

# Correcting for bias in psychology: A comparison of meta-analytic methods

Evan C. Carter\*

U.S. Army Research Laboratory, Aberdeen, MD, USA

Felix D. Schönbrodt\*

Ludwig-Maximilians-Universität, Munich, Germany

Will M. Gervais

University of Kentucky, Lexington, KY, USA

Joseph Hilgard

University of Pennsylvania, Philadelphia, PA, USA

Publication bias and questionable research practices in primary research can lead to badly overestimated effects in meta-analysis. Methodologists have proposed a variety of statistical approaches to correct for such overestimation. However, much of this work has not been tailored specifically to psychology, so it is not clear which methods work best for data typically seen in our field. Here, we present a comprehensive simulation study to examine how some of the most promising meta-analytic methods perform on data that might realistically be produced by research in psychology. We created such scenarios by simulating several levels of questionable research practices, publication bias, heterogeneity, and using study sample sizes empirically derived from the literature. Our results clearly indicated that no single meta-analytic method consistently outperformed all others. Therefore, we recommend that meta-analysts in psychology focus on sensitivity analyses—that is, report on a variety of methods, consider the conditions under which these methods fail (as indicated by simulation studies such as ours), and then report how conclusions might change based on which conditions are most plausible. Moreover, given the dependence of meta-analytic methods on untestable assumptions, we strongly recommend that researchers in psychology continue their efforts on improving the primary literature and conducting large-scale, pre-registered replications. We provide detailed results and simulation code at <https://osf.io/rf3ys> and interactive figures at <http://www.shinyapps.org/apps/metaExplorer/>.

**Manuscript accepted for publication in *Advances in Methods and Practices in Psychology*:**

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 Practices in Psychological Science*, 2. doi:10.1177/2515245919847196

This article has undergone heavy copy editing in the final stage. It is not the copy of record, but all results and conclusions are identical to the final published version.

Also check out the interactive app that visualizes all results: <http://shinyapps.org/apps/metaExplorer/>. Fully reproducible R scripts are available at Github: <https://github.com/nicebread/meta-showdown>.

**Keywords:** meta-analysis, publication bias, *p*-hacking, questionable research practices, bias-correction.

Statistical techniques for analyzing the results from a set of studies in aggregate—often called meta-analysis—are popular in psychology and many other scientific disciplines because they provide high-powered tests, the ability to ex-

amine moderators across studies, and precise effect size estimates that are useful for planning future studies and making policy decisions. However, just as the results from individual studies can be made completely misleading by bias (e.g., Simmons, Nelson, & Simonsohn, 2011), so too can meta-analytic results. To address this, researchers have developed statistical techniques designed to identify and correct for bias. Without having a particular preference in any specific method, we present a neutral comparison (Boulesteix, Wil-

---

Correspondence concerning this article should be addressed to Evan Carter, Email: [evan.c.carter@gmail.com](mailto:evan.c.carter@gmail.com). \*These authors contributed equally to this work.

son, & Hapfelmeier, 2017) of how several promising methods perform when applied to simulated data that could have plausibly been produced by research in psychology. Our goal is to help researchers in psychology know what to expect from different methods when conducting meta-analysis in the face of bias.

### Meta-analysis

Meta-analytic techniques involve synthesizing a set of results from studies investigating the same empirical phenomenon (Borenstein, Hedges, Higgins, & Rothstein, 2011). Most often, the results from the individual studies take the form of effect size estimates, and because meta-analyses are usually applied to studies with dependent variables measured on different scales, effect size estimates are typically standardized. The typical goal of a meta-analysis is to produce a single summary estimate of the hypothetical true underlying effect,  $\delta$ , estimated by each effect size in the dataset. This is usually called fixed-effect meta-analysis (Cooper, Hedges, & Valentine, 2009) and can be modeled as  $d_i = \delta + e_i$ , where  $d_i$  is the observed effect size for study  $i$  that differs from the true underlying effect,  $\delta$ , by some amount of sampling error  $e_i$ , which is normally distributed with a mean of 0 and a variance of  $v_i$ . (See the glossary for the meanings of all symbols.)

Another common model known as random-effects meta-analysis (Cooper et al., 2009) holds that each study provides an estimate,  $d_i$ , of a different, related true effect,  $T_i$ —that is,  $d_i = T_i + e_i$ . This approach allows for the possibility that researchers attempting to study the same phenomenon may nonetheless be studying different underlying effects that vary as a function of, for example, different operationalizations of the independent variable or different populations. In this model,  $T_i = \delta + u_i$ , where  $\delta$  is the mean of the true effects estimated by the individual studies and the  $i$ th study's deviation from this mean,  $u_i$ , is normally distributed with a mean of 0 and a variance of  $\tau^2$ . Applying the random-effects model to an observed set of studies provides an estimate of the average true underlying effect,  $\delta$ , and the amount of between-study heterogeneity,  $\tau^2$ . For this paper, we use random-effects meta-analysis as our baseline for “uncorrected” meta-analysis. It should be noted, however, that determining which uncorrected estimator for the average true underlying effect to use is an active area of study itself (Baker & Jackson, 2013; T. Stanley & Doucouliagos, 2015; Schmid, 2017; Rice, Higgins, & Lumley, 2017; Veroniki et al., 2016).

### Bias

The effects being estimated by meta-analysis can be systematically over- or underestimated in the face of bias, which is caused by factors that affect the analysis and reporting of the individual studies in the meta-analytic dataset. We consider two primary sources of meta-analytic bias in our

simulation study: *publication bias* and *questionable research practices*.

Publication bias occurs when the probability of results entering the published record is affected by the results themselves (Rothstein, Sutton, & Borenstein, 2006). For example, if researchers strongly believe that an effect is real and positive, statistically non-significant or negative estimates of that effect may never be submitted for publication or may be rejected by reviewers and editors (Greenwald, 1975; Sterling, Rosenbaum, & Weinkam, 1995; Rothstein et al., 2006; Ferguson & Heene, 2012). In other words, statistically non-significant results, or those results that contradict accepted theory, are left in the “file-drawer.” Since the data set collected by the meta-analyst depends on the availability of studies on the topic of interest, and published data are much easier to find, publication bias can result in a meta-analytic sample that over-represents statistically significant, theory-consistent studies. This can lead to misleading meta-analytic results in terms of inflated estimates of the average true underlying effect and increased false positives during hypothesis testing. And although we do not focus on it here, it is important to note that such bias also affects estimates of heterogeneity in complex, non-linear ways (e.g., Jackson, 2007; Augusteijn et al., 2018).

A related but independent form of bias is the use of questionable research practices (QRPs; also referred to as the undisclosed use of researcher degrees of freedom or “*p*-hacking”). QRPs are said to occur when researchers favor a specific analytic approach (e.g., removing outliers or covariates) from the variety of potential approaches based on the results it yields. These choices may be justifiable, yet simultaneously arbitrary and motivated (Simonsohn, Simmons, & Nelson, 2015b). Like publication bias, this behavior can result in overestimates of the true effect, as analyses that yield significant results are highlighted and analyses that do not yield such results are censored. All bias-correcting methods that have been applied in this study only address publication bias, but not QRPs.

### Simulation studies of bias correction in meta-analysis

Many simulation studies have been conducted to compare the performance of methods that correct for bias in meta-analysis (e.g., Hedges & Vevea, 1996; Moreno et al., 2009; Rucker, Carpenter, & Schwarzer, 2011; T. Stanley & Doucouliagos, 2014; Simonsohn, Nelson, & Simmons, 2014; van Assen, van Aert, & Wicherts, 2015; R. C. van Aert, Wicherts, & van Assen, 2016; McShane, Böckenholt, & Hansen, 2016; T. Stanley, 2017). However, there is very little overlap between these studies in terms of either the methods they examine or the simulated conditions they explore—different simulation studies implement bias differently, sample sizes are drawn from different distributions, and the value and form of the simulated true underlying ef-

fects varies widely. This lack of overlap is not surprising given that there is an effectively infinite number of possible combinations of different conditions to explore and no way of determining which conditions have actually led to any real dataset. In other words, not only is there an inherent dimensionality problem in these simulation studies, but there is also no ground truth. These problems are often not discussed in simulation studies, and indeed, many of the studies cited above—explicitly or implicitly—recommend the use of a single method, despite the fact that each simulation study examines performance in only a limited subset of possible conditions and among a handful of correction methods.

