How to detect publication bias in psychological research?

A comparative evaluation of six statistical methods

Frank Renkewitz & Melanie Keiner

University of Erfurt

Please address correspondence to:

Frank Renkewitz, Department of Psychology, University of Erfurt

Nordhäuser Str, 63, 99089 Erfurt, Germany

Phone: + 49 361 737 2223

Email: frank.renkewitz@uni-erfurt.de

Author note:

We thank Johannes Hönekopp, Marcel van Assen and Robbie van Aert for their helpful comments on a previous version of this manuscript.

Abstract

Publication biases and questionable research practices are assumed to be two of the main causes of low replication rates. Both of these problems lead to severely inflated effect size estimates in meta-analyses. Methodologists have proposed a number of statistical tools to detect such bias in meta-analytic results. We present an evaluation of the performance of six of these tools. To assess the Type I error rate and the statistical power of these methods, we simulated a large variety of literatures that differed with regard to true effect size, heterogeneity, number of available primary studies, and sample sizes of these primary studies; furthermore, simulated studies were subjected to different degrees of publication bias. Our results show that across all simulated conditions, no method consistently outperformed the others. Additionally, all methods performed poorly when true effect sizes were heterogeneous or primary studies had a small chance of being published, irrespective of their results. This suggests that in many actual meta-analyses in psychology, bias will remain undiscovered no matter which detection method is used.

Keywords: meta-analysis, publication bias, optional stopping, bias detection, heterogeneity.

How to detect publication bias in psychological research?

A comparative evaluation of six statistical methods

The observation that many psychological results are not replicable (Open Science Collaboration, 2015) has caused a growing sense of unease about the reliability of research findings in the field of psychology in general. Failed replications, suggesting that even very prominent effects that are not only widely investigated but repeatedly discussed and cited, are actually non-existent (e.g. Wagenmakers et al., 2016), have further intensified this loss of confidence. If even seminal research findings cannot be replicated, this raises the question as to which findings one can trust at all. Thus, on the one hand, the current situation requires a change in research practices that is suitable to enhance the replicability of psychological findings. On the other hand, it calls for a re-evaluation of already published evidence in order to separate non-replicable results from those that are replicable and thus more trustworthy. Such a re-evaluation needs to be informed by the sources from which replication problems stem.

From a methodological point of view, two central and interrelated causes of the replication crisis are publication biases that favor significant over non-significant results, and questionable research practices (QRPs) that serve to increase the chances of obtaining a publishable significant result even when the null hypothesis is true (i.e., *p*-hacking; Simmons, Nelson, & Simonsohn, 2011). Such *p*-hacking practices involve repeatedly trying different ways to analyze the data (e.g., including or excluding outliers, controlling for different covariates, or collecting additional participants after observing negative test outcomes) until a significant result is found. Both of these problems not only increase the proportion of false positives in the literature but also bring about inflated meta-analytic estimates of true effect sizes and, hence, threaten the validity of meta-analytic results in general. It follows that

statistical methods to detect and correct biases in effect size estimates have been primarily developed out of the literature on meta-analytic methods. In recent years and with growing awareness of the replication crisis, the inventory of such statistical tools has grown rapidly and already extant methods have been refined. As a consequence, the psychological scientist now faces the added challenge of selecting the best possible method from this wider inventory, given the characteristics of the available data.

A large number of simulation studies have evaluated the performance of statistical methods for the detection and correction of biases in meta-analyses (e.g., Sterne, Gavaghan, & Egger, 2000; Macaskill, Walter, & Irwig, 2001; Kromrey & Rendina-Gobioff, 2006; Rücker, Carpenter, & Schwarzer, 2011; Francis, 2013; Simonsohn, Nelson, & Simmons, 2014; Stanley & Doucouliagos, 2014; van Assen, van Aert, & Wicherts, 2015; McShane, Böckenholt, & Hansen, 2016; van Aert, Wicherts & van Assen, 2016; Pustejovsky & Rodgers, 2018). However, these studies included only subsets of the methods that are currently available. Furthermore, their results are difficult or impossible to compare as they diverge with respect to the conditions they investigate: Different studies implement different forms or degrees of publication bias, vary in true effect size, simulate different degrees of heterogeneity in true effect sizes, draw sample sizes of primary studies from different distributions, or differ in the effect size measure they consider. As a consequence, a comprehensive, comparative evaluation of statistical methods for the detection and correction of biases under conditions that are typical for psychological research was missing until recently. This has changed with a paper by Carter and colleagues (2019) that provides an extensive evaluation of the corrected effect size estimates of several different methods. The simulation study we present is conceptually similar to the work of Carter et al. (2019) and, with regard to the correction of biases, corroborated several of their central conclusions. In

this paper, however, we focus on the performance of a large set of statistical tools in detecting biases – an issue that has not been thoroughly investigated so far.

There has been some debate over whether statistical tests for the detection of publication bias should be employed at all. Simonsohn (2013; see also Morey, 2013) argues that such tests are useless as the publication process is known to be biased in general. Thus, the answer to the research question “is there publication bias in this set of evidence?” is known even before a statistical test for it is proposed. Accordingly, Simonsohn, Nelson and Simmons (2014) developed a statistical method for the correction of publication biases in effect size estimates (p -curve), but did not use the rationale of this method to propose a test for publication bias (even though this could be easily done; see van Assen et al., 2015, and their tool p -uniform).

Simonsohn’s argument and the conclusion derived from it, however, are questionable for several reasons. First, there are examples of meta-analytic data collections that comprise (almost) exclusively statistically non-significant findings (e.g., Bürkner et al, 2017; Xie, Zhou, & Liu, 2018) and, hence, can at the very least not be biased in the usual sense of a selection that favors studies rejecting the null hypothesis. That said, in a recent meta-study on 83 meta-analyses published in *Psychological Bulletin*, van Aert, Wicherts and van Assen (2019) observed that only 29% of the included primary effect sizes were statistically significant. This finding points to the possibility that unbiased data collections in psychological meta-analyses may represent more than an exception to the rule.

More importantly, even if the publication process was biased in general, this would not imply that all meta-analyses are affected by this bias. For instance, if the power of the original studies is large (due to a large true effect size and sufficiently large sample sizes), publication bias favoring statistically significant findings will censor very few results. In this case, the bias in the publication process will also add little bias to the meta-analytic effect

size estimate (Ulrich, Miller & Erdfelder, 2018). Of course, tests for publication bias are meant to inform us about systematic distortions in the available evidence. Thus, if publication bias does not give rise to such distortions, it is of little relevance.

Finally, it appears to be widespread practice in meta-analyses to test for publication bias and to assess the validity of the resulting effect size estimate based on the outcome of this test (for an overview of methods used to detect and correct biases in 61 articles published between 1990 and 2013, see van Erp et al., 2017). This practice is in line with recommendations from the Meta-Analysis Reporting Standards (MARS) that request an “assessment of bias including possible data censoring” (APA, 2010, p. 252). Obviously, this assessment should use the best available method: That is, a test that guarantees a proper control of the Type I error rate while maximizing power.

The central objective of our simulation study is to inform the choice of a test for publication bias in psychological meta-analyses. To this end, we compared the performance of different statistical tools under conditions relevant for psychological research. Additionally, by quantifying the power of the statistical tools under a large range of conditions, we aimed to identify research situations in which tests for publication bias can be applied meaningfully and to differentiate them from situations in which any attempt to detect biases will remain a hopeless endeavor. Finally, we investigated how the performance of the statistical tools is affected by factors that vary in psychological meta-analyses (e.g., the mean and the heterogeneity of true effect sizes).

Methods for the Detection of Publication Bias

In the following, we give a brief overview of the statistical tools that we evaluated. These tools and their implementation are described in more detail in the supplemental material.

The tests for publication bias included in this study can be loosely grouped according to the information they use to infer bias. Several methods are based on an assessment of asymmetry in the distribution of observed effect sizes as displayed in a funnel plot (a scatter plot of effect size estimates in primary studies against the standard error of these estimates; see Figure S1 in the supplemental material). Censoring of non-significant results brings about asymmetry in a funnel plot as studies combining rather large standard errors (i.e., small sample sizes) with small or moderate effect sizes will be missing. Begg's rank correlation (Begg & Mazumdar, 1994) and several regression methods use the resulting association between observed effect sizes and their standard errors to gauge this asymmetry. More specifically, Begg's rank correlation (Begg & Mazumdar, 1994) uses Kendall's tau to measure the correlation between (standardized) effect sizes and their variances (the square of the standard errors). Egger's regression (or, equivalently, the precision-effect test (PET); Stanley & Doucouliagos, 2014) is a weighted regression of effects sizes on their standard errors (Egger et al., 1997). A significant regression slope indicates bias. We also included the precision-effect estimate with standard error (PEESE; Stanley & Doucouliagos, 2014), a variant of PET that has been suggested as a means to improve corrected meta-analytic effect size estimates in the presence of bias. With regard to the detection of publication bias, however, both methods yield identical results; therefore, we report only about the performance of PET. Furthermore, we incorporated trim-and-fill (Duval & Tweedie, 2000), a method that assesses asymmetry in the univariate distribution of effect sizes with the help of non-parametric measures. Based on these measures, trim-and-fill estimates the number of excluded studies and diagnoses bias when this number exceeds the significance threshold. It is worth noting that the method relies on an explicit model of the publication process: It assumes that the j studies with the smallest observed effect sizes will be censored. In contrast, the selection mechanisms implemented here simulate a preferential publication of statistically

significant results. That is, we simulate selection processes that are not based on observed effect sizes but their p -values. Thus, with regard to trim-and-fill, our study can be considered as a robustness check when the assumptions of the selection model of the method are not met.

A second group of methods operates on p -values. P -curve (Simonsohn et al., 2014) and p -uniform (van Assen et al., 2015) consider p -values of significant studies only. Both methods rely on a simple fact about the distribution of significant p -values from a set of studies investigating the same hypothesis: When there actually is a true effect in the predicted direction, small p -values close to zero will be more likely than larger p -values close to .05. Thus, a true effect causes a right skew in the distribution of p -values. The degree of right skew depends on the power of the studies, which in turn depends on their sample sizes and the true effect size. P -curve and p -uniform use this relationship to estimate the true effect size from the distribution of significant p -values. Here, we focus on p -uniform, as only this method provides a formalized test for publication bias. This test compares the observed distribution of significant p -values with the distribution implied by the effect size estimate from a fixed-effect meta-analysis.

