



Examining publication bias—a simulation-based evaluation of statistical tests on publication bias

Andreas Schneck

Department of Sociology, Ludwig-Maximilians-Universität München, Munich, Germany

ABSTRACT

Background. Publication bias is a form of scientific misconduct. It threatens the validity of research results and the credibility of science. Although several tests on publication bias exist, no in-depth evaluations are available that examine which test performs best for different research settings.

Methods. Four tests on publication bias, Egger's test (FAT), p-uniform, the test of excess significance (TES), as well as the caliper test, were evaluated in a Monte Carlo simulation. Two different types of publication bias and its degree (0%, 50%, 100%) were simulated. The type of publication bias was defined either as *file-drawer*, meaning the repeated analysis of new datasets, or *p-hacking*, meaning the inclusion of covariates in order to obtain a significant result. In addition, the underlying effect ($\beta = 0, 0.5, 1, 1.5$), effect heterogeneity, the number of observations in the simulated primary studies ($N = 100, 500$), and the number of observations for the publication bias tests ($K = 100, 1,000$) were varied.

Results. All tests evaluated were able to identify publication bias both in the *file-drawer* and *p-hacking* condition. The false positive rates were, with the exception of the 15%- and 20%-caliper test, unbiased. The FAT had the largest statistical power in the *file-drawer* conditions, whereas under *p-hacking* the TES was, except under effect heterogeneity, slightly better. The CTs were, however, inferior to the other tests under effect homogeneity and had a decent statistical power only in conditions with 1,000 primary studies.

Discussion. The FAT is recommended as a test for publication bias in standard meta-analyses with no or only small effect heterogeneity. If two-sided publication bias is suspected as well as under *p-hacking* the TES is the first alternative to the FAT. The 5%-caliper test is recommended under conditions of effect heterogeneity and a large number of primary studies, which may be found if publication bias is examined in a discipline-wide setting when primary studies cover different research problems.

Subjects Ethical Issues, Science Policy, Statistics

Keywords Statistics, Publication bias, Test for excess significance, Caliper test, Monte carlo simulation, p-uniform, Egger's test, FAT

INTRODUCTION

All scientific disciplines try to uncover truth by systematically examining their surrounding environment (*Descartes, 2006: 17*). Natural scientists try to observe regularities in nature, whereas social scientists try to uncover patterns in the social behaviour of humans. The

Submitted 27 June 2017
Accepted 10 November 2017
Published 30 November 2017

Corresponding author
Andreas Schneck,
andreas.schneck@lmu.de

Academic editor
Robert Winkler

Additional Information and
Declarations can be found on
page 15

DOI 10.7717/peerj.4115

© Copyright
2017 Schneck

Distributed under
Creative Commons CC-BY 4.0

OPEN ACCESS

success, as well as the reputation, of science rests on the accuracy and unbiasedness of scientific results. Publication bias, the publication of only positive results confirming the researcher's hypothesis (cf. *Dickersin & Min, 1993*: 135), threatens this validity. Under publication bias, only results showing either statistical significance and/or the desired direction of the effects are published. The published literature in this case is merely a selective (and too optimistic) part of all existing scientific knowledge. Furthermore, science is in the case of publication bias also inefficient as studies that add substantial knowledge to the literature, but contain null-findings remain unpublished.

The study at hand examines the performance of four methods to identify publication bias: Egger's Test/FAT (*Egger et al., 1997*; *Stanley & Doucouliagos, 2014*), p-uniform (PU; *Van Aert, Wicherts & Assen, 2016*; *Van Assen, Aert & Wicherts, 2015*), the test for excess significance (TES; *Ioannidis & Trikalinos, 2007*) and the caliper test (CT; *Gerber & Malhotra, 2008a*; *Gerber & Malhotra, 2008b*). In order to compare the performance of these tests, the false positive rate (α -error, type I error) and the statistical power (true positive rate) were examined in a Monte Carlo simulation study. This makes it possible to assess the performance of the four tests under different conditions of publication bias (*file-drawer* vs. *p-hacking*), as well as study settings (underlying true effect, effect heterogeneity, number of observations in primary studies and in meta-analyses).

The issue of publication bias

The false positive rate of a test (commonly called *p*-value) is the probability of the estimator rejecting H_0 despite this being true. The *p*-value is therefore the probability that the observed estimate is at least as extreme given there is no effect as assumed by H_0 (*Wasserstein & Lazar, 2016*). The larger the *p*-value the higher the risk of assuming an effect if none exists in the data. *p*-values below a certain threshold are called statistically significant, whereas values above the threshold are labelled as non-significant. In the empirical sciences the 5%-significance threshold is mostly used (*Cohen, 1994*; *Labovitz, 1972*; *Nuzzo, 2014*). The difference between 0.049 and 0.051 in the error probability is, however, marginal. Nevertheless, from the standpoint of the 5%-significance threshold the first would be a significant effect, whereas the latter would be a non-significant effect. In both of these two cases, on average around 1 in 20 null-hypotheses of no difference would be rejected, albeit true. If empirical researchers select their data/models until they find, just by chance, significant evidence that seems worth publishing, publication bias is on the rise, leading to inflated or even artificial effects.

Rosenthal (1979) constructs a worst case scenario in which only the 5% of false positive studies that are "significant" solely by pure chance are published. In this case, misinterpreted results shape the scientific discourse and finally result in (medical or political) interventions. Although Rosenthal's example is extreme, a multitude of evidence for publication bias exists in various disciplines and research fields (e.g., *Doucouliagos & Stanley, 2009*; *Jefferson et al., 2012*). *Godlee (2012)* therefore warns that scientific misconduct, under which publication bias is subsumed (*Chalmers, 1990*), may also physically harms patients.

In addition to the societal consequences, publication bias also has severe implications for the evolution of knowledge. Under publication bias no rejection of theories (*Popper, 1968*:

215), on which all scientific progress relies, occurs; this leads to a state of “undead theory” (Ferguson & Heene, 2012: 559) where all existing theories are confirmed irrespective of their truth.

Motivation to commit publication bias

Because statistically significant results stress the originality of research findings (Merton, 1957), Both authors and scientific journals (cf. Coursol & Wagner, 1986; Epstein, 1990; Epstein, 2004; Mahoney, 1977) have large incentives to maximise their significant results to survive in a publish or perish research environment. Authors especially want to increase their publication chances, notably in top-tier journals where low acceptance rates of 5%–10% are quite common (for the top interdisciplinary journals Nature, 2017; Science, 2017; cf. for the political sciences Yoder & Bramlett, 2011: 266). Two distinct strategies to achieve significant results by means of publication bias practices can be pointed out. Firstly, non-significant findings can be suppressed (cf. the classical *file-drawer* effect described by Rosenthal, 1979) and significant results are then searched for in another dataset. Secondly, small bits in the data analysis can be changed (e.g., adding covariates, optional stopping, exclusion of outliers, etc.) until a significant result is obtained—this method is known as *p-hacking* (cf. “fishing” Gelman, 2013; or “researchers degree of freedom” Simmons, Nelson & Simonsohn, 2011: 1359). Whereas the *file-drawer* strategy can be utilised by authors as well as by editors and reviewers, *p-hacking* can only be committed by authors/researchers. Nonetheless, *p-hacking* strategies can be recommended by actors other than authors (e.g., editors, reviewers, etc.). Especially *p-hacking* is almost without any costs, as data analysis tools/packages become increasingly easy to apply (Paldam, 2013).

Evidence on the prevalence of publication bias

So far, there are two strategies for identifying publication bias: the first traces studies through the publication process, the second asks authors, reviewers, or editors about their publication practices via surveys. In the first strategy, most of the analyses trace conference papers or ethics committee decisions if those results get published or remain in the file-drawer. Overall previous findings note, that studies with significant results have a substantially higher chance to get published (cf. Callaham et al., 1998; Coursol & Wagner, 1986; Dickersin, 1990; Easterbrook et al., 1991). Ioannidis (1998) in addition finds that significant studies have, beside their higher publication rate, also a substantially higher publication speed, meaning a shorter time between the completion of the study and the final publication. This results suggest that publication bias is a beneficial strategy in order to maximize academic merits.