Our goal is *not* to identify a single method that should be used in all situations. Instead, we aim to (1) add to the existing literature by exploring a further set of conditions that may plausibly represent real data from research in psychology, (2) compare a larger set of meta-analytic methods that, to our knowledge, have yet to be directly compared, and (3) discuss how our results can facilitate sensitivity analysis in meta-analysis.<sup>1</sup>

### Disclosures

R (R Core Team, 2016) scripts are available at Github <https://github.com/nicebread/meta-showdown>. Furthermore, we provide interactive figures and tables which allow a detailed exploration of the results (<http://www.shinyapps.org/apps/metaExplorer/>). Supplemental material, which includes a comprehensive presentation of our results, is available at <https://osf.io/rf3ys>. We declare that we have no conflicts of interest with respect to the authorship or the publication of this article. We report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study.

ECC, FDS, and JH developed code and managed the simulation study. ECC, FDS, JH, and WMG planned the project and wrote the manuscript. We would like to acknowledge Tyler Yost for helpful comments on an earlier version of the simulation code and manuscript. A preprint of this paper is available at <https://osf.io/rf3ys>.

### Methods

#### Simulation

We simulated the number of meta-analyzed studies,  $k$ , as one of four values (10, 30, 60, 100). All simulated individual studies had a two-group experimental design, so all effect sizes took the form of a standardized mean difference, Cohen's  $d$ . Notably, there is reason to think that this may be the most commonly used effect size measure in psychology (see Table S1 in Fanelli, Costas, & Ioannidis, 2017). Cohen's  $d$  is an estimate of the true underlying effect  $\delta$ , which we chose to simulate as taking one of four values (0, 0.2, 0.5, 0.8) corre-

sponding to the null hypothesis and Cohen's "rule of thumb" values for small, medium, and large effects, respectively.

**Heterogeneity.** As mentioned, the true underlying effect  $\delta$  may vary across studies based on the heterogeneity parameter  $\tau$ . We simulated three values for  $\tau$  (0, 0.2, 0.4)<sup>2</sup> that may plausibly represent research in psychology: In an analysis of 187 meta-analyses using standardized mean differences published in *Psychological Bulletin* from 1990–2013 (van Erp, Verhagen, Grasman, & Wagenmakers, 2017), 50% of all estimates of  $\tau$  were smaller than 0.2, and 80% were smaller than 0.4.<sup>3</sup>

**Study-level data.** Independent samples were randomly generated for the control and experimental group, where observations in the control group were drawn from a normal distribution with mean of 0 and standard deviation of 1, and observations in the experimental group were drawn from a normal distribution with a mean of  $T_i$  and standard deviation of 1.  $T_i$  was defined as the sum of  $\delta$  and  $u_i$ , where  $u_i$  was drawn from a normal distribution with mean 0 and standard deviation  $\tau$ . Note that  $T_i$ , therefore, represented a study-specific true effect that varied randomly if  $\tau$  was greater than 0. Cohen's  $d$  and the variance  $v$  of the effect size were then calculated (see glossary) and a two-tailed independent-samples  $t$ -test was applied to generate a  $t$ -value and  $p$ -value.

Simulated sample sizes were based on an empirical distribution of sample sizes (Marszalek, Barber, Kohlhart, & Holmes, 2011; Marszalek, 2011), which provides per-group sample sizes from 1225 studies from four journals (*Journal of Abnormal Psychology*, *Journal of Applied Psychology*, *Journal of Experimental Psychology: Human Perception and Performance*, and *Developmental Psychology*) from the years 1995 and 2006. After removing sample sizes smaller than 5, the strongly right-skewed per-group sample size distribution had a median of 23 (25% quantile: 14, 75% quantile: 50). We found that an inverse gamma distribution clearly showed the best fit according to the log-likelihood (compared to a negative binomial, log-normal, gamma, or Weibull distribution). For our simulations, we used a truncated inverse gamma distribution (truncated at  $n = 5$  and  $n =$

<sup>1</sup>It is worth noting that our original intent with this study was in fact to identify, if possible, a single best method across many conditions. Further consideration and helpful comments from our peers changed our minds on this.

<sup>2</sup>In terms of the heterogeneity metric  $I^2$ , these values of  $\tau = 0.2$  and 0.4, in combination with the specific primary sample sizes we simulated, are approximately equal to the descriptors proposed by Pigott (2012) for "medium" ( $I^2 = 50%$ ) and "large" ( $I^2 = 75%$ ) heterogeneity: random-effects meta-analysis of the unbiased data (no publication bias, no QRPs, aggregated over  $k$  and  $\delta$ ), yields an average observed  $I^2$  of 46% (SD = 17%) when  $\tau = 0.2$  and 77% (SD = 10%) when  $\tau = 0.4$ .

<sup>3</sup>One should keep in mind that the estimates of  $\tau$  reported by van Erp et al. van Erp et al. (2017) may be over- or under-estimates as a result of bias (Augusteijn et al., 2018).




Figure 1. Comparison of the empirical per-group sample size distribution (histogram) from Marszalek, Barber, Kohlhart, and Holmes (2011) with the best fitting inverse gamma curve (continuous line). X-axis truncated at  $n = 400$  for better visibility. Figure available at <https://osf.io/rf3ys>, under a CC-BY4.0 license.

1905, which was the largest observed per-group sample size in the Marszalek, 2011 data set) with shape = 1.153 and scale = 0.046 to sample our per-group sample sizes. Figure 1 compares the empirical per-group sample size distribution with the best fitting curve.

**Publication bias.** For the simulation of publication bias, we used two censoring functions (“medium publication bias” and “high publication bias”) which compute a probability that a study is published, based on the study’s one-tailed  $p$ -value. If the effect is in the “correct” direction, both functions return a 100% probability of publication when  $p_{\text{onetailed}} < .025$ . “Marginally” significant effects with  $.025 \leq p_{\text{onetailed}} < .05$  are published with an exponentially decreasing probability that reaches a constant level of 20% (medium publication bias) or 5% (high publication bias) for  $.5 > p_{\text{onetailed}} \geq .05$ . If the effect is in the “wrong” direction ( $p_{\text{onetailed}} \geq .5$ ), the censoring function for medium publication bias returns a 50% probability of publication when  $p_{\text{onetailed}} \geq .9995$  (high publication bias: 20%). This probability is exponentially decreasing for  $.9995 > p_{\text{onetailed}} \geq .995$ , until it reaches a constant level of 5% (medium) or 0% (high) for  $.995 > p_{\text{onetailed}} \geq .5$  (see Figure 2). In our simulation, a random Bernoulli draw using the probability computed by these censoring functions determined whether a study was “published”.<sup>4</sup> Studies were continually simulated until the

target number of  $k$  studies had been reached. In the “no publication bias” conditions, all studies were included regardless of  $p$ -value or sign of effect.

To our knowledge, this specific implementation of publication bias is comparable to, but different from, those used in previous simulation studies (e.g., Bayarri & DeGroot, 1991; Guan & Vandekerckhove, 2016). Our primary reason for choosing this approach was that we did not want our publication bias functions to exactly match those assumed by any of the bias-correcting methods being examined (Iyengar & Greenhouse, 1988; McShane et al., 2016, i.e.), as this might give an overly optimistic assessment of performance by those methods (Simonsohn, Simmons, & Nelson, 2017b). Furthermore, it is an open question as to whether methods that have been previously tested on data produced with more straightforward publication bias functions would be robust to our more nuanced implementation.


**Questionable Research Practices (QRPs).** We studied four forms of QRPs: (1) Optional removal of outliers, (2) optional selection between two dependent variables, (3) optional use of moderators, and (4) optional stopping. Each data set that would have QRPs applied to it was designed to simulate a study with a two-by-two (experimental group vs. the control group; level one of the moderator vs. level two of the moderator) design and two dependent variables. Each dependent variable was measured across  $n$  observations. The moderator divided the simulated data set in half in a way that was independent of the dependent variable (i.e., the moderator had no main effect on the dependent variable) and the treatment (i.e., no collinearity between moderator and treatment). The two dependent variables were correlated at  $r = 0.20$ .