In contrast to p -uniform, the test of insufficient variance (TIVA; Schimmack, 2014) uses all available p -values as input. These p -values are converted into z -scores of the standard normal distribution. When studies are published irrespective of their results, these z -scores can generally be expected to have a variance of at least 1. A variance larger than 1 can and will occur when the underlying true effect sizes of the studies vary or sample sizes differ. However, a censoring of non-significant results will cause a restriction in the variance of observed p -values and, hence, in the variance of converted z -scores. Accordingly, TIVA infers bias when the variance in z -scores is significantly smaller than 1.

The last method considered in our study does not fit in any of the aforementioned groups: The test of excess significance (TES; Ioannidis & Trikalinos, 2007; Francis, 2013)

compares the average power in a set of studies to the proportion of significant findings in this set. The power of individual studies is calculated based on the effect size estimate resulting from a fixed-effect meta-analysis of all included studies. TES diagnoses bias when the observed number of significant studies is significantly larger than the number expected according to the average power.

Heterogeneity and Bias Detection Methods

Before we turn to our evaluation of the different detection methods, it is worthwhile to briefly discuss a characteristic that most of these methods share: They were designed to address the problem of publication bias in meta-analyses, in which all included primary studies have the same underlying true effect size and differ in their observed effect sizes only because of sampling error (van Assen, van Aert, Wicherts, 2015; Egger et al., 1997; Ioannidis & Trikalinos, 2007; Begg & Mazumdar, 1994; Duval & Tweedie, 2000).¹ According to two recent surveys (Stanley, Carter & Doucouliagos, 2018; van Erp et al., 2017), however, between-study heterogeneity is the rule, rather than the exception, in actual psychological meta-analyses. In most meta-analyses the true effect sizes of primary studies appear to vary which may be due to, for example, different operationalizations of the independent variable, different measurements or different participant populations.

For all of the detection methods included here, it can either be deduced from their rationale, or has been shown in previous simulation studies, that their performance will be affected by heterogeneity. This is most easily demonstrated for TIVA: Given that variance in true effect sizes will increase the variance in observed effect sizes, it follows that the conversion of observed p -values into z -scores will, in turn, translate to an increased variance of z -scores. To illustrate this with numbers, we appeal to an example from our study: When

¹ The exception here is trim-and-fill that was meant to address the problem of publication bias also with heterogeneous effect sizes. However, previous simulation studies (Terrin, Schmid, Lau & Olkin, 2003; van Assen, van Aert & Wicherts, 2015) showed that the performance of trim-and-fill in correcting biases is adversely affected by heterogeneity when primary studies are selected for publication based on their p -values.

we simulated no publication bias and no heterogeneity, we observed an average variance in z-scores that corresponded almost exactly to the expected value of 1 (aggregated across all other conditions of the design, see below). However, when there was a rather moderate degree of heterogeneity ($\tau = 0.3$), the average variance of z-scores was already approximately 2.6. Thus, publication bias is less likely to induce a reduction in variance below the expected value of 1 when there is heterogeneity.

For *p*-uniform and the regression-based method PET, several simulation studies (e.g., van Aert, Wicherts & van Assen, 2016; Stanley, 2017) demonstrated that, under publication bias and with larger heterogeneity, these tools yield increasingly inflated effect size estimates that more closely resemble the (conventional) meta-analytic effect size estimate. In both cases, this also implies that the signal these methods use to infer the presence of bias (i.e. the inconsistency of the distribution of significant *p*-values with the meta-analytic effect size estimate for *p*-uniform, and the relationship between observed effect sizes and their standard errors for PET) is weakened by heterogeneity. Thus, for *p*-uniform, PET and Begg's rank correlation (which also relies on the association between observed effect sizes and their standard errors), it has to be expected that these methods lose power under heterogeneity.

Finally, the performance of TES is affected by heterogeneity because its estimation of the power of primary studies relies on the assumption that all of these studies share the same true effect size (Ioannidis & Trikalinos, 2007; Francis, 2013). When there actually is heterogeneity, the resulting estimate of the average power of the primary studies is systematically distorted (Yuan & Maxwell, 2005; this issue will be discussed in more detail below). As a consequence, the expected number of significant results among primary studies will no longer correspond to the observed number of significant results even when there is no bias.

Taken together, these considerations suggest that heterogeneity is an obstacle to the detection of publication bias no matter which of the methods considered here is used. Thus, from a more practical perspective, the choice of a detection method – as well as the interpretation of its results – needs to be informed about its (relative) robustness against heterogeneity. We, therefore, incorporate heterogeneity in the design of our study and aim to provide information pertaining to this problem.

Method

To evaluate the six different detection methods outlined above, we performed a Monte Carlo simulation in which we generated primary studies that were subjected to five different forms of selection in the process of publication (the R code for this simulation is given in the supplemental material and is available for download at osf.io/8yefd/). All simulated primary studies had a two-group experimental design, but varied with regard to the (mean of the) underlying true effect sizes (δ) and the heterogeneity of these true effect sizes (τ ; i.e., the standard deviation of true effect sizes). Furthermore, sample sizes of the primary studies were drawn from different distributions (see Table 1 for an overview of all factors that were varied in the simulation). Sets of k studies were then meta-analyzed, using a random-effects model (with the DerSimonian-Laird estimator for the between-study variance τ^2 ; DerSimonian & Laird, 1986).

The different selection mechanisms that were realized affected what proportion of simulated studies was published. Throughout, only published studies fed into meta-analyses, but in all selection conditions primary studies were simulated until the target number of k studies was available for meta-analysis. Four of the selection mechanisms represented different degrees of publication bias. These mechanisms were fully crossed with the remaining four factors given in Table 1, which resulted in 2,400 unique conditions (4

selection mechanisms $\times$ 4 true effect sizes $\times$ 6 degrees of heterogeneity $\times$ 5 sample size distributions $\times$ 5 numbers of studies per meta-analysis). In the fifth selection condition, we combined complete censorship of non-significant results with additional optional stopping, a variant of p -hacking. We simulated two different intensities of optional stopping, but the (initial) sample size of primary studies with both intensities was always set to $n = 20$ per cell (for more information on the implementation of optional stopping, see below). Thus, this fifth selection mechanism was not crossed with the factor affecting sample sizes. A combination of the two intensities of optional stopping with the remaining three factors accounted for 240 additional unique conditions of the simulation design (2 intensities of optional stopping $\times$ 4 true effect sizes $\times$ 6 degrees of heterogeneity $\times$ 5 numbers of studies per meta-analysis).

Table 1: Simulation Design

Experimental factor	Levels
True effect size (δ)	0, 0.2, 0.5, 0.8
Heterogeneity (τ)	0, 0.1, 0.2, 0.3, 0.4, 0.5
Distribution from which sample sizes were drawn (n per cell)	$U(20, 30)$, $U(10, 40)$, $U(45, 55)$, $U(35, 65)$, $U(20, 80)$
Number of studies per meta-analysis (k)	5, 7, 10, 30, 50
Selection mechanism	100% bias, two-step bias, 90% bias, no bias, Optional stopping (increment: 1 or 5) ¹

¹ Optional stopping always started with $n = 20$. Consequently, the two different increments in optional stopping were not crossed with the factor “distribution from which sample sizes were drawn”.

Selection Mechanisms

In the selection condition with no bias, all simulated original studies were published. Hence, this condition served as a control condition which we used to estimate the Type I error rates of the different detection methods. In contrast, in the condition with 100% bias, all simulated non-significant results (with $\alpha = .05$, one-tailed) were excluded from publication.

With two further selection mechanisms we implemented a less strict censorship of non-significant findings: In the two-step bias condition, “marginally significant” studies with p -values between $p = .05$ and $p = .10$ had a 20% chance of being published; studies with $p > .10$ remained unpublished. In the condition with 90% bias, non-significant results were censored with a probability of 90%. In other words, 10% of the simulated primary studies in this condition were published irrespective of their findings.

Finally, in the condition involving not only publication bias but also optional stopping, an initial significance test was performed in all primary studies with a sample size of $n = 20$ per cell. If the result was non-significant, the sample size was increased in either increments of 1 or 5. After each increment, the test for significance was repeated. The procedure was stopped either when a significant result was obtained (in which case the study was published) or when a sample size of $n = 40$ was reached. Studies that did not find a significant result remained unpublished. By varying the increment in sample size after negative test results, we simulated two different intensities of optional stopping. Both these intensities, however, represent an excessive use of this variant of p -hacking as additional data are collected in all studies that initially find a non-significant result.

Mean True Effect Sizes and Degrees of Between-Study Heterogeneity

The four levels of the factor “mean true effect size (δ)” were chosen according to Cohen’s conventions (Cohen, 1988). Heterogeneity was varied in six levels between $\tau = 0$ and $\tau = 0.5$. These levels cover the range of heterogeneity typically observed in psychological meta-analyses (van Erp et al., 2017). In terms of the heterogeneity metric I^2 (i.e., the ratio of variance in true effect sizes to variance in observed effect sizes), Higgins and colleagues (2003) proposed to describe values of 0.25, 0.5 and 0.75 as low, medium and high degrees of heterogeneity. I^2 depends not only on τ but also on the sample sizes of the primary studies. Within our simulation, random-effects meta-analyses of unbiased study sets yielded average

estimates of I^2 of 0.14, 0.26, 0.42, 0.56 and 0.66 for the τ -values between 0.1 and 0.5, when the mean sample size is $n = 25$ (no publication bias, aggregated over δ , k and the two conditions with a mean sample size of $n = 25$). In the three conditions with a mean sample size of $n = 50$, the respective I^2 estimates are 0.18, 0.40, 0.59, 0.72 and 0.80.

Sample Sizes of Primary Studies

Sample sizes (n) for experimental and control groups were identical for each individual primary study and drawn from five different uniform distributions that varied in mean and range. Larger mean sample sizes will obviously increase the proportion of significant primary studies in all conditions in which the true effect size δ is larger than zero. Consequently, they will also reduce the proportion of studies that is excluded from publication and may, therefore, affect the performance of bias detection methods. The range of the sample sizes of the studies included in a meta-analysis is of direct importance for Begg's rank correlation and PET, as these methods assess the covariation between effect sizes and their standard errors. Thus, if there is little variation in sample size (and, therefore, in standard error) the methods are likely to fail. The lower and upper limits of the different uniform distributions in our design roughly correspond to lower and upper quartiles of per-group sample size in four psychological journals as reported by Marszalek and colleagues (2011), e.g., *Journal of Abnormal Psychology*: $Q_1 = 10$ and $Q_3 = 27$ in 1977; *Journal of Applied Psychology*: $Q_1 = 21$ and $Q_3 = 83$ in 2006.

