology*, 15(3), 1-18. Retrieved from <https://doi.org/10.1371/journal.pbio.2000797>
- van Doorn, J., Ly, A., Marsman, M., & Wagenmakers, E.-J. (2018). Bayesian inference for kendall's rank correlation coefficient. *The American Statistician*, 72(4), 303-308. Retrieved from <https://doi.org/10.1080/00031305.2016.1264998> doi:
10.1080/00031305.2016.1264998
- Veldkamp, C. L., Nuijten, M. B., Dominguez-Alvarez, L., Van Assen, M. A., & Wicherts, J. M. (2014). Statistical reporting errors and collaboration on statistical analyses in psychological science. *PloS one*, 9(12), e114876.
- Wang, J., Veugelers, R., & Stephan, P. (2017). Bias against novelty in science: A cautionary tale for users of bibliometric indicators. *Research Policy*, 46(8), 1416-1436. Retrieved from
<http://www.sciencedirect.com/science/article/pii/S0048733317301038>
- Wang, X., Dworkin, J., Zhou, D., Stiso, J., Falk, E., Zurn, P., ... Lydon-Staley, D. M. (2020a). Gendered citation practices in the field of communication.
- Wang, X., Dworkin, J., Zhou, D., Stiso, J., Falk, E. B., Zurn, P., ... Lydon-Staley, D. M. (2020b, Sep). *Gendered citation practices in the field of communication*. PsyArXiv. Retrieved from psyarxiv.com/ywrcq doi: 10.31234/osf.io/ywrcq

- West, R., & McIlwaine, A. (2002). What do citation counts count for in the field of addiction? An empirical evaluation of citation counts and their link with peer ratings of quality. *Addiction*, 97(5), 501–504. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1360-0443.2002.00104.x>
- Zhou, Z. Q., Tse, T. H., & Witheridge, M. (2019). Metamorphic robustness testing: Exposing hidden defects in citation statistics and journal impact factors. *IEEE Transactions on Software Engineering*, 1–1. Retrieved from <https://doi.org/10.1109/TSE.2019.2915065>