QRPs were applied so as to simulate the behavior of a researcher fishing for statistical significance. For the “maximal QRP treatment”, the simulated researcher first tested the main effect of experimental manipulation on the first dependent variable. If this effect was not statistically significant and positive, the simulated researcher removed outliers (defined as observations with an absolute value  $z$ -score greater than 2).<sup>5</sup> If this second test was not positive and significant, the simulated researcher moved to the second dependent variable and repeated the above steps. If no positive and significant effect was found, the researcher moved back to the first dependent variable and tested for an interaction effect between the experimental manipulation and the moderator. In the presence of a significant interaction, the researcher compared the experimental and control groups in only the subgroup defined by the first level of the moderator. This ex-


<sup>4</sup>If the probability of publication was 0.25, for example, we drew one random sample from a Bernoulli distribution where  $p = 0.25$ . If the sample value was 1, the simulated result was “published.”

<sup>5</sup>We also could have deleted only outliers in one direction, which would have made the  $p$ -hacking more efficient.

(A) No publication bias



(B) Medium publication bias



(C) Strong publication bias

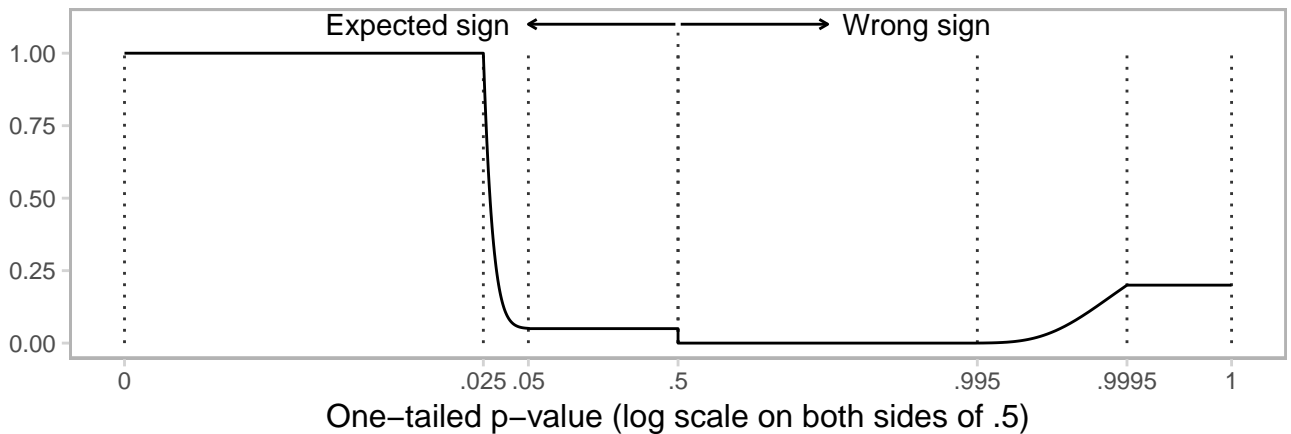


Figure 2. Censoring functions which return the probability of publication given a one-tailed  $p$ -value. The x-axis has a logarithmic scale on both sides of  $p_{onetailed} = .5$  to increase the visibility of the function at the high and low ends of the scale. Figure available at <https://osf.io/rf3ys>, under a CC-BY4.0 license.



amination was conducted first with and then without outliers. If no positive and significant effect was found, the second subgroup defined by the second level of the moderator was assessed in the same way. In the absence of a positive and significant effect at the second level, the researcher moved to the second dependent variable and repeated the same procedure.

Additionally, the simulated researcher could opt to collect some additional amount of data (see below). After each additional collection effort, the QRPs described above were repeated. Thus, for each data collection effort, simulated researchers could potentially apply 12 comparisons. If none of these analyses produced a positive and significant effect, the first test (experimental vs. control on the first dependent variable with outliers untouched and no division by the moderator) was taken as the final result.

To organize different levels of severity, we created three types of *individual QRP strategies* a simulated researcher could adopt: (1) pure (no use of QRPs); (2) moderate (optional dependent variables and the addition of three observations per cell for up to three data collection efforts); and (3) strong (use of optional outliers, optional dependent variables, optional moderators, and the addition of three observations per cell for up to five data collection efforts). Given the sample sizes of our simulated primary studies, the moderate QRP strategy results in an inflated false-positive rate (computed in conditions without heterogeneity, without publication bias, and only counting directionally consistent results) of 9%, the strong strategy in 27%. Note that more aggressive *p*-hacking beyond our “strong” setting is easily possible, for example by examining even more dependent variables, doing only directional outlier exclusion, etc. For example, Simmons et al. (2011) report a false positive rate of 61% by combining certain types of *p*-hacking.

As it is unlikely that every researcher in a field applies QRPs in the same fashion, we defined three *QRP environments* to describe possible prototypical research fields with a specific severity of QRP application. Each QRP environment was characterized by a mixture of simulated researchers with individual QRP strategies: (1) *none* (100% of simulated researchers adopted the pure strategy); (2) *medium* (30% pure, 50% moderate, and 20% strong), leading to a 11% false-positive rate; and (3) *high* (10% pure, 40% moderate, and 50% strong), leading to a 17% false-positive rate.

Not all QRP types have the same distorting impact on a meta-analysis. Furthermore, some QRP types lead bias-correcting techniques to overestimate the true effect, while other types lead to an underestimation (R. C. van Aert et al., 2016). Our goal was not to investigate the differential impact of distinct QRP types on bias correction, but rather to investigate some combinations of QRPs that may be plausible in real settings (John, Loewenstein, & Prelec, 2012). As there are infinite possibilities of how QRP can be done, and

Table 1  
*Simulation parameters*

Experimental factors	Levels
True underlying effect ( $\delta$ )	0, 0.2, 0.5, 0.8
Between-study heterogeneity ( $\tau$ )	0, 0.2, 0.4
Number of studies in the meta-analytic sample ( $k$ )	10, 30, 60, 100
Publication bias ( $PB$ )	None, medium, strong
QRP environment ( $QRP$ )	None, medium, high

an infinite number of ways how these individual researcher strategies can be combined to QRP environments, our results are best considered as a sensitivity analyses that explores the effect of a range of three plausible QRP environments on meta-analytic results. This does not necessarily generalize to other implementation of QRPs.

**Design.** To summarize, we simulated data for 432 unique combinations of five fully-crossed factors (Table 1). We simulated 1,000 meta-analyses for each of the 432 conditions. For a random selection of conditions we also computed 10,000 simulations and computed the Monte Carlo simulation error (Koehler, Brown, & Haneuse, 2009). These comparisons clearly demonstrated that 1,000 replications lead to sufficiently stable estimates (see supplementary material, <https://osf.io/rf3ys>).

### Performance metrics

For the hypothesis test of whether the estimate provided by each meta-analytic method differs from zero, we evaluated the false positive (Type I error) rate at  $\delta = 0$  and the true positive rate (i.e., the statistical power) at  $\delta = 0.2, 0.5,$  and  $0.8$ .

Following the recommendations of Burton, Altman, Royston, and Holder (2006), we measured the bias-adjusting performance of each method in terms of mean error (ME), root mean squared error (RMSE), and 95% coverage probability (see glossary).

### Meta-analytic methods

We examined the performance of seven estimators. Further details on our specific implementations are available in the supplemental material <https://osf.io/rf3ys>.

**Random-effects meta-analysis (RE).** We applied the random-effects meta-analysis as described above using the *metafor* package in R (Viechtbauer, 2010). This approach makes no adjustment for publication bias or QRPs. We used the REML method for estimating between-study variance.

**Trim-and-fill (TF).** Trim-and-fill (Duval & Tweedie, 2000) was introduced as a diagnostic test for publication bias based on funnel plot asymmetry (a scatter plot of effect size estimates against the standard error of those estimates). Publication bias introduces clear rightward asymmetry in a funnel plot (see supplemental material) by censoring non-significant and negative observations. Trim-and-fill iteratively removes (i.e., trims) observations from one side of the funnel plot until a criterion for symmetry is met and then fills observations back into the funnel plot along with imputed observations reflected about the mean. Standard meta-analytic methods can then be applied to a data set including both observed and imputed studies.

Several previous simulation studies suggest that, although trim-and-fill can correct for bias in some cases, not only do other methods tend to outperform it, but it generally fails with increasing heterogeneity (e.g., Terrin, Schmid, Lau, & Olkin, 2003; Peters, Sutton, Jones, Abrams, & Rushton, 2007; Moreno et al., 2009; Idris & Ruzni, 2012; Simonsohn et al., 2014). For example, Terrin et al. (2003) examined the coverage probability of trim-and-fill both with and without heterogeneity. They observed that trim-and-fill coverage decreased as heterogeneity increased, primarily because the method imputed studies that were not missing when effect sizes from large studies (i.e., those near the top of the funnel plot) were far from the average overall effect. In addition to unnecessarily correcting for bias, more recently it has been suggested that trim-and-fill also does not correct enough for bias that does exist (Simonsohn et al., 2014; Simonsohn, Simmons, & Nelson, 2017a; van Assen et al., 2015).