Number of Studies per Meta-Analysis

In choosing the levels of this factor, we partly focused on small study sets ($k = 5, 7$ or 10) as bias detection methods have been applied in rather small samples (e.g., Francis, 2012b; 2013), and because we aimed to assess the minimum number of studies necessary for the different methods to achieve acceptable power. According to data collected by van Erp et al.

(2017), the median number of studies in psychological meta-analyses is 12. The larger sizes of study sets in our design ($k = 30$ and 50) roughly correspond to the mean number of studies in psychological meta-analyses and to the 80th percentile.

Procedure

Observations in the control group of each individual study were randomly drawn from a normal distribution with $\mu = 0$ and $\sigma = 1$. Observations in the experimental group were also generally generated from a normal distribution with $\sigma = 1$. In conditions in which there was no heterogeneity ($\tau = 0$), the mean of this normal distribution was simply equal to the true effect size δ in the given condition. In conditions with heterogeneity, the population mean and, hence, the true effect size for each individual study, was drawn in a preceding step from a normal distribution with $\mu = \delta$ and $\sigma = \tau$. For each simulated study, we computed Hedges' g (as an estimate for the true effect size δ), its variance and standard error (see the supplemental material for details), and determined a one-tailed p -value from an independent samples t -test. Thus, across conditions, studies that were selected for publication based on their p -values generally had a positive effect size.

In each of the 2,640 unique conditions of the complete design, we simulated 1,000 respective meta-analyses. Hence, the maximum standard error of the proportion of significant results, found by the different detection methods in each of these conditions, is $SE = 0.016$ (the standard error of an estimated proportion p is given by $SE_p = \sqrt{p(1-p)/n}$ and will take its maximum value when $p = .50$).

Results

An important reference point for the interpretation of the statistical power of the detection methods is, of course, the size of the bias that they aim to uncover. Before we turn

to the false positive and true positive rates of the detection methods, we therefore briefly describe the bias in meta-analytic effect size estimates that resulted from the different selection mechanisms.

Bias in Meta-Analytic Effect Size Estimates

Figure 1 A) shows the mean estimates of meta-analyses for different true effect sizes in all selection conditions when there is no heterogeneity (95% CIs for all means displayed in Figure 1 have a maximum width of 0.004). In general, when all non-significant studies are excluded from publication (i.e. 100% bias), the bias in effect size estimates is fully determined by the average power of the original studies.² Obviously, when there is no true effect ($\delta = 0$), 95% of all studies will be suppressed. Given the sample sizes of original studies in our simulation, the remaining 5% of studies had a weighted mean effect size of about $\bar{d} = 0.50$, thus giving rise to a very severe overestimation of the true effect size. When the true effect size was $\delta = 0.2$, the average power of simulated original studies was 22% and, consequently, the proportion of censored studies dropped to 78% (see Figure S2 for power values in all conditions). Meta-analyses of the 22% published studies yielded an average effect size estimate of $\bar{d} = 0.52$. Hence, the bias was still large but was also considerably smaller than in the condition with $\delta = 0$. With increasing true effect sizes (and increasing power) the bias decreased further. In the condition with $\delta = 0.8$, it was already almost nil as the average power in this condition exceeded 90% and only relatively few studies were excluded from publication.

Please insert Figure 1 about here.

² The data points in Figure 1 A) and B) result from an averaging across the factors 'number of studies per meta-analysis' and 'distribution of sample sizes'. Whereas the first of these factors did not affect the mean of meta-analytic estimates, sample sizes of primary studies had an impact: Due to the larger power of larger primary studies, the bias in meta-analytic estimates was smaller when the mean sample size was $n = 50$ than when it was $n = 25$. Qualitatively, however, the result pattern described in this section did not change across different distributions of sample sizes, and the effects of true effect size, heterogeneity and selection mechanism on the bias in meta-analytic estimates were stable in their direction and hardly affected in their size.

As Figure 1 A) illustrates, the functional relationship between true effect size and meta-analytic effect size estimates remains unaltered when there is additional optional stopping, or strong but incomplete censoring of marginally significant results. In the latter condition (two-step bias), the bias in effect size estimates was only slightly reduced. This small reduction had to be expected for two reasons: Across all true effect sizes, 92% of published studies in this condition were still significant by the conventional criterion of $\alpha = .05$ (see Figure S2; this proportion of significant results among published studies roughly corresponds to estimates of the proportion of significant results in the psychological literature, Fanelli, 2010; Kühberger, Fritz & Scherndl, 2014). Additionally, marginally significant results are be associated with smaller, but still inflated mean estimates of the true effect sizes.

In contrast, in the condition with additional optional stopping, the bias was slightly amplified. Given the parameter settings of our simulation, optional stopping brought about significant results with smaller sample sizes. The mean sample size of significant studies in the condition with complete publication bias (across all true effect sizes) was $n_1 = n_2 = 40.7$; in the condition with additional optional stopping, this sample size was $n_1 = n_2 = 24.0$. Significant results with smaller sample sizes, of course, have to be associated with larger observed effect sizes.

Finally, when 10% of studies were published irrespective of their results, the bias in meta-analytic effect size estimates was greatly reduced. Furthermore, the largest bias no longer occurred at $\delta = 0$ but at $\delta = 0.2$. These findings may appear surprising at first sight, but can be easily deduced: With $\delta = 0$, the proportion of significant results was 5%. Of the remaining 95% of studies, 10% will be published. Thus, even though 90% of studies are produced under publication bias, the proportion of significant results among published studies

was only $0.05 / (0.05 + 0.95 \times 0.10) = 34.5\%$. The large share of published non-significant results reduced the bias drastically ($\bar{d} = 0.15$). With $\delta = 0.2$, as reported above, the average power of simulated original studies was 22%. The respective significant studies had a weighted mean effect size of $\bar{d} = 0.52$ and, thus, still overestimated the true effect size severely. At the same time, the proportion of significant results among published studies increased to $0.22 / (0.22 + 0.78 \times 0.10) = 73.8\%$. As a consequence, significant studies had a stronger weight on mean estimates, and the resulting bias was larger than with $\delta = 0$ ($\bar{d} = 0.41$; a formal analysis of the relationship between different selection mechanisms and the resulting bias in effect size estimates is given in Ulrich, Miller & Erdfelder, 2018).

Figure 1 B) shows mean meta-analytic estimates for $\tau = 0.3$, illustrating the effect of heterogeneity on the bias in these estimates. We chose $\tau = 0.3$ for this figure (as well as for Figures 2 and 3) as two recent surveys suggest that the estimated degree of heterogeneity in psychological meta-analyses is typically of about this size (van Erp et al., 2017; Stanley, Carter & Doucouliagos, 2018). With growing heterogeneity, the bias in all selection conditions increases. Furthermore, the degree of bias in different selection conditions becomes more similar (see Figure S3 for mean meta-analytic estimates in all heterogeneity conditions). The general rise in bias is mostly driven by the increased spread of observed effect sizes that is caused by heterogeneity in true effect sizes: Smaller observed effect sizes remain censored due to the significance criterion, whereas larger observed effect sizes are published and, in turn, inflate meta-analytic estimates further.

To summarize, the largest and most relevant biases generally occurred with non-existent or small true effect sizes. With larger effect sizes (and sufficient power), however, the bias tended toward zero. Optional stopping inflated the bias marginally. In contrast, a relatively small proportion of studies that were not subject to censoring were sufficient to decrease the size of bias substantially. Finally, heterogeneity exacerbated biases in effect size

estimates resulting from all forms of preferential publication of significant findings. Thus, to avoid seriously wrong conclusions from meta-analytic results, successful bias detection is even more important in the presence of heterogeneity.

Type I Error Rate

Figure 2 shows the proportion of significant results for all detection methods when there is no publication bias. Under homogeneity (Figure 2 A), PET was best at maintaining the nominal α -level, whereas TIVA slightly exceeded it when $\delta = 0.8$ (which is exclusively due to an error rate of 8% with $k = 50$ studies). The false positive rate of Begg's rank correlation was at about 4% with all true effect sizes as the method very rarely indicated bias with $k = 5$ studies (see Figure S4 and S5). All other methods tended to yield conservative error rates. This tendency was especially pronounced when the number of studies per meta-analysis was small ($k \leq 10$) and the true effect size was large ($\delta \geq 0.5$): Under these circumstances, the Type I error rates of p -uniform, TES and trim-and-fill often dropped below 1%.

Please insert Figure 2 about here.

The performance of PET, Begg's rank correlation and trim-and-fill was hardly affected by heterogeneity (see Figure 2 B where $\tau = 0.3$). On the contrary, the false positive rates of TIVA and p -uniform strongly declined with increasing heterogeneity. None of these methods were associated with an appreciably inflated Type I error rate in any combination of true effect size, number of studies and heterogeneity considered here. The striking exception to this rule was TES.

The performance of TES requires some further explanation. When the true mean effect was $\delta = 0$ or $\delta = 0.2$, TES produced an increasingly inflated false positive rate with increasing heterogeneity (see Figure S4). However, with $\delta = 0.5$ or $\delta = 0.8$, its false positive

rate dropped quickly toward zero with increasing heterogeneity. As mentioned before, in TES power is estimated based on the assumption of a fixed-effect model. Thus, when there actually is heterogeneity in true effect sizes, the resulting estimates are systematically distorted. More specifically, when the power estimate of a fixed-effect model is below 50%, the actual power of studies affected by heterogeneity is larger than estimated (Kenny & Judd, 2018; Yuan & Maxwell 2005; this is also discernible by a comparison of the proportion of significant results in the no bias condition under homogeneity and heterogeneity in this simulation, see Figure S2). This implies that we observed more significant results than expected based on a fixed-effect model power analysis in primary studies with $\delta = 0$ and $\delta = 0.2$. From this difference between the observed and expected number of significant results in primary studies, TES frequently inferred bias even though all studies were properly run and published. In contrast, when the power estimate of a fixed-effect model is above 50%, the actual power of studies affected by heterogeneity is smaller than estimated. Thus, with $\delta = 0.5$ and $\delta = 0.8$, we observed less rejections of the null hypothesis than expected, and consequently, TES very rarely detected an excess of significant findings.³

With $\delta = 0$ and $\delta = 0.2$, severely inflated Type I error rates of TES occurred with as little as 10 studies when $\tau = 0.3$, and with 30 studies when $\tau = 0.2$ (see Figure S3). As a possible solution to this problem Francis (2013) suggested using a random-effects model to estimate power and to determine the expected number of significant primary studies according to this power estimation. With that said, this procedure is not only analytically