The second approach asks directly about the publication practices of the involved actors. In a survey of psychologists that used a sensitive question technique up to 50% of the respondents claimed that they exercised publication bias (John, Loewenstein & Prelec, 2012: 525). Franco, Malhotra & Simonovits (2014) also note that most non-significant findings go to the file-drawer right after the analysis and are not even written up. Also, other forms of misbehaviour, like optional stopping (stopping data collection when significance is reached) or erroneous rounding of *p*-values to reach significant results, are alarmingly widespread

Table 1 Publication bias tests in comparison. compares the four evaluated publication bias tests in respect to four criteria, the measurement level, the sample used by the test, its underlying assumptions and its limitations.

Test	Measurement level	Sample	Assumption	Limitation
FAT	Continuous $[-\infty, \infty]$	All	$\text{Cov}(es, se) = 0$	Only one-sided publication bias (PB) detectable
PU	Continuous $[0, 1]$	$p < 0.05$, effects of same sign	Uniform or right skewed Skewness ≥ 0	Only one-sided PB detectable Only on prespecified levels Effect homogeneity (fixed-effect meta-analysis)
TES	Dichotomous $[0, 1]$	All	$E = O$	Only on prespecified levels Effect homogeneity (fixed-effect meta-analysis)
CT	Dichotomous $[0, 1]$	Threshold \pm caliper width	$P(UC) = P(OC)$	Only on prespecified levels

(prevalence rate around 22.5% *John, Loewenstein & Prelec, 2012: 525*). These results are in line with the survey of *Ulrich & Miller (2017: 9)*, who report that researchers in the field of psychology prefer significant over non-significant results, and, furthermore, attribute more value to results with smaller p -values. These estimates may even be conservative because it is known from the survey literature that sensitive behaviours like scientific misconduct may be underreported (*Kreuter, Presser & Tourangeau, 2008: 848*). According to the presented research results *file-drawer* and *p-hacking* behaviour is therefore quite widespread.

METHODS

Publication bias tests in comparison

So far, the presented detection strategies ask either directly for publication preferences or examine the publication fate of conference papers. Both approaches have the weakness that they either rely on the potentially biased answers of the actors involved or require an immense effort to follow the publication process, while publication bias may have happened before the paper is submitted to a conference. Statistical tests on publication bias circumvent this problem by relying only on the published literature. In the paper at hand the regression-based FAT (*Egger et al., 1997; Stanley & Doucouliagos, 2014*), PU (*Van Aert, Wicherts & Assen, 2016; Van Assen, Aert & Wicherts, 2015*), an extended version of p -curve (*Simonsohn, Nelson & Simmons, 2014a; Simonsohn, Nelson & Simmons, 2014b; Simonsohn, Simmons & Nelson, 2015*), the TES (*Ioannidis & Trikalinos, 2007*), and the CT (*Gerber & Malhotra, 2008a; Gerber & Malhotra, 2008b*) were evaluated (see online [Appendix](#) for an in-depth discussion of the tests).¹

In order to compare the different publication bias tests, four different criteria have to be established: the assumptions of the test, the measurement level, the sample used, the test method, and its according limitations (see [Table 1](#)).

The FAT tests basically the relationship between study's precision and its effect size with all available effect sizes from primary studies. If larger effects are observed for studies with low precision (and low N) publication bias is suspected. Nonetheless, alternative reasons may lead to this result: small studies examine specific high risk populations in which treatments may be more effective (*Sterne et al., 2011*); this effect heterogeneity may lead to the diagnosis of publication bias where none exists (*Schwarzer, Antes & Schumacher, 2002*).² The FAT has furthermore the disadvantage that only one-sided publication bias

¹Because for Fail-save-N (*Rosenthal, 1979*) only rules of thumbs (instead of a formal statistical test) exist it was not included in the simulation at hand. Although it is still widely applied (*Banks, Kepes & McDaniel, 2012: 183; Ferguson & Brannick, 2012: 4*), this benchmark is not recommended in the *Cochrane Handbook*, a guideline for conducting meta-analyses (*Higgins & Green, 2008: 321f.*).

²For a similar result, see *Terrin et al. (2003)* for the related Trim and Fill technique (*Duval & Tweedie, 2000*).

either in favour of a positive or negative significant effect can be tested. *Alinaghi & Reed (2016: 10)* show that if significant results of either sign are searched for, the FAT suffers from massively inflated false positive rates. In the Monte Carlo simulation at hand only the FAT is used because of its better statistical power as shown in prior simulations compared to the similar rank correlation test of *Begg & Mazumdar (1994)*³ and the trim and fill technique (*Duval & Tweedie, 2000*).⁴

³For simulations see: *Hayashino, Noguchi & Fukui (2005)*; *Kicinski (2014)*; *Macaskill, Walter & Irwig (2001)*; *Sterne, Gavaghan & Egger (2000)*.

⁴For simulations see: *Bürkner & Doebler (2014)*; *Kicinski (2014)*; *Moreno et al. (2009)*; *Renkewitz & Keiner (2016)*.

PU has the assumption that every left skewness in the distribution of p -values smaller than the significance threshold (e.g., $p < 0.05$) and conditioned on the underlying observed mean effect (pp -value) is caused by publication bias. This assumption is, however, grounded mainly on the fixed-effect estimate of the mean effect, which is very sensitive to effect heterogeneity. PU furthermore limits its test-value only on significant estimates in the direction where publication bias is suspected (*Van Aert, Wicherts & Assen, 2016: 727*). Therefore PU is, as the FAT, only able to identify one-sided publication bias.

The TES (*Ioannidis & Trikalinos, 2007*; also called ic-index see *Schimmack, 2012*) in contrast relies only on a dichotomous classifier, testing if the number of expected significant results and the empirically observed number of significant effects differ. Because the TES relies, as PU, on the fixed-effect estimate of the mean effect of all included studies it is sensible to effect heterogeneity. A large controversy in the literature is not about the TES itself, but on its application. Francis (*2012a, 2012b, 2012c, 2012d, 2012e, 2013*) used the TES to identify singular articles in order to test if they suffer from publication bias. This may invalidate the assumption of independence (*Morey, 2013: 181*) as well as inflate the false positive rate in a similar manner than in primary research (cp. HARKing *Kerr, 1998*) if the TES is used in such an exploratory manner (*Simonsohn, 2013: 175*). Ioannidis, however, responds that if the TES is applied on prespecified research questions with a large and independent number of effect sizes, the TES is even a conservative test on publication bias (*Ioannidis, 2013: 185*).

The CT uses the most limited sample of the included tests that includes only estimates slightly over and under the chosen significance level in a distribution of z -values. In case of publication bias the assumption of a continuous distribution that results in an approximately even distribution in a narrow interval (caliper) is violated by an overrepresentation of just significant results. The broader the interval is set the more it may deviate from the assumed even distribution caused by the true underlying effect. This restrictive sample has the downside that the exclusion of most available values may drastically reduce the statistical power of the test.

In contrast to the FAT, the other tests are only able to test for publication bias on pre-specified levels (e.g., 0.05). Because the TES and the CT focus only on dichotomous classifiers (significant or not in the case of the TES, slightly over or under the threshold for the CT) also tests on two-sided publication bias are possible.

In previous simulation studies with a low number of included studies as well as observations PU was superior to the TES (*Van Assen, Aert & Wicherts, 2015: 303*) and the FAT (*Renkewitz & Keiner, 2016*). However, no evaluations exists based on a larger number of primary studies. In particular, the newer publication bias tests like PU, the TES, and the CT, are in need of an evaluation under different conditions. For the CT

Table 2 Data generating process (DGP) of Monte Carlo simulation. The 100 conditions of the Monte Carlo simulations are described. Two different aspects were varied: the underlying data and the publication bias behaviour of the actors. For the underlying data the true effect size, the number of observations (N) and the number of studies included in the meta-analysis (K) were varied. The behavioural component altered the proportions of authors who are willing to commit publication bias and its actual form as either *p-hacking* or *file-drawer*.