Overall, no conclusion has been reached on the best way to implement trim-and-fill, as performance can vary widely with different versions of the algorithm in different conditions (Peters et al., 2007; Moreno et al., 2009). Therefore, we use the default algorithm provided by the *metafor* package. Notably, trim-and-fill did not always converge. Across all conditions, it returned a valid estimate in 95% of data-sets. Non-convergence mostly happened at  $k \geq 60$  with strong publication bias.

**Weighted average of the adequately powered studies (WAAP-WLS).** T. Stanley and Doucouliagos (2017) proposed the use of an intercept-only weighted least squares (WLS) meta-regression estimator as a replacement for the naive fixed-effect and random-effects meta-analytic models. Simulation studies (T. Stanley & Doucouliagos, 2017; T. Stanley, 2017; T. Stanley, Doucouliagos, & Ioannidis, 2017) suggested that the WLS estimator performed on par with the fixed-effect and random-effects models when the assumptions underlying those models were true, but outperformed both of them when they were violated (e.g., in the presence of publication bias).

Scholars have suggested extending this WLS estimator to reduce the impact of potential publication bias (Ioannidis,

Stanley, & Doucouliagos, 2017; T. Stanley et al., 2017). In this extension, one first performs a WLS meta-analysis on all primary studies to get an (potentially biased) estimate of the true underlying effect. Then one performs a second WLS meta-analysis on only those studies that have 80% statistical power to detect this estimated effect size—that is, a weighted average of adequately powered (WAAP) studies. This approach attempts to avoid bias by discarding underpowered studies, which must overestimate the true effect to find statistical significance. If there are no or only one adequately powered study in the dataset, the WLS estimate for the entire dataset is used as a fallback solution. This conditional estimator, called WAAP-WLS, applies WAAP when there are at least two adequately-powered studies and WLS otherwise.

Across all conditions, WAAP-WLS returned 77% WAAP and 23% WLS estimates. In small- $k$ , small- $\delta$  conditions, there were not enough adequately powered studies, and 100% of estimates used WLS. In large- $k$ , large- $\delta$  conditions, 100% of estimates used WAAP.

Previous simulation studies suggest that WAAP-WLS is comparable to WLS, standard fixed-effects meta-analysis, and random-effects meta-analysis in the absence of heterogeneity and publication bias; however, as those conditions change, WAAP-WLS has outperformed both WLS and standard meta-analysis (T. Stanley et al., 2017). The same simulation study suggested, however, that WAAP-WLS in turn has been outperformed by PET-PEESE (see below) in terms of efficiency and overall bias.

**$p$ -curve.** A  $p$ -curve is the distribution of all statistically significant  $p$ -values from the set of studies of interest (i.e.,  $ps < 0.05$ ; Simonsohn et al., 2014). The shape of the  $p$ -curve is a function of the statistical power of the studies, which is itself a function of the sample sizes and the true effect size. When studies have no statistical power (i.e., when the null is true) then the distribution of significant, independent  $p$ -values is uniform between .00 and .05. With increasing power, this distribution becomes increasingly right-skewed. Because the degree of right skew is a function of the average study power,  $p$ -curve can use the degree of right-skew to (a) test the absence of a real effect and (b) estimate the average effect size corrected for publication bias.

Previous simulations by Simonsohn et al. (2014) demonstrated that some typical QRPs cause  $p$ -curve to underestimate the true effect size. Later work by R. C. van Aert et al. (2016), however, suggested that both upward and downward bias can be present in  $p$ -curve depending on the specific type of QRPs. Additional work demonstrated that  $p$ -curve overestimates the average true underlying effect when there is heterogeneity (Simonsohn et al., 2014; R. C. van Aert et al., 2016).

Critically, Simonsohn et al. (2014) interpret the  $p$ -curve estimate differently from how we use it here: “It is the average effect size one expects to get if one were to rerun all

studies included in the *p*-curve” (p. 667; see also Simmons, Nelson, & Simonsohn, 2018). In our view, however, meta-analysts generally aim to recover the mean effect size of all conducted studies—indeed, that is the purpose of the other estimators we examine. For consistency with all other estimators, we therefore interpret *p*-curve in this fashion; however, in the supplemental material, we also assess *p*-curve’s ability to recover the average true effect size of all studies submitted.

We implemented *p*-curve as recommended by Simonsohn et al. (2014) with the following settings. Only statistically significant and directionally consistent studies were submitted to the *p*-curve. If significant but negative studies were present, they were discarded. Consequently, when no significant studies in the positive direction were in a set, *p*-curve does not return an estimate (0.8% of all simulations). Across all conditions, *p*-curve returned an estimate in 99.2% of all simulated data-sets. Not surprisingly, the method failed to produce an estimate almost exclusively at  $k = 10$ ,  $\delta = 0$ , and no publication bias.

In some cases, *p*-curve can return an estimate with a negative sign, even though all studies entered had positive signs. In their web-app <http://www.p-curve.com>, Simonsohn, Nelson, and Simmons suggest one should only interpret non-negative effect size estimates from *p*-curve. When *p*-curve returns a negative effect size estimate, they suggest setting the negative estimate to zero, which we did in our study.<sup>6</sup>

Concerning the test for the presence of an effect, we relied on the test for evidential value (i.e., the test for right skewness) for the full *p*-curve (Simonsohn, Simmons, & Nelson, 2015a), which is conceptually equivalent—but not statistically equivalent—to a test for  $\mu > 0$ . Furthermore, *p*-curve estimation does not provide confidence intervals, so we could not assess this aspect of estimation (see below).

***p*-uniform.** Like *p*-curve, the *p*-uniform method also considers only the statistically-significant results. It is based on the idea that the distribution of *p*-values is uniform *conditional on the population effect size* (van Assen et al., 2015). Hence, it focuses on the *p*-value distribution under the alternative hypothesis, and it yields a fixed-effects estimate of the true effect by finding the value  $d^*$  which makes the conditional distribution of *p*-values as uniform as possible.

*p*-uniform provides a hypothesis test, an estimate of the bias-corrected effect size, and a confidence interval around that estimate. Computationally, *p*-curve and *p*-uniform only differ by an alternative implementation of the estimation algorithm, so in general *p*-curve and *p*-uniform are expected to have similar strengths and weaknesses (McShane et al., 2016). For the computation of the *p*-uniform estimate, we used the Irwin-Hall estimator as implemented in the *puniform* package (van Aert, 2017) and recommended by R. C. van Aert et al. (2016). We also followed the recommendation of R. C. van Aert et al. (2016) and set the estimate to

zero if the average of all significant *p*-values was larger than .025, because under a positive true effect the average *p*-value is lower than .025.<sup>7</sup>

Just as *p*-curve, *p*-uniform does not return an estimate without significant studies in the positive direction (0.8% of all simulations). Across all conditions, *p*-uniform returned an estimate in 99.2% of all simulated data-sets, where in 10.4% of all computations the estimate was replaced by zero.

**PET, PEESE, and PET-PEESE.** The precision-effect test (PET; T. Stanley & Doucouliagos, 2014) is a meta-regression approach to adjusting for small-study effects (see supplemental material). Small-study effects are said to exist when the observed effect size gets larger as the standard error grows (i.e., the sample size shrinks; see also the closely related Egger’s test for publication bias, Egger, Smith, Schneider, & Minder, 1997). One cause of this pattern is publication bias, although other benign causes also exist (Sterne, Gavaghan, & Egger, 2000; T. Stanley & Doucouliagos, 2014). PET fits a linear regression line to this relationship, then extrapolates to estimate the effect size of a hypothetical study with a standard error of zero (i.e., a study with infinite sample size and perfect measurement). The result is an estimate of the true underlying effect that has been corrected for publication bias and other small-study effects. Of course, if small-study effects have many benign causes, this may represent a substantial over-correction.