³ The false positive rates displayed in Figure 2 (and in Figures S4, S5 and S6 in the supplemental materials) are averaged across the factor ‘distribution of sample sizes’. In general, this factor did not have an impact worth discussing. With regard to TES, however, this assessment is not fully accurate: Assuming a fixed-effect model, $\delta = 0.2$ was associated with an average power below 50% with all distributions of sample sizes simulated. For $\delta = 0.5$ and $\delta = 0.8$, average power was above 50% independently of the sample sizes of primary studies. Hence, when $\tau \geq 0.2$ TES always yielded inflated false positive rates for $\delta \leq 0.2$, whereas its false positive rates were always reduced for $\delta \geq 0.5$. However, the difference between the actual power of a heterogeneous set of primary studies and the TES power estimate is not independent of sample sizes. As a consequence, the Type I error rate of TES was also affected by sample sizes. In particular, variations in mean sample size had an appreciable effect when $\delta = 0$ and $\tau \geq 0.2$. In this case, larger and more strongly inflated false positive rates occurred when the mean sample size was $n = 50$ than when it was $n = 25$.

much more complex, but we are not aware of any actual application of it. More importantly, power calculation for a random-effects model requires an estimate of τ . These estimates are notoriously imprecise, unless the set of studies is large (Viechtbauer, 2007). Additionally, estimates of τ can and will also be systematically distorted by publication bias (Augusteijn, van Aert & van Assen, 2019; Jackson, 2007; see also Figure S16). Given the ubiquity of heterogeneity in psychological meta-analyses (Stanley, Carter & Doucouliagos, 2018), we therefore warn against the use of TES in sets of more than 10 studies.

There is one additional, noteworthy restriction to the results depicted in Figure 2: For the evaluation of Begg's rank correlation and PET, we transformed the effect sizes of the primary studies in a way that stabilized their variance (Hedges, 1981; see the supplemental materials for details on this transformation). When these methods were applied to Hedges' g values directly, they produced inflated false alarm rates for true effect sizes larger than zero (see Figure S6). This can be explained as follows: Begg's rank correlation and PET rely on estimates of effect sizes and their standard errors in primary studies. Estimates of the standard error of Hedges' g depend on the estimate of g itself, in such a way that larger values of g^2 are associated with larger standard errors (Hedges, 1981). Thus, correlations between observed effect sizes and their standard errors are guaranteed to occur when the true effect size is different from zero. The resulting inflated Type I error rate for Begg's rank correlation and PET, with g (or r) as effect size measure, has been described before (e.g. Kromrey & Rendina-Gobioff, 2006; Pustejowski & Rodgers, 2018), but seems to be oftentimes neglected in actual applications of these methods. As a general rule, methods assessing the association between observed effect sizes and their standard errors should only be used after a variance-stabilizing transformation.

Power

Figure 3 shows the statistical power of each method as a function of selection mechanism, study number, true effect size (only levels $\delta = 0$ and $\delta = 0.5$ shown), and heterogeneity (only levels $\tau = 0$ and $\tau = 0.3$ shown). Power estimates for all simulated effect sizes and degrees of heterogeneity are reported in Figures S7 to S12. In all of these figures, we present power estimates that are averaged across the factor ‘distribution of sample sizes of primary studies’. Mostly, the effect of this factor on the performance of the different detection methods was small. There were, however, exceptions to this rule that will be described and discussed below.

We organize the description of the statistical power of the detection methods according to selection mechanisms and start out with the 100% bias condition. Results for other selection mechanisms will be described subsequently.

Please insert Figure 3 about here.

100% bias. As expected, the statistical power of all methods increased with the number of available studies (k) and with smaller true effect sizes (i.e., with larger biases in mean estimates). Under homogeneity ($\tau = 0$), TIVA and p -uniform outperformed the remaining methods when the true effect size was $\delta = 0$. The advantage of these methods was especially pronounced when the number of studies was small ($k \leq 10$): TIVA and p -uniform achieved a power of more than 70% with as little as $k = 5$ studies. Thus, with these methods there is a very reasonable chance to detect the largest and most relevant biases even in very small study sets. With larger true effect sizes, the performance of all methods declined but their power diminished at different rates. Most importantly, when $\delta > 0$, TIVA achieved higher power than p -uniform. With effect sizes $\delta \geq 0.5$, however, and study numbers $k \geq 30$, TIVA was outperformed by TES. With that said, and as was mentioned above, in study sets

of this size (i.e. $k > 10$), TES should only be used when heterogeneity can be excluded (due to its inflated Type I error rate with even moderate degrees of heterogeneity). As this will rarely be the case, TIVA is quite clearly the winner of this part of this contest. When one is ready to miss the smaller biases that are associated with larger true effect sizes, p -uniform is a viable alternative: Its power was identical or similar to TIVA's when $\delta \leq 0.2$ and was clearly smaller only when $\delta \geq 0.5$ (see Figure S7).

Given that homogeneity appears to be rather rare in meta-analyses (van Erp et al., 2017; Stanley, Carter & Doucouliagos, 2018), it is crucial whether TIVA's superiority generalizes to situations with (some) heterogeneity. The graphs for $\tau = 0.3$ and 100% bias in Figure 3 illustrate that heterogeneity diminishes the statistical power of all methods, with the exception of trim-and-fill. Hence, even though TIVA's power suffered from heterogeneity, it generally prevailed over the other methods given small true effect sizes ($\delta \leq 0.2$) in small study sets ($k \leq 10$; see Figure S10). When the true effect size increased, the detection rate of all methods quickly dropped toward zero in small study sets, rendering the detection of publication biases a futile endeavor. Thus, given $k \leq 10$, TIVA appears to be the best available method under 100% publication bias. However, considering that the power of all detection methods at $k = 10$ and $\tau = 0.3$ was already below 40% when $\delta = 0$, and below 30% when $\delta = 0.2$, any attempt to detect even the most severe biases in study sets of this size will only be promising when there is little or no heterogeneity.

In larger study sets ($k \geq 30$), TIVA was outperformed by trim-and-fill, the only method that was hardly affected by heterogeneity. As the power of trim-and-fill did not trail far behind the power of TIVA under homogeneity in large study sets, trim-and-fill may be considered the preferable option when $k \geq 30$. This assessment requires a word of caution: As described above, trim-and-fill assumes that the j studies with the smallest effect sizes are excluded from publication. Despite this difference between the selection model of trim-and-

fill and the selection based on p -values simulated here, the false positive rate of trim-and-fill did not exceed the nominal α -level in any condition of our design, and its power in larger study sets was acceptable under homogeneity yet comparatively high under heterogeneity. Therefore, trim-and-fill is a viable option not only for the detection of effect size based biases (Duval & Tweedie, 2000), but also for the detection of significance based biases when $30 \leq k \leq 50$. This, however, should not obscure the fact that trim-and-fill consistently and drastically underestimated the number of excluded studies. For example, when $\delta = 0$ and $\tau = 0$, it took (on average) 600 studies to obtain $k = 30$ significant results. Trim-and-fill's mean estimate of the number of missing studies in this condition, however, was not 570 but $j = 9.4$. As a consequence, the bias-corrected effect size estimates also offered by trim-and-fill very severely overestimated the true effect sizes (our results on corrected effect size estimates are partly documented in osf.io/3p65a/; very similar results regarding trim-and-fill are reported by Carter et al., 2019). The method may, therefore, correctly detect biases but, at the same time, falsely create the impression that these biases do not greatly inflate effect size estimates. Thus, while trim-and-fill may be useful for the detection of biases caused by the exclusion of non-significant studies, it is definitely not useful for the correction of such biases – and its effect size estimates should simply be ignored.

A final remark on the results under 100% bias primarily concerns the performance of Begg's rank correlation and PET. As already alluded to, the graphs in Figure 3 present results that are collapsed across the five conditions with primary studies that have different mean sample sizes and different degrees of variation in these sample sizes. The statistical power of all methods differed between these conditions but this did not affect the relative performance of trim-and-fill, TES, TIVA and p -uniform. For the latter three methods effects on the statistical power were mainly driven by variations in mean sample size. Like larger true effects, larger mean sample sizes imply greater power in primary studies. Accordingly, the

performance of TES, TIVA and p -uniform tended to decline when mean sample sizes increased. On the contrary, the performance of the methods assessing the association between observed effect sizes and their standard errors was primarily and quite strongly affected by the variation in sample sizes of primary studies. For example, with $\delta = 0$ and $\tau = 0$, the average power of Begg's rank correlation and PET was about 20% when sample sizes were drawn from the distributions $U(20, 30)$ and $U(45, 55)$, but it rose to approximately 60% when sample sizes were drawn from $U(10, 40)$, $U(35, 65)$ or $U(20, 80)$. Both methods continued to be outperformed by other procedures in our simulation even when the variation in sample sizes was largest, but in general the power of these methods may catch up when there is more variation in sample sizes and the number of available studies is large (separate power values of Begg's rank correlation and PET, for different degrees of variation in sample sizes, are given in Figures S13 to S15).

Optional stopping. With additional optional stopping, the statistical power of all methods increased. Under homogeneity ($\tau = 0$), now all methods, with the exception of trim-and-fill, achieved a statistical power of close to 100% with $k = 10$ studies when the true effect size was small ($\delta \leq 0.2$; see Figure S7). In even smaller study sets, TIVA still prevailed. With larger true effects ($\delta \geq 0.5$ and $\tau = 0$), it became more apparent that optional stopping affected the sensitivity of the methods differently. More specifically, the performance of Begg's rank correlation and PET benefitted most from additional optional stopping, and these methods outperformed all other methods with $k \geq 10$ studies (again, excluding TES due to its inflated Type I error rate under heterogeneity).

When there was heterogeneity ($\tau = 0.3$), the pattern of results was governed by effects we already observed: Heterogeneity diminished the power of all methods (again, except for trim-and-fill), but power was still larger than in a situation with mere publication bias. Additionally, the power gains induced by optional stopping were particularly large for Begg's

rank correlation and PET. Under these conditions, these two methods dominated all other methods (excluding TES) with as little as $k = 7$ studies.