Conditions	Values	Functional form	N (conditions)
Data setup:			
1. True effects β :	$\beta = 0; 0.5; 1; 1.5; \text{Het}$ $\varepsilon = NV(0, 10)$ $\sigma_x = 2$	$y = \beta x + \varepsilon$	5
2. Number of observations N :	$\mu_N = 100; 500$	$ N(\mu_N) N > 30$	2
3. Number of studies K :	$K = 100; 1,000$		2
Behavioural setup:			
4. Publication bias (PB)	$PB = 0; 0.5; 1.0$	$\beta > 0 \ \& \ p < 0.05$	$1 + 2 * 2 = 5$
4.1. <i>File-drawer</i>	Draw new sample size N	(max. 9 additional samples)	
4.2. <i>p-hacking</i>	Run new analyses with same dataset	$y = \beta x + \gamma_j z_j + \varepsilon$ $z = 0.5x + 0.5y + \varepsilon$ (max. 3 z 's = 7 combinations)	
			$5 * 2 * 2 * 5 = 100$

also no studies exist regarding the best caliper width to use. Despite the existence of some simulation studies on publication bias tests, so far no direct comparison exists that evaluates the performance of all four publication bias tests, especially under effect heterogeneity.

Simulation setup

In order to examine the performance of the four publication bias tests, a Monte Carlo simulation approach is used. For the simulation two different processes have to be distinguished: firstly, the data generation process (DGP), and, secondly, the meta-analytical estimation method (EM). The DGP provides the ground for the hypothetical data used by the simulated actors, as well as the results they report, whereas the EM applies the tests on publication bias reported in the previous section. The central advantage of using Monte Carlo simulations is that controlling the DGP allows to identify which simulated studies suffer from publication bias and which do not. Similar to the case in experiments, different conditions can be defined to ensure a controlled setting. The performance of the estimators can then be examined under the different conditions.

Data setup of the primary studies and meta-analyses

The first step of the DGP defines different effect size conditions that underlie the analyses of the simulated actors (see Table 2). As a first condition, the underlying true effect was specified by a linear relationship with $\beta = 0, 0.5, 1.0, 1.5$. Analogous to a linear regression model, this means for $\beta = 0.5$ that an increase of one unit of the independent variable x increases the dependent variable y by 0.5. The specified linear relationship between the dependent variable y and the independent variable x had a normally distributed regression error term of $\varepsilon = N(0, 10)$, while the variation of the independent variable was defined as $\sigma_x = 2$ (for a similar setup see Alinaghi & Reed, 2016; Paldam, 2015). The regression coefficients can also be transformed in the Pearson correlation coefficient

⁵A total of 16.6% of the simulated studies were adequately powered with at least 80% power (Cohen, 1988: 56). The setting produced by the DGP also reflects the results of Ioannidis, Stanley & Doucouliagos (2017: 245), who report that only 10% of the studies in Economics are adequately powered.

yielding approximately $r = 0, 0.1, 0.2, 0.3$. This results are equivalent to low or medium effect sizes in terms of Cohen (1992) and cover about 75% of the empirical observed effects in psychology (Bosco et al., 2015: 436).⁵ In addition to the homogenous conditions with a common effect size, a heterogeneous condition was added that assumes no fixed distribution of an underlying effect but a uniform mixture of all four effect sizes, as defined above, plus an additional effect of $\beta = 2.0$ ($r = 0.4$) in order to ensure enough variation.

As the FAT is based on study precision, which is mainly driven by the number of observations (N) of the primary studies, N was computed as a second condition by an absolute normal distribution with a mean of 100 (small N) or 500 (large N) and a standard deviation of 150. In order to ensure an adequate statistical analysis for the primary studies, N s equal to or smaller than 30 were excluded. This procedure resulted in a right skewed distribution with a mean N of roughly 500 for the large N , and 165 for the small N condition. The small N condition reflects the observed number of observations in leading economics journals (mean: 152, own computations from the publicly available dataset of Brodeur et al., 2016) as well as of typical trials included in Cochrane reviews (mean: 118, Mallett & Clarke, 2002: 822). Because both studies refer mainly to an experimental literature, the large N condition reflects the more common number of observations especially in ex-post-facto designs (e.g., survey studies).

The heterogeneity of effects in each of the meta-analyses was measured by I^2 , the share of systematic variation in respect to the overall variation consisting of the systematic and random variation (Higgins & Thompson, 2002). In case of the small N 68.62% and for a large N 86.86% of the variation was systematic in the heterogeneous effect condition. In terms of Higgins & Thompson (2002: 1553), an I^2 larger than 50% has to be modelled explicitly in meta-analyses and cannot be ignored.

In addition to the number of observations in the primary studies (N) the number of primary studies that were included in the meta-analysis and form the basis of the publication bias tests (K) was varied in the third condition. A setting with 100 studies was used as a lower condition, whereas 1,000 studies were set as an upper condition. Although on average the number of trials in a meta-analysis is usually much lower than 100 studies (median 28 studies in the meta-meta-analysis by Elia et al., 2016: 5). One hundred studies were chosen because in this setting every publication bias test evaluated is at least partially applicable. In other research areas like Economics, where meta-regression models are more widely used to model effect heterogeneity, higher numbers of included trials estimates are also quite common (e.g., 1,474 effect estimates in Doucouliagos & Stanley, 2009).

Behavioural setup of publication bias

Building on this data setup stage of the DGP the behavioural setup adds publication bias to the simulation in a fourth step (see Table 2). Publication bias was defined as the willingness to collect new data or run additional analyses if statistical significance failed ($p \geq 0.05$) or a negative effect occurred. In the simulation only one-sided publication bias was modelled because both the FAT and PU are not able to model two-sided publication bias that focuses only on significant results irrespective of its sign. It is important to note that only the intent to commit publication bias was varied in the simulation setup. The actual publication bias

depends on the data setup itself: how large is the true effect size (β) and the number of observations (N) in the primary studies? Or, in short: is there already a significant positive result which does not need a publication bias treatment?

Five different publication bias conditions have to be distinguished. Firstly, the condition without publication bias: in this ideal case all estimates (βx) are estimated by a bivariate ordinary least squares (OLS) model and afterwards published. Publishing in terms of the simulation model means that all estimates enter the final meta-analysis. Therefore, in the condition without publication bias either 100 or 1,000 regression results were estimated and enter the meta-analysis.

In the second and third conditions publication bias was present with a 50% probability. That means that 50% of the actors were willing to run additional analyses in order to obtain significant results. These conditions seem closest to the behavioural benchmark of the empirical studies presented.

If a non-significant result was obtained, actors operating under the second condition chose to collect new data in order to obtain significant results that can be published. This second condition therefore modelled publication bias under the *file-drawer* scenario, because the datasets not used remained unpublished. An actor tried to run analyses on the basis of up to nine additional datasets and only stopped earlier if a significant result with a positive sign was obtained. If none of the 10 datasets yielded a significant relationship with a positive sign, the estimate which was closest to the significance threshold has been published. This rule served two purposes: firstly, it seemed plausible that an actor who has tried that many analyses wants to get the results published in the end to compensate for the invested effort and to avoid sunk costs (*Thaler, 1980*). Secondly, from a technical point of view, this allowed to keep the number of observations in a meta-analysis K constant across all simulation conditions.

In the third condition an actor did not try to achieve significant results by running the same bivariate analysis on different samples, but rather tried to run different model specifications on the same data by including control variables (z_j) to achieve statistical significance of the coefficient of interest (βx). The third condition therefore modelled publication bias as a form of *p-hacking*, because the existing dataset was optimised to receive a significant p -value. The actor was able to add three different control variables to the model. The control variables were defined as collider variables that are both an effect of x as well as y , which biases the effect of interest (*Cole et al., 2009; Greenland, Pearl & Robins, 1999*). The effect of x and y on z_j was, however, only small ($\gamma = 0.5$). The error term of the equation defining z was normally distributed $N(0,10)$. With three available control variables z_j an actor had seven different combinations to improve the research results in order to obtain a significant effect of x on y .