PET is a weighted-least-squares regression where effect size is regressed on its standard error:  $d_i = b_0 + b_1 se_i + e_i$ , where  $b_0$  and  $b_1$  are the intercept and slope terms describing the linear relationship between the  $i$ th effect size estimate  $d_i$  and its associated standard error  $se_i$ . The regression model is weighted by the inverse of the variance (i.e., the squared standard errors) of the effect size estimates. The intercept  $b_0$  represents the estimated effect size when the standard error is zero.

The precision-effect estimate with standard error (PEESE) (T. Stanley & Doucouliagos, 2014) fits a quadratic relationship between effect size and standard error. The rationale is that if there is some true effect, low-precision studies are poorly powered and publishable only when the effect is badly overestimated. On the other hand, high-precision studies will be well-powered and routinely publishable without such overestimation. Thus, publication bias (and the observed small-study effect) is expected to be stronger when the standard error is larger. A quadratic relationship can model such differences in bias across standard errors. PEESE is the weighted-least-squares regression model where ef-

<sup>6</sup>Technically, we constrained the numerical optimizer to values  $\geq 0$ . In 10.3% of all cases the estimate was  $< 0.0001$ .

<sup>7</sup>Although in this special case no *p*-value and no CI are provided, we treated these cases as “do not reject  $H_0$ ”. Hence, the special case is utilized in the computation of the false-positive error rate, but not in the coverage probability metric.



fect size is regressed on the square of the standard error:  $d_i = b_0 + b_1 se_i^2 + e_i$ . As in PET, the weights are the inverse of the variances and the intercept is interpreted as an estimate of the true underlying effect that is uninfluenced by small-study effects.

Simulation studies suggest that PET outperforms PEESE when the true underlying effect is zero, whereas PEESE outperforms PET when the true underlying effect is non-zero (T. Stanley & Doucouliagos, 2014). In an attempt to offset the opposite biases of PET and PEESE, T. Stanley and Doucouliagos (2014) proposed the conditional estimator PET-PEESE. PET-PEESE considers the statistical significance of the PET estimate to decide whether PET or PEESE is taken as the final estimate. When the estimate from PET is statistically non-significant (i.e., the estimated true effect is not distinguishable from zero) in a one-sided test at  $\alpha = 5\%$ , the PET estimate is taken. In contrast, when the estimate from PET is statistically significant, the PEESE estimate is used as the value for the conditional PET-PEESE estimator. As suggested by T. Stanley and Doucouliagos (2014), we will only focus on the conditional PET-PEESE estimator, but report the performance of the separate PET and PEESE estimators in the supplemental material.

Although initial simulation results indicated promising performance by PET-PEESE (T. Stanley & Doucouliagos, 2014), two later simulations detailed some weaknesses. One observed that the standard random-effects meta-analysis estimator outperformed PET and PEESE in some ways, such as greater estimation efficiency (lower mean squared error) when heterogeneity was present (Reed, Florax, & Poot, 2015). Another found that PET-PEESE performance was unacceptable under conditions that seem common in psychology – a small number of studies, small samples across all studies, and high heterogeneity (T. Stanley, 2017).

**Selection model (3PSM).** Selection model approaches to mitigation of bias in meta-analysis have been in use for some time (Hedges, 1984; Iyengar & Greenhouse, 1988; Hedges & Vevea, 1996). Here, we employ the three-parameter selection model (3PSM) as developed by Iyengar and Greenhouse (1988) and recently discussed by McShane et al. (2016). This model’s three parameters represent the average true underlying effect  $\delta$ , the heterogeneity of the random effect sizes  $\tau^2$ , and the probability  $p_1$  that a non-significant effect enters the literature. The last parameter,  $p_1$ , is modeled by a step function with a single cutpoint at  $p = .025$  (one-tailed), which corresponds to a two-tailed  $p$ -value of .05. This cutpoint divides the range of possible  $p$ -values into two bins: significant and non-significant. The three parameters are estimated using maximum likelihood.

We implemented 3PSM using the default function in the *weightr* package (Coburn & Vevea, 2017). If no  $p$ -value is present in one of the bins, the probability  $p_1$  cannot be estimated. In this case, the *weightr* package uses a plu-

gin value of .01, which makes it possible to estimate the model (Veeva & Woods, 2005). However, even with this plugin value some models could not be estimated due to non-convergence. Across all conditions, 3PSM returned an estimate in 91.5% of all simulated data-sets. Estimation mostly failed at  $k = 100$ ,  $\delta = 0.2$  and at least medium publication bias and/or QRPs.

Several previous simulation studies of selection models exist (e.g., Terrin et al., 2003; Hedges & Vevea, 1996). However, to our knowledge, only one examines the specific method we implement here: McShane et al. (2016) compared 3PSM to  $p$ -uniform and  $p$ -curve, both of which can be understood as single parameter selection models (i.e., only  $\mu$  is estimated, publication bias is set to 100% and heterogeneity is set to 0). In that study, 3PSM clearly outperformed both  $p$ -uniform and  $p$ -curve in terms of estimation and hypothesis testing in conditions where  $\delta \leq 0.30$  and  $\tau > 0$  or incomplete bias allowed some non-significant studies to be published.

## Results

### Presentation of results

We simulated 1,000 meta-analyses under 432 unique conditions (Table 1) and analyzed each with seven different meta-analytic methods. Here, we avoid an exhaustive presentation of the results and give instead a more focused report. However, all of our findings, including information on convergence probabilities and exact values for ME, RMSE, and coverage probabilities for all conditions, are available in supplemental material (<https://osf.io/rf3ys>). We also provide several interactive figures and tables (<http://www.shinyapps.org/apps/metaExplorer/>) as a means of encouraging researchers to explore combinations of conditions that they find to be particularly relevant to their own work.

In the following sections we provide figures only for conditions in which  $\delta = 0$  (i.e.,  $H_0$  or “under the null”) or  $\delta = 0.5$  (i.e.,  $H_1$  or “under the alternative”) and how this behavior changes as heterogeneity is added ( $\tau = 0.2$ ). Furthermore, we focus only on  $k = 10$  and  $k = 60$  and for cases with no publication bias or strong publication bias. The influence of QRPs is discussed separately at the end of this section.

Rather than providing exact values for ME, RMSE, and coverage, we display the distributions of effect size estimates in terms of means and the 95% quantile ranges in Figure 4. In our view, these values are more intuitive, and they allow a visual assessment of ME (the means) and RMSE (the range). Figure 3 shows both the false-positive rates (when  $\delta = 0$ ) and the statistical power (when  $\delta = 0.50$ ) of each method.

Some methods did not always return an estimate. This includes both intended cases (i.e., when no significant, directionally consistent studies are available for  $p$ -curve and  $p$ -uniform), as well as failures to produce an estimate, for example due to non-convergence. Consequently, we report

the summaries of all computations that did return an estimate, but readers should be aware that this implies a conditional interpretation: The reported ME, RMSE, and error rates are conditional on the method providing a result. If a method performs well when it returns an estimate, but does not return an estimate the majority of the time, this should be taken in to consideration when comparing that method to others. In Figures 3 and 4, symbols indicate if a method did not return an estimate in more than 25%, 50% or 75% of the 1000 simulation runs. Table 1 in the supplemental material reports the exact percentages of returned estimates for each method in each condition.

### No publication bias, no QRPs

**Type I error rate.** Under the null, most methods had appropriate Type I error rates, although error rates for 3PSM and  $p$ -uniform were more conservative (0-3%) than the nominal alpha rate.  $p$ -curve and  $p$ -uniform sometimes failed to provide an estimate due to the scarcity of statistically significant results, especially when  $k$  was small.

The addition of heterogeneity lead to a small increase (<10%) in Type I error rates for random-effects meta-analysis and trim-and-fill. Type I error rates raised moderately (10-30%) for WAAP-WLS and PET-PEESE. Type I error rates for  $p$ -curve and  $p$ -uniform raised considerably (8-56%). Heterogeneity did not increase Type I error rates for 3PSM, which remained excessively conservative.

**Power.** At  $k = 10$ , random-effects meta-analysis offered the best power of all methods, followed closely by trim-and-fill and 3PSM. Power was poorer for  $p$ -curve, PET-PEESE, and WAAP-WLS, and the power of  $p$ -uniform was markedly poorer in small  $k$ . Despite this, all methods had 100% power by  $k = 60$ .

The addition of heterogeneity at  $k = 10$  slightly increased the power of  $p$ -curve and  $p$ -uniform and slightly reduced the power of trim-and-fill. The power of 3PSM, WAAP-WLS, and PET-PEESE fell by more, 5–15 percentage points. Despite this, all methods had >95% power at  $k = 60$ .