The steep increase in power observed for most of the methods can, of course, be explained by the ubiquitous and intense optional stopping we simulated. This raises the question as to what extent these findings may generalize to different implementations of optional stopping or other forms of p -hacking. In principle, the degree to which p -hacking impacts the distribution of significant (and, hence, published) p -values depends on the correlation between subsequent p -values generated in the process of sequential testing (Bishop & Thompson, 2016). A high correlation between subsequent p -values implies that a significant p -value, which may be eventually obtained via the process of p -hacking, differs only by a small margin from the preceding non-significant p -value. Thus, a high correlation between subsequent p -values will induce a strong clustering of published p -values just below the significance threshold and, therefore, produce a left-skew in their distribution.⁴

As p -uniform infers bias from an insufficient right-skew, its power will increase with this correlation. Obviously, by simulating optional stopping with repeated tests after 1 or 5 additional participants, we created particularly strong correlations between subsequent p -values. Thus, less intense optional stopping or other forms of p -hacking (e.g., testing several correlated dependent variables, including or excluding outliers, or controlling for one or more covariates) should also increase the power of p -uniform, but to a lesser extent depending on the correlation between subsequent p -values that they produce.

⁴ Obviously, this only applies when p -hacking is conceived of as a process of repeated sequential testing in the first place. Sequential testing is stopped when a significant result is obtained or when some stopping criterion is reached. It should be noted that there are several simulation studies that implemented p -hacking in a different way: For instance, Ulrich and Miller (2015) assumed that multiple tests on several correlated dependent variables are run *in parallel* and that only the variable providing the *smallest significant* p -value is reported (see also van Aert, Wicherts & van Assen, 2016). Such a procedure will not cause a clustering of p -values just below the significance threshold but, on the contrary, increase the prevalence of particularly small p -values. As a consequence, the impact of this procedure on the bias in meta-analytical effect size estimates and on the detectability of this bias will differ from our results on the effects of optional stopping. Thus, all of the following considerations are valid only when researchers p -hack their data by engaging in some form of sequential testing.

A similar logic also applies to TES and TIVA. With TIVA, a stronger clustering of published p -values just below the significance threshold will reduce the variance in back-transformed z -values and, therefore, increase the power of the method. For TES, a stronger clustering of p -values just below the significance threshold means that there are more primary studies with an estimated power just above 50%. Consequently, these studies will reduce the expected number of significant results and, hence, boost the power of TES. In line with this reasoning, we observed larger power across all these methods when testing was repeated after 1 additional participant as compared to 5 additional participants. However, due to our rather weak manipulation of the correlation between subsequent p -values, this effect was generally small (we, therefore, report mean power values across both optional stopping conditions throughout).

For Begg's rank correlation and PET, the situation is slightly different. A stronger clustering of p -values just below the significance threshold also implies an increased correlation between effect size estimates and their standard errors. Thus, as with the other detection methods, all forms of p -hacking can be expected to increase the power of Begg's rank correlation and PET to the degree to which they produce a high correlation between subsequent p -values. However, our implementation of optional stopping not only created high correlations between subsequent p -values but also affected the correlation between effect size estimates and their standard errors directly: As the initial sample size in all primary studies was $n_1 = n_2 = 20$, and testing was repeated after each additional participant, studies that reached significance at a later point in testing (i.e. with larger n and a smaller standard error) were almost guaranteed to also have a smaller observed effect size. Thus, our implementation of optional stopping produced an extremely strong correlation between sample size and effect size (that was reduced only slightly when testing was repeated after increments of 5 additional participants). Obviously, this correlation will decrease with

increasing variation in initial sample sizes. Furthermore, other forms of p -hacking cannot necessarily be expected to affect the correlation between effect sizes and sample sizes to the same degree as optional stopping. Therefore, we do not assume that the large gains in power for Begg's rank correlation and PET that we observed generalize to other implementations of optional stopping or other forms of p -hacking.

The performance of all detection methods will, of course, not only be affected by the intensity of p -hacking (as indicated by the correlation between subsequent p -values) but also by the proportion of p -hacked studies in the available study set. Given the parameter settings of our simulation, this proportion is fully determined by the power of the first statistical test. As we began optional stopping with an initial sample size of $n_1 = n_2 = 20$, power for the four true effect sizes was 5.0%, 15.3%, 46.3%, and 79.9%. Thus, the graphs depicted in Figure 3 show statistical power for the different detection methods when 95%, 84.7%, 53.7%, or 20.1% of the available studies were intensely p -hacked. In general, the power of the detection methods will decline to the degree that a smaller proportion of studies is p -hacked, or that p -hacking in these studies is less intense.

Two-step bias. Compared to the condition with 100% bias, the power of all methods was diminished when studies with p -values between $p = .05$ and $p = .10$ had a 20% chance of being published. This loss of power was mostly slight. However, for p -uniform and TES it was more pronounced when the true effect size was small and there were only a few studies available ($\delta \leq 0.2$ and $k \leq 10$). The reason for this is two-fold: First, the proportion of non-significant studies included in meta-analyses was largest when the true effect size was $\delta = 0$ (see Figure S2). Second, non-significant results affected the performance of p -uniform and TES more heavily in small study sets. For p -uniform, this occurred because the method solely considers significant findings for the detection of bias. Thus, the inclusion of non-significant results in small study sets left only a few studies for the p -uniform analysis. For TES, a

smaller difference between the observed and expected proportion of significant findings, of course, led to a larger loss of power in smaller study sets.

Beyond this peculiarity, the pattern of results in the two-step bias condition closely resembled the pattern in the 100% bias condition at weakly reduced power levels. Most importantly, the performance of TIVA was comparatively robust to the inclusion of marginally significant results. This strengthens our recommendation for this method. Under homogeneity, TIVA was the only method that achieved more than 80% power for the detection of the most severe biases (at $\delta = 0$) with as little as $k = 7$ studies. It consistently outperformed all other methods at all true effect sizes (again, with the exception of TES in large study sets) as long as $\tau = 0$. Again, the validity of this recommendation is limited by situations in which there is heterogeneity: Starting from $\tau = 0.3$ and with large study sets ($k \geq 30$), the power of TIVA was consistently exceeded by the power of trim-and-fill (see also Figure S9).

90% bias. When 10% of studies were published irrespective of their statistical significance, the pattern of results clearly changed: Under homogeneity and with $\delta = 0$, the power of all methods but TES dropped below 30% even with $k = 50$ studies. The power of Begg's rank correlation, PET, trim-and-fill and TIVA was approximately as large as, or even below, the nominal α -level. With $k = 10$ studies, TES achieved a power of about 20%. Thus, with a censoring of 90% of studies, the resulting bias in meta-analytic estimates at $\delta = 0$ will mostly remain undetected with study sets of this size no matter which method is applied. In larger study sets ($k \geq 30$), the power of TES was actually extremely high, but the applicability of this method remains restricted due to its susceptibility to increased Type I error rates under heterogeneity.

With an effect size of $\delta = 0.5$ (and $\tau = 0$), the power of most methods was only weakly reduced as compared to the 100% bias condition. This is due to the fact that at $\delta = 0.5$, the

proportion of significant results among published studies was already 95% (see Figure S2). Hence, the data basis in the 90% bias condition closely resembled the data basis in the 100% bias condition with $\delta \geq 0.5$. Interestingly, the power of trim-and-fill was nonetheless very adversely affected by the inclusion of only a few uncensored studies.

It is worth noting, that p -uniform achieved its maximum power at an effect size of $\delta = 0.2$. With this effect size, its power clearly exceeded 60% in large study sets ($k \geq 30$; see Figure S7). The large gain in power, as compared to the condition with $\delta = 0$, was caused by the large increase in the proportion of significant results among published studies (34.5% with $\delta = 0$ and 73.8% with $\delta = 0.2$). In contrast, for TIVA the mixture of significant and non-significant results at $\delta = 0.2$ did not bring about a restriction in the variance of test statistics that could be reliably detected, with performance of the method peaking at $\delta = 0.5$. Thus, in general, it was possible to uncover biases in meta-analytic estimates with p -uniform and TIVA in larger study sets when 90% of studies were censored. However, given the sample sizes of primary studies in this simulation, p -uniform preferentially detected biases at $\delta = 0.2$, whereas TIVA preferentially detected biases at $\delta = 0.5$. While these biases are of similar size as the bias at $\delta = 0$ (or even slightly larger), they are hardly judged as more relevant.

When there is heterogeneity, the results in the 90% bias condition can be easily summarized: With increasing heterogeneity the power of all methods ensuring a proper control of Type I errors quickly collapsed. Already at $\tau = 0.2$, the power of all of these methods was consistently at or below 30% (see Figure S9).

Discussion

This simulation study investigated the detectability of different degrees of publication bias across a large range of conditions and evaluated the performance of six statistical detection methods. It identified a number of circumstances under which biases can be

uncovered reliably but it also determined a variety of scenarios under which bias detection is unlikely or even bound to fail. With regard to the performance of the detection methods, a first central result was that none of the methods kept the nominal α -level across all simulated conditions. However, only one method was associated with systematically inflated Type I error rates: TES produced too many false positives when there was heterogeneity. The statistical power of the detection methods followed a complex pattern. None of the methods consistently outperformed all others. Hence, choosing an optimal method would require knowledge about underlying parameters that meta-analysts cannot have.

In the following, we will first summarize our results on the statistical power of detection methods and then derive several recommendations for the application of such methods from these findings. Afterwards, we will discuss constraints on the generalizability of these recommendations as well as some additional limitations of our study. We end with some considerations on the interpretation of results from bias tests and the usefulness of these methods in general.

When and How Can Publication Biases be Detected?

Two factors that appear to be of central importance for the detectability of biases are the degree of heterogeneity and the degree of censoring of non-significant results. Therefore, we structure the discussion of results on statistical power along two distinctions: one between meta-analyses with homogeneous and heterogeneous effect sizes; one between literatures where non-significant results are (almost) fully censored (represented here by the conditions with 100% bias, optional stopping and two-step bias) and literatures where non-significant results have an appreciable chance of being published (here, 90% bias).