In contrast to the second and third conditions, where 50% of the actors had the intention to commit publication bias, in the fourth and fifth conditions all actors had the intention to engage in publication bias practices, once again either through *file-drawer* (fourth condition) or *p-hacking* behaviour (fifth condition). Part from the higher degree of intention to engage in publication bias practices the settings remained the same. Although the two conditions where all actors had the intent to engage in publication bias are far

⁶In order to specify the number of replications that are necessary to achieve a sufficient statistical power of at least 80% (Cohen, 1988: 56), a power analysis was conducted for the statistical power as well as the false positive rate estimates. For the false positive rate, a small deviation of 1 percentage point from the set 5%-false positive rate has to be correctly identified with at least an 80% chance. To achieve this goal, every condition without publication bias had to be supported with 3,729 runs. As deviations in power are, though important but not as essential as the false positive rate (Cohen, 1988: 56) a difference of 3 percentage points is set as acceptable. In order to identify a 3 percentage point deviation from the target power of 80% each of the 80 conditions with existing publication bias needed 1,545 runs. In total, 198,080 runs were necessary, resulting in nearly 109 million primary studies that in the case of publication bias contained up to 10 different regression models.

too pessimistic, they allow to evaluate the performance of the tests in the most extreme publication bias environment. Tests that are not able to detect publication bias even under such extreme conditions are of low utility to the research community.

The resulting design matrix had 100 different combinations resulting from 20 data setup conditions multiplied by the five publication bias conditions. In order to obtain reliable estimates similarly to an experiment (Carsey & Harden, 2013: 4f.), every single cell of the design matrix had to be replicated multiple times.⁶

The aim of the simulation study at hand was to compare the performance of the four tests in respect of: (A) their capability to detect publication bias if present (true positive, statistical power), as well as (B) consistent false positive classification (α -error). Because the conditions with and without publication bias are known in a simulation study, the power of the tests and the false positive rate is computable (Mooney, 1997: 77–79). In a first step, a dummy variable (s) was constructed, with the value 1 for a significant test result below the significance threshold (5% significance level; $s = 1$ if $p < 0.05$). The statistical power, was defined as the proportion of significant results s in respect to all runs (r) with publication bias ($\sum_{i=1}^r s_i/r$ if $PB > 0$). The false positive rate was computed equivalently but in conditions without publication bias ($\sum_{i=1}^r s_i/r$ if $PB = 0$).

RESULTS

Prevalence of publication bias

Because publication bias in the experimental setup was implemented as the intent to commit publication bias, three variables are useful to address the actual publication bias and its impact on the overall bias. Firstly, the share of actual studies per meta-analysis that suffer from publication bias (if $p < 0.05$ or negative result are obtained as a first result), secondly the share of studies that achieve their goal of a significant positive result by publication bias, and thirdly the impact of publication bias on the p -value of a fixed-effect meta-analysis (deflation factor of the p -value). Because the heterogeneous effect condition of the simulation does not allow an absolute bias measure the p -value deflation factor was used for all conditions.

In a first step the focus is on how the opportunity structures of the simulation conditions shape the committed publication bias. In the first two columns of Table 3 the actual committed publication bias is shown dependent on study characteristics like the mean number of observations in the primary studies (N) and the underlying effect (β), including effect heterogeneity. As expected, around 50% respectively 100% of the studies committed publication bias in case of an underlying null-effect because only 2.5% of the results had the right positive sign and were significant just by chance. In case of an underlying effect the share of committed publication bias decreased because an already significant finding made a publication bias treatment unnecessary. For $\beta = 0.5$ in the 50% publication bias condition only 35%; for $\beta = 1$, 15% for $\beta = 1.5$ only 9%; and in the heterogeneous condition 22% of the studies employed publication bias practices. The 100% publication bias condition approximately doubled the prevalence rates of the 50% condition as expected.

Table 3 Risk factors for publication and its impact on bias in the simulated data (OLS regression). The first two columns in Table 3 show that actual committed publication bias behaviour depended largely on the opportunity structure of the underlying data. Despite the defined 50% or 100% willingness of the authors to commit publication bias, only those actors who face insignificant effects (caused by small effects and sample sizes) engaged in publication bias practices. The success of publication bias in terms of significant results is shown in column three, dependent on the opportunity structure and form of publication bias. Conditions under p-hacking were slightly less effective in obtaining significant results than conditions under file-drawer publication bias. Column four shows the deflating impact of publication bias on meta-analytic p -values. For an average publication bias this p -values halved or even quartered.

	Publication bias committed (50% intention)	Publication bias committed (100% intention)	Publication bias successful (in relation to committed)	Deflation of p -value
$N = 500$	-0.105***	-0.211***	0.179***	
(ref. $N = 100$)	(0.000)	(0.001)	(0.001)	
$\beta = 0.5$	-0.196***	-0.391***	0.471***	
(ref. $\beta = 0$)	(0.001)	(0.001)	(0.001)	
$\beta = 1$	-0.389***	-0.777***	0.513***	
	(0.001)	(0.001)	(0.001)	
$\beta = 1.5$	-0.451***	-0.899***	0.503***	
	(0.001)	(0.001)	(0.002)	
$\beta = \text{heterogeneous}$	-0.319***	-0.636***	0.358***	
	(0.001)	(0.001)	(0.001)	
p-hacking			-0.100***	0.197***
(ref. file-drawer)			(0.001)	(0.002)
Committed PB [+10ppts]				-0.018***
(ref. mean = 32.5%)				(0.001)
Successful PB [+10ppts]				-0.077***
(ref. mean = 18.8%)				(0.001)
Constant	0.541***	1.080***	0.313***	0.225***
	(0.000)	(0.001)	(0.001)	(0.002)
Observations	61,760	61,760	115,843	123,520
R^2	0.939	0.958	0.648	0.168

Notes.

Standard errors in parentheses.

* $p < 0.05$.

** $p < 0.01$.

*** $p < 0.001$.

Besides the necessity of publication bias to achieve significant results also the success probability in respect to committed publication bias depended on the conditions of the primary studies as shown by the third column of Table 3. For small studies ($N = 100$) with an underlying null-effect ($\beta = 0$) the success-probability in respect to the committed publication bias was about 31.3%. Publication bias got more effective if larger studies ($N = 500$) provide the primary study with more statistical power. The success probability of publication bias rose dramatically around 50 percentage points if a specific underlying empirical effect existed. Also, in case of effect heterogeneity, the success probability increased about 35.8 percentage points. Slight differences could be observed in the effectivity

of the publication bias mechanism, as *p-hacking* was with 10 percentage points less effective than the *file-drawer* condition to achieve significant results.

As publication bias deflates *p*-values and therefore biases meta-analytical effect estimates the impact of the actual observed publication bias (the share of committed and successful publication bias) on the meta-analytical *p*-value is presented. The fourth column of [Table 3](#) shows that with an average proportion of publication bias committed (32.6%) as well successfully implemented (18.8%) in a meta-analysis with 100 studies ($K = 100$) the *p*-value of the meta-analysis more than quartered. This is further aggravated if the share of committed as well as successful publication bias rose by 10 percentage points. The actual impact of successful publication bias deflated the *p*-values by 7.7 percentage points and was more pronounced than the deflation caused by non-successfully committed publication bias (deflation by 1.7 percentage points). The deflation also was less severe if *p-hacking* procedures, as implemented in the simulation at hand, were used. Nonetheless, the meta-analytical *p*-value in case of *p-hacking* is still less than half the size (42.2%) of the unbiased estimate.

False positive rate of publication bias tests

For the evaluated tests on publication bias consistent false positive rates are most important. In the simulation none of the tests should exceed the prespecified 5% error probability in any condition. The false positive rate of the test was fixed in the simulation setting to 0.05, so all false positive rates should be equal to, or even smaller than, 0.05. Positive deviations from 0.05 point to inflated false positive rates, which lead to more false conclusions than expected.

[Table 4](#) shows the false positive rate in dependence of the number of studies included ($K = 100, 1,000$) and effect heterogeneity measured by I^2 . In the constant condition of a meta-analysis with $K = 100$ and no effect heterogeneity none of the tests had larger false positive rates than the expected 0.05. In particular, the TES, the 3% and 5% CTs were very conservative. A larger meta-analytical sample increased the false positive rates for the TES and the CTs. The broadest 15% CT missed the expected significance threshold of 5%, with 7.8% clearly. The false positive rates for PU in contrast were slightly lower. Increasing effect heterogeneity resulted in more conservative false positive rates for PU, the 15% CT, and to a smaller extent also for the FAT, the TES and the 10% CT. The narrower 3% and 5% CTs were unaffected by effect heterogeneity.