**ME.** All methods were generally unbiased, with some exceptions. At  $k = 10$ , trim-and-fill and PET-PEESE demonstrated slight downward bias (about -0.02); these downward biases were mitigated by larger  $k$ .  $p$ -curve and  $p$ -uniform both exhibited upward bias under the null. For  $p$ -uniform, the mean estimate was 0.30 at  $k = 10$  and 0.20 at  $k = 60$ . For  $p$ -curve, the mean estimate was 0.89 at  $k = 10$  and 0.55 at  $k = 60$ . However, both were unbiased when the null was false. This upward bias is likely caused by these methods' truncation from below at zero — there are no negative underestimates to cancel out positive overestimates.

Heterogeneity had little effect on the bias of random effects, WAAP-WLS, and 3PSM. It added some upward bias to  $p$ -curve and  $p$ -uniform. It added a very slight downward bias to trim-and-fill and PET-PEESE.

**RMSE.** Random-effects meta-analysis provided the most efficient results and smallest RMSE. RMSE was slightly greater for trim-and-fill and WAAP-WLS, moderately greater for 3PSM and PET-PEESE, and noticeably greater for  $p$ -uniform and  $p$ -curve. RMSE was particularly poorer for  $p$ -curve and  $p$ -uniform when the null was true, consistent with their upward bias (and also potentially because these methods use only statistically significant results, of which there are fewer). 3PSM also had relatively high RMSE when the null was true and  $k$  was small.

Heterogeneity only slightly increased the RMSE of random effects. Heterogeneity caused slightly larger increases in RMSE for trim-and-fill, WAAP-WLS, 3PSM, and PET-PEESE. Heterogeneity tended to slightly increase the RMSE of  $p$ -curve and  $p$ -uniform, presumably by causing bias. However, when the null was true or  $k$  was small, heterogeneity improved the RMSE of these methods, presumably by increasing the number of significant results to draw upon.

**95% CI coverage.** Coverage rates of random-effects meta analysis, PET-PEESE, and trim-and-fill were ideal at 95%. WAAP-WLS,  $p$ -uniform, and 3PSM also had approximately correct coverage rates ( $\pm 2\%$ ).

Heterogeneity caused slightly poorer coverage for all methods, particularly  $p$ -uniform. For random-effects meta-analysis and 3PSM, greater  $k$  helped to restore coverage rates. For PET-PEESE,  $p$ -uniform, and WAAP-WLS, greater  $k$  exacerbated the undercoverage. (Recall that  $p$ -curve does not give a confidence interval.)

### Strong publication bias, no QRPs

In the face of strong publication bias, sets of meta-analyzed results often consisted of only significant results, especially when  $k$  was small.

**Type I error.** Random-effects meta-analysis suffered from false-positive rates of 98% and higher. Trim-and-fill had slightly lower, but still unacceptable Type I error rates in excess of 70%. WAAP-WLS had poor Type I error rates at low  $k$  (>85%), but unlike random-effects and trim-and-fill methods, Type I error rates decreased with increasing  $k$  (45% by  $k = 60$ ). 3PSM had lower Type I error rates (31%) at  $k = 10$ , but these errors increased with increasing  $k$  (82% by  $k = 60$ ).  $p$ -curve,  $p$ -uniform, and PET-PEESE had approximately conservative Type I error rates ranging from 3-10%.

The addition of heterogeneity slightly reduced the Type I error rate for random-effects, but error rates still approached 100% with increasing  $k$ . Trim-and-fill was generally not affected. Heterogeneity reduced Type I error in WAAP-WLS when  $k$  was small but increased it at large  $k$ ; error rates were still quite high (50+%). Heterogeneity substantially increased the Type I error rates of  $p$ -curve and  $p$ -uniform, leading to error rates of 40+% at  $k = 10$  and 98+% at  $k = 60$ . Heterogeneity also substantially increased the Type I error rate of PET-PEESE (+33 percentage points or more). Het-







erogeneity increased 3PSM's type I error rates at low  $k$  (+14 percentage points) but reduced it substantially at higher  $k$  (-48 percentage points).

**Power.** Random-effects meta-analysis and trim-and-fill had approximately 100% power at all levels of  $k$ . Other methods had slightly poorer, but still good (89+%) power, even at  $k = 10$ . Listed in order of ascending power, they are: PET-PEESE,  $p$ -uniform, WAAP-WLS, 3PSM, and  $p$ -curve. Note that "good power" in this case is accompanied by high type I error—that is, the null is frequently rejected because methods overestimate the true underlying effect.

Heterogeneity had little influence on power. When  $k = 10$ , it reduced the power of PET-PEESE by 13 percentage points, WAAP-WLS by 7 percentage points, and 3PSM by 5 percentage points. It slightly increased the power of  $p$ -curve (1 percentage points) and  $p$ -uniform (4 percentage points). Power of all methods still approximated 100% at  $k = 60$ .

**ME.** Under the null, many methods were badly biased upwards. Random effects meta-analysis estimated the null effect as about 0.33. Trim-and-fill did not substantially reduce this bias, estimating the effect as about 0.22. WAAP-WLS was quite biased at small  $k$  (0.27) and became less biased with greater  $k$  (0.13). At  $k = 10$ ,  $p$ -curve,  $p$ -uniform, and 3PSM overestimated null results ( $\sim 0.15$ ), but less so than other methods, and the bias decreased with increasing  $k$  (to 0.05-0.11), with the strongest benefits observed with  $p$ -curve and  $p$ -uniform. PET-PEESE was unbiased.

When the null was false, random effects still slightly overestimated the true effect ( $\sim 0.05$ ). Trim-and-fill and WAAP-WLS both reduced this slight bias, yielding estimates that were only very slightly biased.  $p$ -curve,  $p$ -uniform, and 3PSM were unbiased in estimating the true effect. PET-PEESE tended to slightly underestimate the true effect, although this tendency was ameliorated with increasing  $k$  (-0.06 at  $k = 10$ , -0.02 at  $k = 60$ ).

Adding heterogeneity to a true effect of zero tended to create upward bias. Random-effects meta-analysis, trim-and-fill, WAAP-WLS, and PET-PEESE all demonstrated slight increases in upward bias under heterogeneity.  $p$ -curve and  $p$ -uniform demonstrated substantial upward bias under heterogeneity ( $\sim 0.25$ ). In contrast, 3PSM did not demonstrate greater bias under heterogeneity.

Adding heterogeneity to a true nonzero effect also tended to move estimates slightly upwards. This caused slightly greater upward bias in random effects, trim and fill, and WAAP-WLS; slightly reduced downward bias in PET-PEESE; and created slight upward bias in  $p$ -curve and  $p$ -uniform ( $\sim 0.06$ ). 3PSM remained unbiased.

**RMSE.** Given the null, RMSE was considerably elevated for random effects. All adjustments lead to some improvement in RMSE, with the exception of  $p$ -curve when  $k = 10$ . WAAP-WLS and trim and fill lead to modest improvements, and  $p$ -uniform,  $p$ -curve, 3PSM, and PET-PEESE lead

to successively greater improvements. The difference in benefits between TF and WAAP-WLS vs 3PSM,  $p$ -curve,  $p$ -uniform, and PET-PEESE was particularly pronounced at higher  $k$ .

Given a true effect, RMSE of random effects meta-analysis was not terrible. Most methods again provided some benefit to RMSE, with the exception that, given small  $k$ ,  $p$ -uniform and  $p$ -curve did not improve RMSE, and PET-PEESE increased RMSE. 3PSM, WAAP-WLS, and trim-and-fill provided modest improvements in RMSE. Given larger  $k$ , all methods yielded improvements in RMSE, with trim-and-fill, WAAP-WLS, and 3PSM being slightly more efficient than  $p$ -uniform,  $p$ -curve, and PET-PEESE.

Adding heterogeneity to a true effect of zero caused modest increases in RMSE for all methods, and all methods provided some benefit relative to unadjusted random-effects. The greatest improvement was provided by 3PSM, followed by PET-PEESE,  $p$ -curve,  $p$ -uniform, and trim-and-fill. WAAP-WLS improved with  $k$ , performing worse than  $p$ -curve and  $p$ -uniform at  $k = 10$  but better than those methods at  $k = 60$ .