Let us first consider the combination of homogeneity with complete or severe censorship. Under these circumstances, properly chosen methods often succeeded in detecting the resulting bias in meta-analytic effect size estimates. While several methods

performed well, TIVA appeared to be the best available method. In the 100% bias condition, TIVA achieved high statistical power with even as little as 5 studies when the true effect size was $\delta = 0$ and the distortion in effect size estimates was most severe. As had to be expected, the power of all methods increased in larger study sets but declined with larger true effect sizes (and, hence, smaller overestimation). As a consequence, when the true effect was of medium size ($\delta = 0.5$), a reliable detection of biases was only possible in our simulation when at least 30 studies were available. However, TIVA was still the most recommendable option to accomplish this task. When simulated primary studies were additionally affected by optional stopping, the power of all methods increased further. The large gains in power that we observed were certainly due to the specific implementation of optional stopping we realized here. But, as already discussed above, increases in power are likely to also occur when only a portion of researchers is willing to engage in *p*-hacking or when *p*-hacking is less intense. TIVA was outperformed by other methods under optional stopping, but this occurred only when the true effect size was large ($\delta \geq 0.5$). It should be noted that in practice, meta-analysts can know neither the true effect size nor the degree to which available studies are affected by *p*-hacking. Furthermore, most meta-analysts will likely be most interested in biases that are associated with small or non-existent true effects. Therefore, the observation that other methods performed better when large effects were combined with intense optional stopping gives little reason to prefer these methods. Finally, TIVA's performance proved relatively robust against the inclusion of some "marginally significant" results in meta-analytic study sets. While its power declined, it nevertheless prevailed over the other methods, and still detected the most severe biases (at $\delta \leq 0.2$) with high power, when at least 10 studies were available.

The limitations, however, of this positive view, regarding both the detectability of publication biases and TIVA's performance, become apparent when we follow the

distinctions introduced above. Heterogeneity affected the performance of detection methods adversely even when censorship was complete or severe. This occurred even though heterogeneity exacerbated the bias in meta-analytic effect size estimates. For the analysis of real data sets this is of great importance as heterogeneity appears to be the rule rather than the exception in psychological meta-analyses. In a collection of 187 meta-analyses using a standardized mean difference as effect size metric (van Erp et al., 2017), 50% of τ estimates were larger than 0.2, and 30% were larger than 0.3. Under 100% bias and with $\tau = 0.2$, TIVA still achieved acceptable power when the true effect size was small ($\delta \leq 0.2$) and at least $k = 10$ studies were available. With $\tau = 0.3$, its power exceeded 50% only when there was no true effect ($\delta = 0$) and the meta-analytic data set comprised 30 or more studies. Furthermore, with a heterogeneity of $\tau \geq 0.2$, TIVA was almost always outperformed by trim-and-fill in large study sets ($k \geq 30$).

This illustrates that there is already no single best method for the detection of publication biases that are caused by a complete suppression of non-significant results. Even when no other selection mechanism is considered, an optimal choice of a detection method would require knowledge about the (mean) true effect size and the degree of heterogeneity. Additionally, our results suggest that the detection of biases may quickly become difficult with increasing heterogeneity. Even with moderate degrees of heterogeneity, the power of all methods is rather low in small study sets.

Our second distinction, between an (almost) complete censoring and a censoring that gives non-significant results a small chance of being published, is important as it largely separates situations in which bias detection can succeed from situations where it is most likely to fail. In the 90% bias condition, detection methods mostly had small power. When there was a moderate degree of heterogeneity ($\tau = 0.2$), the power of all methods that yielded a proper control of Type I errors was consistently below 30% irrespective of true effect size

and number of available studies. Under homogeneity and with $\delta = 0$, TES was the only method that achieved good and even excellent power. Interestingly, TES was clearly the best available option in this specific constellation. Applications of TES, however, will remain doubtful due to the inflated Type I error rate of the method under heterogeneity.

The drop of power observed in the 90% bias condition had to be expected as the resulting bias in meta-analytic effect size estimates was also considerably smaller than in the 100% bias condition. This, however, does not imply that the remaining bias in effect size estimates was irrelevant: In fact, at $\delta = 0$ and $\tau = 0$, the mean effect size estimate was $\bar{d} = 0.15$. A meta-analytic effect of this size may still give rise to the false impression that there actually is a true effect. Moreover, with some heterogeneity ($\tau = 0.2$ or $\tau = 0.3$), the meta-analytic effect was inflated to a size of about $\bar{d} = 0.30$. Notably, effect estimates of this size appear to be quite common in psychological meta-analyses (Stanley, Carter, & Doucouliagos, 2018) and are most likely to lead to the conclusion that a true effect is present. With that said, the power of all detection methods (excluding TES) was reduced to less than 10% even in large study sets. Thus, the combination of a censoring that gives non-significant results a small chance of being published with a moderate degree of heterogeneity creates a situation in which the bias in meta-analytic estimates can be both highly relevant yet also mostly undetected.

Four Recommendations for the Application of Tests for Publication Bias

Given the complexity of the results, one may wonder which practical advice can be given to meta-analysts that aim to assess whether their collection of data is affected by bias. We think that our studies allows for at least four recommendations.

The first of these recommendations is that meta-analysts should strive to avoid heterogeneity before applying methods for the detection of publication bias. Of course, the actual degree of heterogeneity in a collection of studies cannot be known. Furthermore, it can

also not be estimated properly in the presence of publication bias. To demonstrate this, Augusteijn, van Aert & van Assen (2019) showed analytically that publication bias can have severe and rather complex effects on heterogeneity estimates: Depending on the combination of the actual degree of heterogeneity, true effect size and degree of censorship heterogeneity may either be under- or overestimated. Within the parameter settings of our study, the actual τ was oftentimes underestimated drastically in all four selection conditions realized (see Figure S16). Hence, even low estimates of τ do not preclude that heterogeneity might actually be large. Consequently, there is no way for meta-analysts to ensure that heterogeneity is small or absent.

However, meta-analysts should be aware that their inclusion criteria for primary studies can partly control the degree of heterogeneity that has to be expected. Primary studies that differ widely in design, procedure, measures, or participant population are likely to be inflicted by more heterogeneity than studies that are methodologically more similar. Our results illustrate that, under publication bias, combining methodologically diverse studies may be a way to both inflate meta-analytic effect size estimates further as well as disguise this bias at the same time. Thus, meta-analysts should decide on plausible moderator variables on theoretical grounds, and apply tests for publication bias to subsets of studies, defined by the different levels of these moderators. The gain in power of the detection methods achieved by reducing heterogeneity will often outweigh the loss induced by the lower number of primary studies available at different moderator levels. Of course, the number of primary studies in a subset should not become too small. That said, there may be solutions to this problem: At least for p -uniform it should be possible to adapt the method in a way that allows to incorporate different meta-analytic effect size estimates per subset, while including all available primary studies in a single test for publication bias.⁵

⁵ We thank Marcel van Assen for pointing this possibility out to us.

Our second recommendation concerns the choice of a test for the detection of publication bias. Within our study, TIVA either performed best or at similar levels as any other method as long as censorship was severe and heterogeneity absent or small. With larger degrees of heterogeneity, it was often outperformed by trim-and-fill, especially in large study sets. As will be discussed below, these results may be restricted to the parameter settings simulated here. That said, at least whenever these parameter settings appear plausible, it seems reasonable to apply both these tests. If either of these methods yields a significant result this should be considered a red flag, indicating that meta-analytical results may not be valid. Considering the low Type I error rates of both of these methods in almost all conditions, there is little reason to worry that this multiple testing gives rise to inflated false alarm rates. Lastly, with regard to trim-and-fill, it should be emphasized again that the method may be useful for the detection of biases, but is not recommended for the estimation of true effect sizes when bias is present (Carter et al., 2019; van Aert, Wicherts & van Assen, 2016).

Finally, if a sizeable proportion of the available primary studies is non-significant and there is strong reason to assume that these studies are homogeneous (for instance, because the studies can be considered as close replications of each other), it may be advisable to apply TES: TES was the only method providing a reasonable chance to detect bias when censorship was less extreme (in the 90% bias condition). In other words, with the exception of the TES method, biases that may occur even though some studies are published irrespective of their results will mostly remain undetected by all other methods.

The third recommendation pertains to the interpretation of non-significant results of bias tests. In view of the rather low power that the detection methods achieve under heterogeneity or with less severe censoring, it is clear that a non-significant result of a bias test provides little information, especially when the number of available studies is small.

Hence, a non-significant result should not be interpreted as evidence for the absence of bias, but simply as the absence of evidence.

Our final recommendation is again justified by the oftentimes low power of detection methods: To raise power, meta-analysts should consider to test for bias at an increased significance level of 10%, at the very least when the test is applied to a small set of studies. Similar suggestions have been made by a number of authors with regard to a variety of detection methods (e.g, Begg & Mazumdar, 1994; Ioannidis & Trikalinos, 2007; Francis, 2013).

Limitations

Even though our simulation investigated a large number of conditions, it still captures only a small sample of the multi-dimensional space representing the ways in which data may be generated, analyzed and selected for publication. This may pose restrictions to the generalizability of our results to data that were generated or censored under conditions not simulated here.

We used a simple two group design to simulate data in primary studies and measured the effect size in these primary studies as a standardized mean difference. Obviously, real world research uses a multitude of other (and often more complex) designs and may also employ different effect size measures even in two group comparisons. There are, however, reasons to assume that neither the design of primary studies nor the specific effect size measure is of great importance for our results. Of course, different effect size measures may be converted into each other without affecting effect direction or statistical significance in primary studies. Hence, detection methods that rely on p -values or on power estimates (p -uniform, TIVA and TES) are by definition independent of specific effect size measures. Stable p -values with different effect metrics also imply that the ratio of effect estimates to their standard errors is stable. This ensures that the results of Begg's rank correlation and

PET remain unaltered across different suitable metrics. Finally, trim-and-fill considers exclusively the rank order of effect sizes – a property that will be invariant across any meaningful transformation of measures.

With regard to design, it is helpful to distinguish between primary studies and meta-analyses summarizing these primary studies. While primary studies may use all kinds of complex multi-factorial designs, most meta-analyses focus on simple contrasts between two conditions. Thus, when extracting information from primary studies, meta-analysts typically aim to compute effect sizes that describe the difference between the two focal conditions, and which are comparable even though the data were collected in different designs. Accordingly, most meta-analyses use standardized mean differences (Cohen's d or Hedges' g) or (point-biserial) correlations as effect measures (van Erp et al., 2017; Fanelli, Costas & Ioannidis, 2017). As such, in principle our results should be applicable to all meta-analyses proceeding in this way.

There are, however, constraints on the generalizability of our findings that we deem as more relevant: A recurring theme of our summary of the results has been that a performance assessment of detection methods does not generalize across different simulated selection mechanisms, true effect sizes or degrees of heterogeneity. The results also demonstrate that the performance of the methods depends on the power of primary studies – and, therefore, on their sample sizes. With regard to true effect sizes and degrees of heterogeneity, the parameter settings simulated here may cover a broad range of values that appear plausible in psychology. This is, however, certainly not true for selection mechanisms and sample sizes.