The overall influence of the varied conditions on the false positive rate was small, as can be seen by the small share of explained variance ($R^2 < 1.7\%$). Looking at the false positive rates by each condition ([Table A1](#) in the online [Appendix](#)) only the 10%- and 15%-caliper showed increased false positive rates because the underlying true effect rather than publication bias elicited an overrepresentation of just significant values. Note however, that the 3% and 5% CTs showed no increased false positive rates.

Statistical power of publication bias tests

The following regression model ([Table 5](#)) addresses the statistical power conditional on the type of publication bias and its occurrence (committed as well as successful publication

Table 4 Conditional false positive rates of the publication bias tests (OLS regression). Table 4 displays the false positive rates of the publication bias tests conditional on the number of studies included in the meta-analysis (K) as well as the between study heterogeneity (I^2). The FAT had the most consistent false positive rate. The 15% CT missed the 5%-level clearly while the 10% CT showed a large variability and gets close to it. The 10% and 15% CT are therefore problematic because they may suffer from inflated false positive rates.

	PU	FAT	TES	3% CT	5% CT	10% CT	15% CT
$K = 1,000$ (ref. $K = 100$)	-0.005*** (0.001)	0.000 (0.002)	0.006*** (0.001)	0.026*** (0.001)	0.023*** (0.001)	0.026*** (0.001)	0.050*** (0.002)
I^2 [+10 percentage points]	-0.003*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	0.000 (0.000)	0.000 (0.000)	-0.001*** (0.000)	-0.003*** (0.000)
Constant ^a	0.023*** (0.001)	0.049 (0.001)	0.010*** (0.001)	0.003** (0.001)	0.008*** (0.001)	0.019*** (0.001)	0.028*** (0.001)
Observations	73,960	74,560	74,560	62,644	66,546	69,718	70,936
R^2	0.005	0.000	0.002	0.010	0.007	0.006	0.017

Notes.

^aTest H0: constant = 0.05.

Standard errors in parentheses.

* $p < 0.05$.

** $p < 0.01$.

*** $p < 0.001$.

bias). Starting from the baseline condition of a meta-analysis with $K = 100$, a mean share of publication bias committed (32.6%), as well as successfully applied (18.8%) via a *file-drawer* procedure and no effect heterogeneity, the FAT had a superior power of 56.9%, followed by the TES (51.5%) and the PU (48.3%). The CTs performed worst and yielded only a power of 0.0%–38.6%.

The underperformance of the CTs is largely explained by the small number of studies in the meta-analyses. With $K = 100$ hardly any study falls within the small caliper around the significance threshold. This limitation on just significant or non-significant effects also led to missing values, because without observations in the caliper no CT could be performed. The underperformance of the CT changed if 1,000 studies were included, which improved the estimated power substantially, by 30.7–57.3 percentage points, while smaller calipers profited most. The FAT, the TES, and the PU, profited moderately from an increased number of studies, by 24.4, 23.8 and 16.5 percentage points, respectively. When focussing on the influence of heterogeneity in the meta-analyses the PU and the TES showed a drastic drop in power, by 6.5 and 6.4 percentage points, if the heterogeneity measured by I^2 rose by 10 percentage points. This decrease in power shows that neither PU nor TES were able to cope with heterogeneity. In contrast, the FAT and the CTs actually showed a slight increased statistical power. Varying the publication bias procedure from *file-drawer* to *p-hacking*, which is less related to the standard error of the effect estimates, increased the power of PU, TES, and the CTs. The CTs profited most, increasing the statistical power by around 18 percentage points. The TES and PU showed a smaller increase of power, by 7.5 and 4.8 percentage points. The FAT, in contrast lost about 11 percentage points of its power under *p-hacking* compared to the *file-drawer* condition. Although the differences in power are dependent especially on the operationalization of the *p-hacking* condition in the simulation, this result points on a weakness of the FAT under non *file-drawer* conditions

Table 5 Conditional statistical power of the publication bias tests (OLS regression). shows the statistical power of the publication bias tests conditional on the number of studies included in the meta-analysis (K) and the between study heterogeneity (I^2). In contrast to Table 4, also the share of committed as well as successful publication bias and its form as either file-drawer or *p-hacking* was controlled. Overall the FAT had the largest power but was not able to detect *p-hacking* as good as the TES. The CTs were underpowered if a low number of studies was included in a meta-analysis but performed well in studies with large K s. Both, PU and the TES, were not able to detect publication bias under effect heterogeneity.

	PU	FAT	TES	3% CT	5% CT	10% CT	15% CT
$K = 1,000$	0.165***	0.244***	0.238***	0.573***	0.513***	0.382***	0.307***
(ref. $K = 100$)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
I^2 [+10 percentage points]	-0.065***	0.001*	-0.064***	0.006***	0.005***	0.008***	0.010***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
<i>p-hacking</i>	0.048***	-0.110***	0.075***	0.179***	0.187***	0.186***	0.177***
(ref. file-drawer)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Comitted PB [+10pppts]	0.051***	0.030***	-0.065***	-0.035***	-0.053***	-0.073***	-0.084***
(ref. mean = 32.6%)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Successful PB [+10pppts]	0.103***	0.099***	0.221***	0.162***	0.193***	0.224***	0.234***
(ref. mean = 18.8%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Constant ^a	0.483***	0.569***	0.515***	-0.002	0.125***	0.300***	0.386***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	123,520	123,520	123,520	107,736	111,315	115,243	117,207
R^2	0.572	0.306	0.473	0.497	0.483	0.457	0.446

Notes.

^aTest H0: constant = 0.8.

Standard errors in parentheses.

* $p < 0.05$.

** $p < 0.01$.

*** $p < 0.001$.

that are less related to the standard error of the estimate but are still detectable in the distribution of z - or p -values.

The structural difference between tests based on a continuous effect distribution (FAT, PU) and tests that focus only on a dichotomous classification (TES, CTs)⁷ becomes evident looking at the effect of the proportion of studies that underwent a publication bias treatment in the simulation and the proportion of studies that had a successful outcome after publication bias. Increasing the share of studies under publication bias lifted the power by 3.0 (FAT) and 5.1 (PU) percentage points. A 10 percentage point increase in studies successfully applying publication bias increases the power by 9.9 (FAT) and 10.3 percentage points (PU). The TES and the CTs, however, were only able to detect successful publication bias. An increase only in studies committing publication bias (whether successful or not) in contrast reduced the statistical power. Both tests were therefore not able to detect all possible outcomes of publication bias. This is especially problematic as non-successful publication bias may also inflate the overall estimated effect in meta-analyses. All effects presented are statistically significant ($p < 0.05$).

In contrast to the influence of the varied conditions on the false positive rate, the influence on statistical power was substantial, varying from 30.6% in the case of the FAT to 57.2% for the PU. This finding underlines the fact that all publication bias tests have their strengths and weaknesses in specific conditions.

⁷Significant or not (TES) over- or under-caliper (CTs).

DISCUSSION & CONCLUSION

In the simulation at hand, the performance of four different tests (PU, FAT, TES, CTs) were evaluated in a Monte Carlo simulation. Different conditions were varied: the underlying true effect size, including effect heterogeneity, the number of observations in the primary studies, the number of studies in the meta-analyses, the degree of publication bias and its form as either *file-drawer* or *p-hacking*.

Limitations

In order to compare the tests in a realistic setting that is nonetheless at least from the assumptions of all four tests applicable, four central limitations have to be pointed out:

Firstly the simulation and its according publication bias procedures rest on the assumptions that all the correlation between the study's precision and its effect size is caused by publication bias. In case that studies with larger effects are, for example after a pre-study power analysis (*Lau et al., 2006*) conducted with a lower number of observation especially the FAT may yield increased false positive rates (*Schwarzer, Antes & Schumacher, 2002*).

Secondly the number of observations included in the meta-analyses either set to $K = 100$ or $1,000$ is large compared to the average meta-analysis (*Elia et al., 2016*). The results however showed that even in such large meta-analyses and especially in the more realistic condition in which 50% of the actors are willing to commit publication bias, the tests hardly yielded an adequate statistical power under most conditions. Increasing the number of included studies is therefore important to assure an adequately powered test on publication bias.