Adding heterogeneity to a true nonzero effect caused moderate increases in RMSE for all methods. Again, all methods provided some benefit relative to unadjusted random-effects, and these benefits were comparable across methods; the one exception was that PET-PEESE caused an increase in RMSE when  $k$  was small.

**95% CI coverage.** Given the null, because of the considerable publication bias, 95% CI coverage was very poor without adjustment (<3%). All methods lead to some improvement in coverage. Benefits of trim-and-fill were slight (28% at  $k = 10$ , 0% at  $k = 60$ ). Benefits of WAAP-WLS increased with  $k$  but did not improve beyond 55% coverage. Benefits of 3PSM, on the other hand, decreased with  $k$ , yielding 68% coverage at  $k = 10$  but 18% coverage at  $k = 60$ .  $p$ -uniform and PET-PEESE were the only methods to yield good coverage rates (90–94%).

Given a true effect, unadjusted random effects CI coverage was much better (86%). However, as  $k$  increased, coverage fell, presumably due to some combination of bias and insufficient CI width (e.g., at  $k = 10$ , coverage was 86%, but at  $k = 60$ , coverage was only 42%.) All adjustments achieved approximately correct coverage rates (90–97%) with the exception of PET-PEESE, which demonstrated undercoverage that grew worse with increasing  $k$ .

Given the null and heterogeneity, random effects CI coverage was slightly better but still very poor. All adjustments slightly improved CI coverage, with the greatest benefits observed from 3PSM (52% coverage at  $k = 10$ , 65% coverage at  $k = 60$ ) and PET-PEESE (60% coverage at  $k = 10$ , 36% coverage at  $k = 60$ ). These are still suboptimal coverage rates.

Given a true effect and heterogeneity, random effects CI

coverage substantially worsened relative to homogeneity. All adjustments somewhat improved CI coverage. At low  $k$ , 3PSM, WAAP-WLS, and  $p$ -uniform brought the greatest improvements, reaching coverage of 78%, 81%, and 87%, respectively. At high  $k$ , coverage of  $p$ -uniform and WAAP-WLS coverage fell noticeably ( $\sim 68\%$  coverage) while 3PSM improved to 93% coverage.

### The influence of QRPs

**Influence of QRPs on naive meta-analysis.** QRPs generally led to an increase in bias in random-effects meta-analysis. This increase was greatest when there was medium publication bias and the null was true; under these circumstances, high use of QRPs could double ME from 0.15 to 0.32. At small  $k$  and medium bias, QRPs could also approximately double the Type I error rate from 51% to 95%; at higher  $k$ , publication bias alone was enough to cause Type I error rates approaching 100%. The effect of QRPs was smaller given strong publication bias and a true null: an increase from about 0.32 to 0.44. Here the damage is presumably lesser because publication bias had already inflicted considerable mean error. Additionally, when there was a true effect of  $\delta = 0.5$ , or when there was no publication bias, the influence of QRPs on ME and RMSE was minimal.

However, when  $k$  was large, QRPs did lead to noticeable increases in Type I error even in the absence of publication bias. For example, high QRPs increased Type I error to 19%, even though the ME remained a mere 0.03. ( $k = 60$ ,  $\delta = 0$ ,  $\tau = 0$ , no publication bias, high QRPs). QRPs similarly tended to harm confidence interval coverage.

In summary, for our simulation design and for random-effects meta-analysis, QRPs can exacerbate the effects of publication bias when there is no effect; however, in our simulation, the effects of QRPs on ME are modest when (1) there is no publication bias, (2) publication bias is already quite strong, or (3) there is a true effect. Thus, QRPs, such as we have implemented them, seem to play a small role in meta-analytic bias on their own. In the company of moderate publication bias, however, QRPs can lead to a considerable amplification of problems.

**Influence of QRPs on bias-corrected meta-analysis.** The effect of QRPs in our simulation varied as a function of both method and performance metric. In this section, we focus chiefly on bias as the metric. We do this because the effects of QRPs on bias are the most straightforward and communicable. The effects of QRPs on RMSE, error rates, and coverage were generally a function of whether QRPs caused an increase or decrease in bias: Where QRPs reduced the absolute value of ME, RMSE and coverage probability generally improved; where QRPs increased the absolute value of ME, RMSE, coverage rates, and Type I and II error rates were accordingly poorer. In some cases, a curvilinear effect was observed; as QRPs increased, an initial positive bias

would be reduced, then become a negative bias, causing these metrics to first improve, then deteriorate. The influence of QRPs was generally strongest when there was a null or small effect, presumably because studies with medium or large true effects required less  $p$ -hacking to publish.

The effect of the QRPs we modeled on trim-and-fill were similar to its effects on random-effects analysis—bias increased when the null was true and there was medium or strong publication bias. This also led to elevated Type I error rates (except when there was heterogeneity and no publication bias, in which case Type I error decreased slightly). The effect of the QRPs we modeled on WAAP-WLS were similar, but increases in bias were smaller compared to trim-and-fill. A curious exception is that QRPs in our simulation reduce Type I error in WAAP-WLS when there is medium bias and large  $k$ , although these reductions still lead to unacceptable ( $\geq 40\%$ ) Type I error rates. Perhaps in these cases, WAAP-WLS switches from the better-powered WLS test to the poorer-powered WAAP test.

In contrast, QRPs—as we implemented them—*reduced* ME in  $p$ -curve and  $p$ -uniform. In the context of homogeneity, this led to underestimation of the effect size and an increase in Type II error. In the context of heterogeneity, this reduced overestimation due to heterogeneity and reduced the absolute value of ME, reducing Type I errors and increasing Type II errors. We consider this as two simple effects of opposite sign: Heterogeneity causes upward ME, and QRPs cause downward ME, so they may yield a smaller absolute ME when both present than when only one or the other is present. QRPs also helped to reduce the upward bias in the average  $p$ -curve/ $p$ -uniform estimates when the null was true, perhaps by increasing the number of significant studies available.

The QRPs we modeled nudged PET-PEESE estimates downwards. When PET-PEESE was biased upwards in our simulation, as in the case of small or null effects with publication bias and heterogeneity, this yielded a less biased estimate and slightly better Type I error rates. When PET-PEESE was unbiased or downward biased, as in the case of nonzero true effects, this led to greater downward bias. This downward bias could be quite strong when the null was true. PET-PEESE reported statistically significant effects of opposite sign in many analyses; these Type I error rates tended to grow with increasing QRPs, publication bias, and  $k$ , with rates ranging from 9% (medium QRPs, high censoring,  $k = 10$ ) to 62% (high QRPs, medium censoring,  $k = 100$ ). Researchers have, at times, considered a significant and negative PET-PEESE estimate as evidence that the PET-PEESE estimate is incorrect, choosing to interpret instead less extreme adjustments such as trim-and-fill (see, e.g., Bediou et al., 2018). In these cases, such an extreme estimate from PET-PEESE probably is incorrect, but researchers should be aware that there is likely to be some combination of QRPs

and publication bias in this circumstance and, perhaps, a null effect.

QRPs, as we have implemented them, generally led to a slight downward bias in 3PSM. This bias was stronger when heterogeneity was present. At worst, given high heterogeneity ( $\tau = 0.4$ ), ME caused by QRPs could be as large as  $-0.32$ . QRPs therefore tended to reduce Type I error and increase Type II error in 3PSM.

What's the worst that can happen to each estimator as a consequence of the QRPs we implemented? In random-effects meta-analysis, trim-and-fill, and WAAP-WLS, QRPs can exacerbate the effects of publication bias. Under moderate publication bias, QRPs increased the bias in these estimators substantially (for random effects, from  $ME = 0.14$  to  $ME = 0.31$ ; for trim-and-fill, from  $ME = 0.08$  to  $0.22$ ; for WAAP-WLS, from  $0.11$  to  $0.15$ ). These changes in bias caused corresponding increases in Type I error rates and could increase error rates by as much as 40 percentage points. Under strong publication bias, QRPs could cause an additional  $0.12$  of bias, but this was less dramatic relative to the bias already inflicted by publication bias. PET-PEESE,  $p$ -curve, and  $p$ -uniform each could demonstrate downward bias up to  $-0.14$ . This could cause a loss of power of up to 17 percentage points in PET-PEESE, 17 percentage points in  $p$ -curve, and 37 percentage points in  $p$ -uniform. High application of QRPs also caused PET-PEESE to frequently mistake null effects for significant negative effects (up to 62% at  $\tau = 0$ , high QRPs, medium publication bias). 3PSM could underestimate  $\delta = 0.5$  as being as little as  $d = 0.31$ , and loss of power could reach 32 percentage points.