Sample sizes in a given area of research may easily differ, in their mean or variance, from what we simulated here. Furthermore, the distribution of sample sizes of psychological studies appears to be right-skewed (Marszalek et al., 2011). Hence, actual meta-analyses may easily include one or a few studies with unusually large sample sizes. Larger mean sample

sizes will generally decrease the detectability of biases which may, additionally, occur at different rates for different detection methods, depending on a variety of other factors (e.g., heterogeneity). We already touched upon the relevance of variance in sample sizes when we presented the results for Begg's rank correlation and PET. These methods attain higher statistical power when the variation in sample sizes increases (which may, of course, also be caused by a few outliers). In contrast, TIVA's power is likely to be reduced by larger variation in sample sizes as this variation will also increase the variance of converted z -scores. The different levels of variation realized here had surprisingly little impact on TIVA's power. Still, it has to be expected that larger variation would diminish TIVA's performance and may cause Begg's rank correlation and PET to be the more sensitive methods.

The mechanisms by which real world studies are selected for publication may differ from the selection mechanisms that we simulated in a number of ways: Censoring of non-significant results may be less severe, the functional relationship between p -values and publication probabilities may take a different form, and publication probabilities may be affected by factors other than p -values. It is possible that even small deviations from the selection mechanisms simulated here can alter the performance of detection methods severely. For example, the comparatively good performance of trim-and-fill that we observed in many conditions with large study sets may very well be limited to the strict one-sided censoring of studies simulated here. Specifically, it is known that trim-and-fill fails to detect bias when a single unusually small or negative effect size is included in the meta-analytic data set (Duval & Tweedie, 2000; Duval, 2005). In real data sets, such an outlier in the distribution of effect sizes may, of course, occur due to reasons that are totally unrelated to the p -value of the respective study. More generally, many of our results may not generalize to selection mechanisms that assign a non-zero probability of publication to negative significant results. This is easily illustrated with PET. When negative significant results are published

with a small probability they will appear as outliers in a funnel plot. Obviously, these outliers will affect the regression slope that PET uses to diagnose bias.

Overall, the ways in which characteristics of primary studies and features of the publication process impact the performance of detection methods are likely to be even more complex than suggested by our results. It should be noted that this also qualifies our recommendation for TIVA and trim-and-fill: These methods can only be expected to show a comparatively good performance in situations in which it seems plausible that the conditions simulated here in fact apply.

A final limitation of our study concerns the detection methods that we evaluated. The performance of most of these methods deteriorated under heterogeneity. This suggests that methods might fare better if they explicitly take into account that primary studies may have heterogeneous effect sizes. Indeed, there is a class of methods that do exactly this: Selection model approaches (e.g., Hedges, 1984; Iyengar & Greenhouse, 1988; Hedges & Vevea, 1996; Copas & Shi, 2000) combine a model of the data generation process with a model of the selection process to mitigate bias. Maximum likelihood techniques are used to estimate the parameters of these models. For instance, the parameters of the three-parameter selection model (Iyengar & Greenhouse, 1988) are δ , τ and the probability p that a non-significant study gets published. Thus, the method provides corrected estimates of the true effect size and the degree of heterogeneity, and allows for a statistical test of the hypothesis that bias is present as based on the parameter p . We did not include selection model approaches here because one of our goals was to assess the minimum number of studies necessary to detect bias. This led us to focus on meta-analyses that include only a small number of primary studies ($k \leq 10$). In such small study sets, selection model approaches can be expected to have severe convergence problems (Field & Gillett, 2010) and will often not yield any results at all. Thus, it is unlikely that they outperform the methods considered here when only a small

number of studies is available. In larger study sets, however, they are certainly a viable option. Notably, a recent simulation study (Pustejowsky & Rodgers, 2018) found that the three-parameter selection model performed comparatively well and may be the preferable choice in sets including 20 or more studies.

How to Interpret a Significant Result of a Bias Test?

We already pointed out that non-significant results of bias tests will typically provide little information. However, the interpretation of significant results of bias tests also entails some problems that should not be neglected.

Under severe publication bias, all of the detection methods considered here attained the highest statistical power when there was no true effect and the resulting distortion in the meta-analytic effect size estimate was particularly pronounced. However, biases may and will also be detected when there is a true effect and the distortion in meta-analytic estimates is reduced. In the 90% bias condition, TIVA and p -uniform detected even more biases with small or moderate true effect sizes than with $\delta = 0$. Obviously, then, a significant bias test does not allow for the conclusion that there is no true effect. Furthermore, a significant result that occurs with a medium effect size of $\delta = 0.5$ will (mostly) correctly indicate that the meta-analytic effect size estimate is not valid. At the same time, the distortion in this estimate will often be relatively small and may be of little relevance for any substantive interpretation. Thus, it remains unclear what should be concluded from a significant bias test with regard to the effect under investigation. For this reason, methods that are not only capable to detect but also correct biases in meta-analytic effect size estimates are desirable. The results of simulation studies on such corrected estimates, however, are not very encouraging (Carter et al., 2019; McShane & Böckenholt, 2016). Thus, there is good reason to assume that methods for the correction of bias, when applied to real world meta-analyses, will rarely converge into an estimate that appears trustworthy and reliable. The only justifiable conclusion from a

significant result of a publication bias test will, therefore, often be that the meta-analytic effect size estimate is not valid and does not allow for any substantive claim about the effect under investigation.

The second caveat for the interpretation of a significant result of a bias test stems from the fact that none of the detection methods actually tests “directly” for publication bias. All of the detection methods search for some signal in the available data that occurs when bias is present (e.g., an association between effect sizes and their standard errors, or an irregularity in the distribution of p -values). Therefore, a significant result can only be interpreted as an indication of bias to the degree to which alternative explanations for the occurrence of the respective signals can be excluded. In principle, there are always alternative explanations to publication bias, no matter which method was used to detect such bias. Within our simulation, this is apparent for TES. Here, the alternative explanation for a significant result is simply heterogeneity. In real data sets, however, there may also be benign causes for an association between observed effect sizes and their standard errors and, thus, for a significant result of Begg’s rank correlation or PET. Such an association may occur, for instance, when true effect sizes vary systematically and researchers successfully adapt their sample sizes to these true effect sizes (Sterne, Gavaghan, & Egger, 2000). Similarly, it has been shown that no specific distribution of p -values implies a specific selection process (Ulrich & Miller, 2015). Thus, while significant results for p -uniform and TIVA do in fact suggest selective reporting, i.e. that not all results of primary studies are included in the meta-analysis, it does not necessarily imply that the selection process had a preference for significant or small p -values in primary studies. In other words, significant results for these detection methods are not by default a function of excluding non-significant primary studies; moreover, they may also not cause an overestimation of true effect sizes (Erdfelder & Heck, this issue). Thus, it is important to emphasize that, in general, no detection method can yield “proof” of bias.

Therefore, a careful interpretation of the results of these methods requires a consideration of alternative explanations.

Conclusion

To conclude, we may ask what can be achieved by detection methods. As we have seen, these methods can uncover biases with high statistical power when censoring of non-significant results is very severe and true effect sizes are homogeneous. Properly chosen methods may even accomplish this in small study sets. This provides a sufficient reason to use such methods at all. For instance, TIVA or p -uniform may be used to identify multi-study papers that test a hypothesis in several experiments and present a doubtful conclusion. However, across all data generation and selection processes that appear to be realistic in the psychological literature, there is neither a single “best” detection method nor will any method be able to separate biased data collections from unbiased data collections reliably. Under many scenarios, relevant biases in meta-analytic effect size estimates are likely to remain undetected. On the other hand, biases that are detected may be small and of little relevance for any substantive interpretation. This implies that not only detection methods, but also the meta-analyses themselves, will not yield any definite conclusions as long as publication biases and questionable research practices have to be expected in the psychological literature. Thus, our results may be understood as a plea for open science practices: If there is doubt on the trustworthiness of the available evidence on a specific research question, then any attempt to resolve this will ultimately have to rely on large-scale, pre-registered replications. In the future, we should produce and publish evidence that is of better quality in the first place. Pre-registrations and openly available data should become standard for all primary studies. In essence, publication decisions should no longer depend on statistical results but on theoretical and methodological soundness.

References

- American Psychological Association (2010). Meta-Analysis Reporting Standards (MARS). In A. P. Association, *Publication Manual of the American Psychological Association* (6th ed., pp. 245-254). Washington.
- Augusteijn, H. E. M., van Aert, R. C. M. & van Assen, M. A. L. M. (2019). The effect of publication bias on the Q test and assessment of heterogeneity. *Psychological Methods*, 24(1), 116-134. <http://doi.org/10.1037/met0000197>
- Begg, C. B., & Mazumdar, M. (1994). Operating characteristics of a rank correlation test for publication bias. *Biometrics*, 50(4), 1088-1101. <https://doi.org/10.2307/2533446>
- Bishop, D. V. M. & Thompson, P. A. (2016). Problems in using *p*-curve analysis and text-mining to detect rate of *p*-hacking and evidential value. *PeerJ* 4:e1715. <https://doi.org/10.7717/peerj.1715>
- Bürkner, P. C., Bittner, N., Holling, H., & Buhlmann, U. (2017). D-cycloserine augmentation of behavior therapy for anxiety and obsessive-compulsive disorders: A meta-Analysis. *PLoS ONE*, 12(3), 1-19. <https://doi.org/10.1371/journal.pone.0173660>
- Carter, E. C., Schönbrodt, F. D., Gervais, W. M., & Hilgard, J. (2019). Correcting for bias in psychology: A comparison of meta-analytic methods. *Advances in Methods and Practice in Psychological Science*, 2(2). <https://doi.org/10.1177/2515245919847196>
- Cohen, J. (1988). Statistical power analysis for the behavioral sciences (2nd ed.). Hillsdale, NJ: Erlbaum.
- Copas, J. & Shi, J. Q. (2000). Meta-analysis, funnel plots and sensitivity analysis. *Biostatistics*, 1(3), 247-262. <https://doi.org/10.1093/biostatistics/1.3.247>

DerSimonian R, Laird N (1986). Meta-analysis in clinical trials. *Controlled Clinical Trials*, 7(3), 177-188. [https://doi.org/10.1016/0197-2456\(86\)90046-2](https://doi.org/10.1016/0197-2456(86)90046-2)

Duval, S. (2005). The trim and fill method. In H. R. Rothstein, A. J. Sutton, & M. Borenstein, *Publication bias in meta-analysis: Prevention, assessment and adjustments* (S. 127-144). Chichester: Wiley.