Thirdly the analysis focused only on one specific form of *p-hacking* that could occur in both small ($N = 100$) or large studies ($N = 500$). Especially for studies where N is small, other strategies like optional stopping may also be applied. Further research on publication bias should therefore focus on the different impact of other *p-hacking* practices.

As a fourth limitation only one-sided publication bias against insignificant or negative results was simulated. By assumption especially PU limits only on the negative or positive signed studies that were supposed to be affected by publication bias. Also the FAT is not able to detect two-sided publication bias because the funnel in this situation may still be perfectly symmetric. The suggestions for applications if two-sided publication bias is suspected are therefore limited to the TES and the CTs only.

CONCLUSION

The following five conclusions can be derived from the results: Firstly, for homogenous research settings and with publication bias favouring only effects in one direction (one-sided publication bias) the FAT is recommended due to its most consistent false positive rate as well as its superior statistical power. Secondly, if there are concerns whether there are any correlations between the precision of the study and its effect size for other reasons than publication bias (see first limitation) and if *p-hacking* is suspected, the TES should be preferred to the FAT under effect homogeneity. As the 5% CT offers more relaxed assumptions it is therefore the first alternative for the FAT under effect heterogeneity if a large number of studies is included in the meta-analysis.

Despite the analysis focussed only on one-sided publication bias, also two-sided publication bias, favouring significant results with either sign may also be present. As PU and the FAT are not able to identify one-sided publication bias only the TES and the CTs remain for two-sided publication bias. Therefore, fourthly, the TES is recommended under effect homogeneity because of its larger statistical power compared to the CTs. Fifthly, in the case of heterogeneous effect sizes and a sufficient number of observations in the meta-analysis the 5% CT provides the best trade-off between a conservative false positive rate and a decent statistical power.

The 5% CT is therefore best used to identify publication bias in an effect heterogeneous discipline-wide setting which relies per definition on completely different underlying effects but offers enough studies to compensate for the low statistical power. Because the wider 10% and 15% CTs yield inflated false positive rates, at least in some conditions, they are not recommended to identify publication bias.

Identifying publication bias in substantial meta-analyses as well as focussing on publication as a general problem within the scientific domain is necessary in order to establish and retain trust in scientific results. Further research, however, should not only focus on the diagnosis of publication bias just stating a problem that is well known ([Morey, 2013](#)). Beyond the nonetheless important diagnosis of the scientific “disease” a further examination of the risk factors, either on the side of the involved actors or with regard to the incentive structure within the discipline (see for example [Auspurg & Hinz, 2011](#)) seem essential. This includes also the evaluation of possible interventions (e.g., an open data policy). Research on publication bias is inevitable to maintain trust in scientific results and avoid wasted research funds that also limit the efficiency of science as a whole.

Beside the diagnosis of publication bias and its risk factors, also estimators of the unbiased effect, that are beyond the scope of this paper, like the effect estimates provided by PU and the PET (for example the PET/PEESE procedure of [Stanley & Doucouliagos, 2014](#)) should be evaluated comparatively. This is at most important for meta-analyses with a heterogeneous effect that try to uncover the underlying true effect rather than test for publication bias alone.

ACKNOWLEDGEMENTS

I thank Katrin Auspurg, the two reviewers, Joseph Hilgard, Bob Reed and the responsible editor Robert Winkler for their valuable comments on this article.

ADDITIONAL INFORMATION AND DECLARATIONS

Funding

The authors received no funding for this work.

Competing Interests

The authors declare there are no competing interests.

Author Contributions

- Andreas Schneck conceived and designed the experiments, performed the experiments, analyzed the data, contributed reagents/materials/analysis tools, wrote the paper, prepared figures and/or tables, reviewed drafts of the paper.

Data Availability

The following information was supplied regarding data availability:

Further information on the analyses is contained in the readme.txt contained in the [Supplemental zip File](#).

Supplemental Information

Supplemental information for this article can be found online at <http://dx.doi.org/10.7717/peerj.4115#supplemental-information>.

REFERENCES

- Alinaghi N, Reed WR. 2016.** Meta-analysis and publication bias: how well does the FATPET-PEESE procedure work? Available at <https://ideas.repec.org/p/cbt/econwp/16-26.html>.
- Auspurg K, Hinz T. 2011.** What fuels publication bias? theoretical and empirical analyses of risk factors using the caliper test. *Journal of Economics and Statistics* **231**:636–660.
- Banks GC, Kepes S, McDaniel MA. 2012.** Publication bias: a call for improved meta-analytic practice in the organizational sciences. *International Journal of Selection and Assessment* **20**:182–196 DOI [10.1111/j.1468-2389.2012.00591.x](https://doi.org/10.1111/j.1468-2389.2012.00591.x).
- Begg CB, Mazumdar M. 1994.** Operating characteristics of a bank correlation test for publication bias. *Biometrics* **50**:1088–1101 DOI [10.2307/2533446](https://doi.org/10.2307/2533446).
- Bosco FA, Aguinis H, Singh K, Field JG, Pierce CA. 2015.** Correlational effect size benchmarks. *Journal of Applied Psychology* **100**:431–449 DOI [10.1037/a0038047](https://doi.org/10.1037/a0038047).
- Brodeur A, Le M, Sangnier M, Zylberberg Y. 2016.** Star wars: the empirics strike back. *American Economic Journal-Applied Economics* **8**:1–32 DOI [10.1257/app.20150044](https://doi.org/10.1257/app.20150044).
- Bürkner P-C, Doebler P. 2014.** Testing for publication bias in diagnostic meta-analysis: a simulation study. *Statistics in Medicine* **33**:3061–3077 DOI [10.1002/sim.6177](https://doi.org/10.1002/sim.6177).
- Callahan ML, Wears RL, Weber EJ, Barton C, Young G. 1998.** Positive-outcome bias and other limitations in the outcome of research abstracts submitted to a scientific meeting. *Journal of the American Medical Association* **280**:254–257 DOI [10.1001/jama.280.3.254](https://doi.org/10.1001/jama.280.3.254).
- Carsey TM, Harden JJ. 2013.** *Monte Carlo simulation and resampling methods for social science*. Los Angeles: Sage Publications.
- Chalmers L. 1990.** Underreporting research is scientific misconduct. *Journal of the American Medical Association* **263**:1405–1408 DOI [10.1001/jama.1990.03440100121018](https://doi.org/10.1001/jama.1990.03440100121018).
- Cohen J. 1988.** *Statistical power analysis for the behavioral sciences*. Hillsdale: L. Erlbaum Associates.
- Cohen J. 1992.** A power primer. *Psychological Bulletin* **112**:155–159 DOI [10.1037/0033-2909.112.1.155](https://doi.org/10.1037/0033-2909.112.1.155).

- Cohen J. 1994.** The earth is round ($p < .05$). *American Psychologist* **49**:997–1003
[DOI 10.1037/0003-066x.49.12.997](https://doi.org/10.1037/0003-066x.49.12.997).
- Cole SR, Platt RW, Schisterman EF, Chu H, Westreich D, Richardson D, Poole C. 2009.** Illustrating bias due to conditioning on a collider. *International Journal of Epidemiology* **39**:417–420 [DOI 10.1093/ije/dyp334](https://doi.org/10.1093/ije/dyp334).
- Coursol A, Wagner EE. 1986.** Effect of positive findings on submission and acceptance: a note of meta-analysis bias. *Professional Psychology: Research and Practice* **17**:136–137
[DOI 10.1037/0735-7028.17.2.136](https://doi.org/10.1037/0735-7028.17.2.136).
- Descartes R. 2006.** *A discourse on the method of correctly conducting one's reason and seeking truth in the sciences*. Oxford: Oxford University Press.
- Dickersin K. 1990.** The existence of publication bias and risk factors for its occurrence. *Journal of the American Medical Association* **263**:1385–1389
[DOI 10.1001/jama.1990.03440100097014](https://doi.org/10.1001/jama.1990.03440100097014).
- Dickersin K, Min Y-I. 1993.** Publication bias—the problem that Won't go away. *Annals of the New York Academy of Sciences* **703**:135–148
[DOI 10.1111/j.1749-6632.1993.tb26343.x](https://doi.org/10.1111/j.1749-6632.1993.tb26343.x).
- Doucoulagos H, Stanley TD. 2009.** Publication selection bias in minimum-wage research? a meta-regression analysis. *British Journal of Industrial Relations* **47**:406–428
[DOI 10.1111/j.1467-8543.2009.00723.x](https://doi.org/10.1111/j.1467-8543.2009.00723.x).
- Duval S, Tweedie R. 2000.** Trim and fill: a simple funnel-plot-based method of testing and adjusting for publication bias in meta-analysis. *Biometrics* **56**:455–463
[DOI 10.1111/j.0006-341X.2000.00455.x](https://doi.org/10.1111/j.0006-341X.2000.00455.x).
- Easterbrook PJ, Berlin JA, Gopalan R, Matthews DR. 1991.** Publication bias in clinical research. *Lancet* **337**:867–872 [DOI 10.1016/0140-6736\(91\)90201-Y](https://doi.org/10.1016/0140-6736(91)90201-Y).
- Egger M, Smith GD, Schneider M, Minder C. 1997.** Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal* **315**:629–634
[DOI 10.1136/bmj.315.7109.629](https://doi.org/10.1136/bmj.315.7109.629).
- Elia N, Von Elm E, Chatagner A, Pöpping DM, Tramèr MR. 2016.** How do authors of systematic reviews deal with research malpractice and misconduct in original studies? a cross-sectional analysis of systematic reviews and survey of their authors. *BMJ Open* **6**(3):e010442 [DOI 10.1136/bmjopen-2015-010442](https://doi.org/10.1136/bmjopen-2015-010442).
- Epstein WM. 1990.** Confirmation bias among social-work journals. *Science Technology & Human Values* **15**:9–38 [DOI 10.1177/016224399001500102](https://doi.org/10.1177/016224399001500102).
- Epstein WM. 2004.** Confirmation bias and the quality of the editorial processes among american social work journals. *Research on Social Work Practice* **14**:450–458 [DOI 10.1177/1049731504265838](https://doi.org/10.1177/1049731504265838).
- Ferguson CJ, Brannick MT. 2012.** Publication bias in psychological science: prevalence, methods for identifying and controlling, and implications for the use of meta-analyses. *Psychological Methods* **17**:120–128 [DOI 10.1037/a0024445](https://doi.org/10.1037/a0024445).
- Ferguson CJ, Heene M. 2012.** A vast graveyard of undead theories: publication bias and psychological science's aversion to the null. *Perspectives on Psychological Science* **7**:555–561 [DOI 10.1177/1745691612459059](https://doi.org/10.1177/1745691612459059).

- Francis G. 2012a.** Evidence that publication bias contaminated studies relating social class and unethical behavior. *Proceedings of the National Academy of Sciences of the United States of America* **109**:E1587 DOI [10.1073/pnas.1203591109](https://doi.org/10.1073/pnas.1203591109).
- Francis G. 2012b.** The psychology of replication and replication in psychology. *Perspectives on Psychological Science* **7**:585–594 DOI [10.1177/1745691612459520](https://doi.org/10.1177/1745691612459520).
- Francis G. 2012c.** Publication bias and the failure of replication in experimental psychology. *Psychonomic Bulletin & Review* **19**:975–991 DOI [10.3758/s13423-012-0322-y](https://doi.org/10.3758/s13423-012-0322-y).
- Francis G. 2012d.** The same old new look: publication bias in a study of wishful seeing. *i-Perception* **3**:176–178 DOI [10.1068/i0519ic](https://doi.org/10.1068/i0519ic).
- Francis G. 2012e.** Too good to be true: publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review* **19**:151–156 DOI [10.3758/s13423-012-0227-9](https://doi.org/10.3758/s13423-012-0227-9).
- Francis G. 2013.** Publication Bias in ‘Red, Rank, and Romance in Women Viewing Men,’ by Elliot et al. (2010). *Journal of Experimental Psychology: General* **142**:292–296 DOI [10.1037/a0027923](https://doi.org/10.1037/a0027923).
- Franco A, Malhotra N, Simonovits G. 2014.** Publication bias in the social sciences: unlocking the file drawer. *Science* **345**(6203):1502–1505 DOI [10.1126/science.1255484](https://doi.org/10.1126/science.1255484).
- Gelman A. 2013.** Too good to be true. Slate. Available at http://www.slate.com/articles/health_and_science/science/2013/07/statistics_and_psychology_multiple_comparisons_give_spurious_results.html (accessed on 1 June 2017).
- Gerber AS, Malhotra N. 2008a.** Do statistical reporting standards affect what is published? publication bias in two leading political science journals. *Quarterly Journal of Political Science* **3**:313–326 DOI [10.1561/100.00008024](https://doi.org/10.1561/100.00008024).
- Gerber AS, Malhotra N. 2008b.** Publication bias in empirical sociological research. *Sociological Methods & Research* **37**:3–30 DOI [10.1177/0049124108318973](https://doi.org/10.1177/0049124108318973).
- Godlee F. 2012.** Research misconduct is widespread and harms patients. *BMJ* **344**:e14 DOI [10.1136/bmj.e14](https://doi.org/10.1136/bmj.e14).
- Greenland S, Pearl J, Robins JM. 1999.** Causal diagrams for epidemiologic research. *Epidemiology* **10**:37–48 DOI [10.1097/00001648-199901000-00008](https://doi.org/10.1097/00001648-199901000-00008).
- Hayashino Y, Noguchi Y, Fukui T. 2005.** Systematic evaluation and comparison of statistical tests for publication bias. *Journal of Epidemiology* **15**:235–243 DOI [10.2188/jea.15.235](https://doi.org/10.2188/jea.15.235).
- Higgins JPT, Green S. 2008.** *Cochrane handbook for systematic reviews of interventions*. Chichester: Wiley Online Library.
- Higgins JPT, Thompson SG. 2002.** Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine* **21**:1539–1558 DOI [10.1002/sim.1186](https://doi.org/10.1002/sim.1186).
- Ioannidis JP. 1998.** Effect of the statistical significance of results on the time to completion and publication of randomized efficacy trials. *Journal of the American Medical Association* **279**:281–286 DOI [10.1001/jama.279.4.281](https://doi.org/10.1001/jama.279.4.281).
- Ioannidis JPA. 2013.** Clarifications on the application and interpretation of the test for excess significance and its extensions. *Journal of Mathematical Psychology* **57**:184–187 DOI [10.1016/j.jmp.2013.03.002](https://doi.org/10.1016/j.jmp.2013.03.002).