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(2), 455–463. <https://doi.org/10.1111/j.0006-341X.2000.00455.x>

Egger, M., Davey Smith, G., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple graphical test. *British medical journal*, 315(7109), 629-634. <https://doi.org/10.1136/bmj.315.7109.629>

Erdfelder, E. & Heck, D. W. (2019). Detecting evidential value and p -hacking with the p -curve tool: A word of caution. *Zeitschrift für Psychologie*, xx(xx), xx-xx.

Fanelli D. (2010). “Positive” results increase down the hierarchy of the sciences. *PLoS ONE*, 5(4). <https://doi.org/10.1371/journal.pone.0010068>

Fanelli, D., Costas, R. & Ioannidis, J. P. A. (2017). Meta-assessment of bias in science. *PNAS – Proceedings of the National Academy of Science of the United States of America*, 114(14), 3714-3719. <https://doi.org/10.1073/pnas.1618569114>

Field, A. F. & Gillet, R. (2010). How to do meta-analysis. *British Journal of Mathematical and Statistical Psychology*, 63(3), 665-694. <https://doi.org/10.1348/000711010X502733>

Francis, G. (2012). The same old New Look: Publication bias in a study of wishful seeing. *i-Perception*, 3(3), 176-178. <https://doi.org/10.1068/i0519ic>

Francis, G. (2013). Replication, statistical consistency, and publication bias. *Journal of Mathematical Psychology*, 57(5), 153-169. <https://doi.org/10.1016/j.jmp.2013.02.003>

Hedges, L. V. (1981). Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics*, 6(2), 107-128.
<https://doi.org/10.3102/10769986006002107>

Hedges, L. V. (1984). Estimation of effect size under nonrandom sampling: The effects of censoring studies yielding statistically insignificant mean differences. *Journal of Educational and Behavioral Statistics*, 9(1), 61-85.
<https://doi.org/10.3102/10769986009001061>

Hedges, L. V., & Vevea, J. L. (1996). Estimating effect size under publication bias: Small sample properties and robustness of a random effects selection model. *Journal of Educational and Behavioral Statistics*, 21(4), 299-332. <http://dx.doi.org/10.2307/1165338>

Higgins, J. P. T., Thompson, S. G., Deeks, J. J., & Altman, D. G. (2003). Measuring inconsistency in meta-analyses. *British Medical Journal*, 327(7414), 557-560.
<https://doi:10.1136/bmj.327.7414.557>

Ioannidis, J. P. A. & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical Trials: Journal of the Society for Clinical Trials*, 4(3), 245-253.
<http://doi.org/10.1177/1740774507079441>

Iyengar, S., & Greenhouse, J. B. (1988). Selection models and the file drawer problem. *Statistical Science*, 3(1), 109-117. <http://doi.org/doi:10.1214/ss/1177013012>

Jackson, D. (2007). Assessing the implications of publication bias for two popular estimates of between-study variance in meta-analysis. *Biometrics*, 63(1), 187-193.

<https://doi.org/10.1111/j.1541-0420.2006.00663.x>

Kenny, D. A., & Judd, C. M. (2018, September 15). The Unappreciated Heterogeneity of Effect Sizes: Implications for Power, Precision, Planning of Research, and Replication.

Retrieved from <http://davidakenny.net/recp.htm>

Kromrey, J.D., & Rendina-Gobioff, G. (2006). On knowing what we do not know. An empirical comparison of methods to detect publication bias in meta-analysis. *Educational and Psychological Measurement*, 66(3), 357-373. <https://doi.org/10.1177/0013164405278585>

Kühberger, A., Fritz, A. & Scherndl, T. (2014). Publication bias in psychology: A diagnosis based on the correlation between effect size and sample size. *PLoS ONE*, 9(9), e105825. <https://doi.org/10.1371/journal.pone.0105825>

Marszalek, J. M., Barber, C., Kohlhart, J., & Holmes, C. B. (2011). Sample size in psychological research over the past 30 years. *Perceptual and motor skills*, 112(2), 331–348. <https://doi.org/10.2466/03.11.PMS.112.2.331-348>

Macaskill, P., Walter, S., & Irwig, L. (2001). A comparison of methods to detect publication bias in meta-analysis. *Statistics in Medicine*, 20(4), 641–654. <https://doi.org/10.1002/sim.698>

McShane, B. B., Böckenholt, U., & Hansen, K. T. (2016). Adjusting for publication bias in meta-analysis: An evaluation of selection methods and some cautionary notes. *Perspectives on Psychological Science*, 11(5), 730–749. <https://doi.org/10.1177/1745691616662243>

Morey, R. D. (2013). The consistency test does not - and cannot - deliver what is advertised: A comment on Francis (2013). *Journal of Mathematical Psychology*, 57(5), 180-183. <https://doi.org/10.1016/j.jmp.2013.03.004>

Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251). <https://doi.org/10.1126/science.aac4716>

Pustejovsky, J. E., Rodgers, M., A. (2018). Testing for funnel plot asymmetry of standardized mean differences. *Research Synthesis Methods*, 10(1), 57-71. <https://doi.org/10.1002/jrsm.1332>

Rücker, G., Carpenter, J. R., & Schwarzer, G. (2011). Detecting and adjusting for small-study effects in meta-analysis. *Biometrical Journal*, 53(2), 351–368. <https://doi.org/10.1002/bimj.201000151>

Schimmack, U. (2014, December 30). The Test of Insufficient Variance (TIVA): A New Tool for the Detection of Questionable Research Practices [Blog Post]. Retrieved from <http://replicationindex.wordpress.com/2014/12/30/the-test-of-insufficient-variance-tiva-a-new-tool-for-the-detection-of-questionable-research-practices/>

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>

Simonsohn, U. (2013). It really just does not follow, comments on Francis (2013). *Journal of Mathematical Psychology*, 57(5), 174-176. <https://doi.org/10.1016/j.jmp.2013.03.006>

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve and effect size: Correcting for publication bias using only significant results. *Perspectives on Psychological Science*, 9(6), 666–681. <https://doi.org/10.1177/1745691614553988>

Stanley, T. D., Carter, E. C., & Doucouliagos, H. (2018). What meta-analyses reveal about the replicability of psychological research. *Psychological Bulletin*, 144(12), 1325-1346. <http://doi.org/10.1037/bul0000169>

Stanley, T. D. (2017). Limitations of PET-PEESE and other meta-analysis methods. *Social Psychological and Personality Science*, 8(5), 581-591. <https://doi.org/10.1177/1948550617693062>

Stanley, T. D. & Doucouliagos, H. (2014). Meta-regression approximations to reduce publication selection bias. *Research Synthesis Methods*, 5(1), 60–78. <https://doi.org/10.1002/jrsm.1095>

Sterne, J., 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(11), 1119–1129. [https://doi.org/10.1016/S0895-4356\(00\)00242-0](https://doi.org/10.1016/S0895-4356(00)00242-0)

Terrin, N., Schmid, C. H., Lau, J. & Olkin, I. (2003). Adjusting for publication bias in the presence of heterogeneity. *Statistics in Medicine*, 22(13), 2113-2126. <https://doi.org/10.1002/sim.1461>

Ulrich, R., Miller, J., & Erdfelder, E. (2018). Effect size estimation from *t* statistics in the presence of publication bias: A brief review of existing approaches with some extensions. *Zeitschrift für Psychologie*, 226(1), 56-80. <https://doi.org/10.1027/2151-2604/a000319>

Ulrich, R., & Miller, J. (2015). P-hacking by post hoc selection with multiple opportunities: detectability by skewness test? Comment on Simonsohn, Nelson and Simmons (2014). *Journal of Experimental Psychology: General*, 144(6), 1137-1145.

<https://doi.org/10.1037/xge0000086>

van Aert, R. C. M., Wicherts, J. M. & van Assen, M. A. L. M. (2019). Publication bias examined in meta-analyses from psychology and medicine: A meta-meta-analysis. *PLoS ONE* 14(4): e0215052. <https://doi.org/10.1371/journal.pone.0215052>

van Aert, R. C. M., Wicherts, J. M., & van Assen, M. A. L. M. (2016). Conducting meta-analyses based on p values: Reservations and recommendations for applying p-uniform and p-curve. *Perspectives on Psychological Science*, 11(5), 713-729.

<https://doi.org/10.1177/1745691616650874>

van Assen, M. A. L. M., van Aert, R. C. M., & Wicherts, J. M. (2015). Meta-analysis using effect size distributions of only statistically significant studies. *Psychological Methods*, 20(3), 293. <https://doi.org/10.1037/met0000025>

van Erp, S., Verhagen, J., Grasman, R. P. P. P., & Wagenmakers, E.-J. (2017). Estimates of between-study heterogeneity for 705 meta-analyses reported in *Psychological Bulletin* From 1990–2013. *Journal of Open Psychology Data*, 5(1), 4.

<https://doi.org/10.5334/jopd.33>

Viechtbauer, W. (2007). Confidence intervals for the amount of heterogeneity in meta-analysis. *Statistics in Medicine*, 26(1), 37-52. <https://doi.org/10.1002/sim.2514>

Wagenmakers, E. J., et al. (2016). Registered replication report: Strack, Martin, and Stepper (1988). *Perspectives on Psychological Science*, 11(6), 917–928.

<https://doi.org/10.1177/1745691616674458>

Xie, H., Zhou, Z., & Liu, Q. (2018). Null effects of perceptual disfluency on learning outcomes in a text-based educational Context: a meta-analysis. *Educational Psychology Review*, 30(3), 745–771. <https://doi.org/10.1007/s10648-018-9442-x>

Yuan, K. H. & Maxwell, S. (2005). On the post hoc power in testing mean differences. *Journal of Educational and Behavioral Statistics*, 30(2), 141–167.

<https://doi.org/10.3102/10769986030002141>

Figure Captions

Figure 1. Mean meta-analytic estimates in different selection conditions. In A) true effect sizes are fixed ($\tau = 0$), in B) true effect sizes have a standard deviation of $\tau = 0.3$. The dotted black lines indicate optimal estimates that are almost exactly matched when there is no publication bias (not shown).

Figure 2. Type I error rates of all evaluated detection methods in different selection conditions. In A) true effect sizes are fixed ($\tau = 0$), in B) true effect sizes have a standard deviation of $\tau = 0.3$. The numbers denote Type I error rates for TES.

Figure 3. Statistical power of all evaluated detection methods in different selection conditions when true effect size is $\delta = 0$ or $\delta = 0.5$ and heterogeneity is $\tau = 0$ or $\tau = 0.3$.

Figure 1.

Figure 2.

Figure 3.