- Ioannidis JPA, Stanley TD, Doucouliagos H. 2017.** The power of bias in economics research. *The Economic Journal* **127**:F236–F265 DOI [10.1111/ecoj.12461](https://doi.org/10.1111/ecoj.12461).
- Ioannidis JPA, Trikalinos TA. 2007.** An exploratory test for an excess of significant findings. *Clinical Trials* **4**:245–253 DOI [10.1177/1740774507079441](https://doi.org/10.1177/1740774507079441).
- Jefferson T, Jones MA, Doshi P, Del Mar CB, Heneghan CJ, Hama R, Thompson MJ. 2012.** Neuraminidase inhibitors for preventing and treating influenza in healthy adults. *Cochrane Database of Systematic Reviews* Epub ahead of print Jan 18 2012 DOI [10.1002/14651858.CD008965.pub3](https://doi.org/10.1002/14651858.CD008965.pub3).
- John LK, Loewenstein G, Prelec D. 2012.** Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science* **23**:524–532 DOI [10.1177/0956797611430953](https://doi.org/10.1177/0956797611430953).
- Kerr NL. 1998.** HARKing: hypothesizing after the results are known. *Personality and Social Psychology Review* **2**:196–217 DOI [10.1207/s15327957pspr0203_4](https://doi.org/10.1207/s15327957pspr0203_4).
- Kicinski M. 2014.** How does under-reporting of negative and inconclusive results affect the false-positive rate in meta-analysis? a simulation study. *BMJ Open* **4**:1–8 DOI [10.1136/bmjopen-2014-004831](https://doi.org/10.1136/bmjopen-2014-004831).
- Kreuter F, Presser S, Tourangeau R. 2008.** Social desirability bias in CATI, IVR, and web surveys: the effects of mode and question sensitivity. *Public Opinion Quarterly* **72**:847–865 DOI [10.1093/poq/nfn063](https://doi.org/10.1093/poq/nfn063).
- Labovitz S. 1972.** Statistical usage in sociology: sacred cows and ritual. *Sociological Methods & Research* **1**:13–37 DOI [10.1177/004912417200100102](https://doi.org/10.1177/004912417200100102).
- Lau J, Ioannidis JPA, Terrin N, Schmid CH, Olkin I. 2006.** The case of the misleading funnel plot. *BMJ* **333**:597–600 DOI [10.1136/bmj.333.7568.597](https://doi.org/10.1136/bmj.333.7568.597).
- Macaskill P, Walter SD, Irwig L. 2001.** A comparison of methods to detect publication bias in meta-analysis. *Statistics in Medicine* **20**:641–654 DOI [10.1002/sim.698](https://doi.org/10.1002/sim.698).
- Mahoney MJ. 1977.** Publication prejudices: an experimental study of confirmatory bias in the peer review system. *Cognitive Therapy and Research* **1**:161–175 DOI [10.1007/BF01173636](https://doi.org/10.1007/BF01173636).
- Mallett S, Clarke M. 2002.** The typical cochrane review. How many trials? how many participants? *International Journal of Technology Assessment in Health Care* **18**:820–823 DOI [10.1017/S0266462302000624](https://doi.org/10.1017/S0266462302000624).
- Merton RK. 1957.** Priorities in scientific discovery: a chapter in the sociology of science. *American Sociological Review* **22**:635–659 DOI [10.2307/2089193](https://doi.org/10.2307/2089193).
- Mooney CZ. 1997.** *Monte Carlo simulation*. Thousand Oakes: Sage Publications.
- Moreno SG, Sutton AJ, Ades AE, Stanley TD, Abrams KR, Peters JL, Cooper NJ. 2009.** Assessment of regression-based methods to adjust for publication bias through a comprehensive simulation study. *BMC Medical Research Methodology* **9**:2 DOI [10.1186/1471-2288-9-2](https://doi.org/10.1186/1471-2288-9-2).
- Morey RD. 2013.** The consistency test does not—and cannot—deliver what is advertised: a comment on Francis (2013). *Journal of Mathematical Psychology* **57**:180–183 DOI [10.1016/j.jmp.2013.03.004](https://doi.org/10.1016/j.jmp.2013.03.004).
- Nature. 2017.** Getting published in nature: the editorial process. Available at http://www.nature.com/nature/authors/get_published/ (accessed on 1 May 2017).

- Nuzzo R.** 2014. Scientific method: statistical errors. *Nature* **506**:150–152 DOI 10.1038/506150a.
- Paldam M.** 2013. Regression costs fall, mining ratios rise, publication bias looms, and techniques get fancier: reflections on some trends in empirical macroeconomics. *Econ Journal Watch* **10**:136–156.
- Paldam M.** 2015. Simulating an empirical paper by the rational economist. *Empirical Economics* **50**:1–25 DOI 10.1007/s00181-015-0971-6.
- Popper KR.** 1968. *Conjectures and refutations: the growth of scientific knowledge*. New York: Harper & Row.
- Renkewitz F, Keiner M.** 2016. How to detect publication biases from published data? In: *A Monte Carlo simulation of different methods. 50 Kongress der Deutschen Gesellschaft für Psychologie*. Leipzig: Pabst Science Publishers.
- Rosenthal R.** 1979. The file drawer problem and tolerance for null results. *Psychological Bulletin* **86**:638–641 DOI 10.1037/0033-2909.86.3.638.
- Schimmack U.** 2012. The ironic effect of significant results on the credibility of multiple-study articles. *Psychological Methods* **17**:551–566 DOI 10.1037/a0029487.
- Schwarzer G, Antes G, Schumacher M.** 2002. Inflation of type I error rate in two statistical tests for the detection of publication bias in meta-analyses with binary outcomes. *Statistics in Medicine* **21**:2465–2477 DOI 10.1002/sim.1224.
- Science.** 2017. The science contributors FAQ. Available at http://www.sciencemag.org/site/feature/contribinfo/faq/#pct_faq (accessed on 1 May 2017).
- Simmons JP, Nelson LD, Simonsohn U.** 2011. False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science* **22**:1359–1366 DOI 10.1177/0956797611417632.
- Simonsohn U.** 2013. It really just does not follow, comments on Francis (2013). *Journal of Mathematical Psychology* **57**:174–176 DOI 10.1016/j.jmp.2013.03.006.
- Simonsohn U, Nelson LD, Simmons JP.** 2014a. P-curve and effect size: correcting for publication bias using only significant results. *Perspectives on Psychological Science* **9**:666–681 DOI 10.1177/1745691614553988.
- Simonsohn U, Nelson LD, Simmons JP.** 2014b. P-curve: a key to the file drawer. *Journal of Experimental Psychology-General* **143**:534–547 DOI 10.1037/a0033242.
- Simonsohn U, Simmons JP, Nelson LD.** 2015. Better P-curves: making p-curve analysis more robust to errors, fraud, and ambitious P-Hacking, A Reply To Ulrich and Miller (2015). *Journal of Experimental Psychology-General* **144**:1146–1152 DOI 10.1037/xge0000104.
- Stanley TD, Doucouliagos H.** 2014. Meta-regression approximations to reduce publication selection bias. *Research Synthesis Methods* **5**:60–78 DOI 10.1002/jrsm.1095.
- Sterne JAC, Gavaghan D, Egger M.** 2000. Publication and related bias in meta-analysis: power of statistical tests and prevalence in the literature. *Journal of Clinical Epidemiology* **53**:1119–1129 DOI 10.1016/S0895-4356(00)00242-0.
- Sterne JAC, Sutton AJ, Ioannidis JPA, Terrin N, Jones DR, Lau J, Carpenter J, Rücker G, Harbord RM, Schmid CH, Tetzlaff J, Deeks JJ, Peters J, Macaskill P, Schwarzer G, Duval S, Altman DG, Moher D, Higgins JPT.** 2011. Recommendations for

examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *BMJ* **343**: DOI [10.1136/bmj.d4002](https://doi.org/10.1136/bmj.d4002).

Terrin N, Schmid CH, Lau J, Olkin I. 2003. Adjusting for publication bias in the presence of heterogeneity. *Statistics in Medicine* **22**:2113–2126 DOI [10.1002/sim.1461](https://doi.org/10.1002/sim.1461).

Thaler R. 1980. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization* **1**:39–60 DOI [10.1016/0167-2681\(80\)90051-7](https://doi.org/10.1016/0167-2681(80)90051-7).

Ulrich R, Miller J. 2017. Some properties of p-curves, with an application to gradual publication bias. *Psychological Methods* Epub ahead of print Apr 2017 DOI [10.1037/met0000125](https://doi.org/10.1037/met0000125).

Van Aert RCM, Wicherts JM, Van Assen MALM. 2016. Conducting meta-analyses based on *p* values. *Perspectives on Psychological Science* **11**:713–729 DOI [10.1177/1745691616650874](https://doi.org/10.1177/1745691616650874).

Van Assen MALM, Van Aert RCM, Wicherts JM. 2015. Meta-analysis using effect size distributions of only statistically significant studies. *Psychological Methods* **20**:293–309 DOI [10.1037/met0000025](https://doi.org/10.1037/met0000025).

Wasserstein RL, Lazar NA. 2016. The ASA's statement on *p*-values: context, process, and purpose. *American Statistician* **70**:129–131 DOI [10.1080/00031305.2016.1154108](https://doi.org/10.1080/00031305.2016.1154108).

Yoder S, Bramlett BH. 2011. What happens at the journal office stays at the journal office: assessing journal transparency and record-keeping practices. *Ps-Political Science & Politics* **44**:363–373 DOI [10.1017/S1049096511000217](https://doi.org/10.1017/S1049096511000217).