As a very broad summary, the reactions of estimators to QRPs may be considered in two clusters. For RE, trim-and-fill, and WAAP-WLS, QRPs cause overestimation, particularly of null effects. For PET-PEESE,  $p$ -curve,  $p$ -uniform, and 3PSM, QRPs cause underestimation of true effects and noticeable loss of power. Critically, the results described above are specific to the way we simulated QRPs. There are likely a variety of changes that could be made to our approach that would provide very different results. We see this topic as an important area for future research.

## Discussion

We inspected and compared the efficacy of meta-analytic adjustments for bias across hundreds of thousands of simulated literatures representing a range of true effect sizes, degrees of heterogeneity, degrees of publication bias, and degrees of questionable research practices. We assessed the results according to both the ability to reject a null effect/detect a true effect and the ability to estimate the mean of the distribution of true underlying effects. The following provides a coarse summary of the three overall patterns we observed, as well as some general recommendations.

First, random-effects meta-analysis, trim-and-fill, and

WAAP-WLS showed alarmingly high false-positive rates (Figure 3) and overestimation (Figure 4) in the face of publication bias in combination with a zero or small true effect size. Generally, WAAP-WLS outperformed both random-effects meta-analysis and trim-and-fill. Second, both  $p$ -curve and  $p$ -uniform had reasonable false-positive rates and little bias under homogeneity. With increasing heterogeneity, however, both showed increasing false-positive rates (Figure 3) and overestimation (Figure 4), particularly for a zero or small true effect size. This poor performance was actually mitigated by increasing QRPs, and was primarily independent of changes in publication bias and sample size. Again, we note that the original authors of  $p$ -curve argue that its performance should not be evaluated using the average true underlying effect, as is the usual approach in the meta-analytic literature, but rather, using the average of the effects submitted to it (Simonsohn et al., 2014). Those results are available in the supplemental material, but in general, it should be noted that  $p$ -curve showed good performance when estimating this quantity. If no QRPs are present,  $p$ -curve recovers that quantity with very low ME, regardless of the level of publication bias and  $\delta$ . QRPs as we have modeled them induced a downward bias, which was stronger for smaller  $\delta$ s and for less heterogeneity. These results are consistent with previous simulation results (Simonsohn et al., 2014).

Third, PET-PEESE and 3PSM both showed mostly reasonable false-positive rates but suffered from low power at smaller sample sizes, higher heterogeneity, and less publication bias and QRPs (Figure 3). These two methods also showed similar patterns in terms of estimation error (Figure 4): Both methods tended to show underestimation with decreasing sample sizes and increasing QRPs, publication bias, and heterogeneity. Although the two methods produced similar overall patterns of results, 3PSM almost always outperformed PET-PEESE.

Furthermore, it is worth noting that random-effects meta-analysis, WAAP-WLS, and PET-PEESE always returned at least some estimate, whereas the other methods sometimes failed to converge or could not be applied because of lack of significant studies. Information on convergence rates is available in the supplemental material and should be considered alongside the performance we report above—for example, it may be that consistent failures to converge in certain simulated conditions indicate that the method will be less applicable in real world usage.

Based on our results, we believe we can confidently make five general recommendations.

1. If publication bias is highly unlikely (e.g., data are from a multi-laboratory pre-registered replication), one should prefer random-effects meta-analysis to any of the other methods we examined.

2. When there may be publication bias, do not rely on random-effects meta-analysis alone. Publication bias can

quickly accumulate in even small sets of published studies, leading to overestimated effects and high Type I error rates.

3. Recognize that the popular trim-and-fill adjustment, while efficient, reduces bias and Type I error rates only slightly. To achieve stronger reductions in bias, adjustments such as PET-PEESE,  $p$ -curve/ $p$ -uniform, and 3PSM should be considered. However, one should also recognize that these adjustments are often inefficient and individual estimates may be poor, even if they are unbiased in the long run.

4. Take steps to ensure the completeness and transparency of the original literature. An ounce of registered report is worth a pound of  $p$ -curve.

5. Do not use  $p$ -curve or  $p$ -uniform if heterogeneity is expected or if many studies yield non-significant results.

### Limits on generalizability

Simulation studies necessarily require making assumptions that might limit the generalizability of their results to real data. Although we simulated a data generation process that might plausibly underlie how real world research in psychology is done, there are several limits that should be kept in mind when considering our findings. First, in the absence of QRPs, we simulated the data generation process as a two group design, despite the fact that real research designs are rarely this simple. However, the vast majority of meta-analyses use effect size measures, such as correlations and standardized mean differences, that ignore such design complexities (see, for example, table S1 in, Fanelli et al., 2017). For example, to meta-analyze data from an experiment with a two-by-two (or more) factorial design, one would typically calculate a standardized mean difference ( $d$ ) by either discarding the factors that are not of primary interest or collapsing across them. So for a two-by-two with a per-group sample size of 20, the comparison entered into the meta-analysis would either have a total  $N$  of  $2*20 = 40$  (when the second factor is discarded) or a total  $N$  of  $2*2*20 = 80$  (when the second factor is collapsed). Therefore, although most designs are more complex than the two-group case we simulate, data in meta-analyses are often reduced to this simple form. As a result, our findings generalize to meta-analyses where the meta-analyst has chosen to handle more complex designs by discarding non-focal factors. In cases where the meta-analyst collapses across the other factors, our simulation likely underestimates sample size on average. However, it should be noted that the choice to collapse is problematic given the required assumption that the other factors do not interact with the comparison of interest (i.e., there is a true interaction effect of 0). Thus, our findings generalize to the least problematic case.

A second, related issue exists if real world data are generated by single-sample designs (e.g., correlational studies). If such studies tend to have larger or smaller sample sizes than

those with factorial designs, it's possible our simulation either under- or over-estimates sample size, respectively. Critically, the generalizability issue here is *only* related to sample sizes, not to the fact that different study designs tend to be summarized with different effect size measures. At the level of the study, one can translate between most effect size measures without changing statistical significance or effect direction. Given that bias acts at the study level through these two features, *the generalizability of our results holds regardless of whether data originally take the form of standardized mean differences or correlations.*

A third point to consider is whether our implementation of bias mirrors bias in real-world data. We implemented publication bias using specific functions with specific parameter values. Of course, it would be entirely possible to use different functions or different parameters. What is unclear, however, is the degree to which different choices at this level would produce different results—for example, we intentionally modeled publication bias in a way that differed from what 3PSM,  $p$ -curve, and  $p$ -uniform are designed for, so it may be that these methods would show improved performance under different specifications of publication bias. Ultimately this is an empirical question and should be the focus of future research. Additionally, our implementation of QRPs is extremely specific and might limit the generalizability of our results. Because the kinds of QRPs that can be applied depend entirely on the design of the specific study, there is an infinite number of possible ways to simulate QRPs (Hartgerink, van Aert, Nuijten, Wicherts, & Van Assen, 2016). Thus, our results based on the QRPs we examine likely won't generalize to designs that are dramatically different than those we simulated.

Fourth, it is impossible to perfectly mirror how real data are generated. However, it is our hope that researchers can use the framework we have described here to close this gap and tackle some of the issues we mention above. It would be relatively easy, for example, to modify our code to use larger or smaller sample sizes—thereby addressing the first two issues described above—and then assess whether this substantially changes how the methods perform.

Finally, bias-correction in meta-analyses is an active field of research and multiple new methods were published after our simulations were done, for example an extension of  $p$ -uniform, called  $p$ -uniform\*, which now estimates heterogeneity and includes non-significant results (R. C. M. van Aert & van Assen, 2018), or a Bayesian fill-in method called “BALM” (Du, Liu, & Wang, 2017).

### Method Performance Checks and Sensitivity Analysis for Meta-Analysis in Psychology

Several authors have suggested the value of a sensitivity analysis in order to evaluate the robustness of conclusions from a meta-analysis (e.g., APA Publications and Commu-



nications Board Working Group on Journal Article Reporting Standards, 2008; McShane et al., 2016; van Aert, 2017; R. C. van Aert et al., 2016). If results do not substantially change across a range of different methods and assumptions, the conclusions can be considered to be robust. However, the set of employed methods in a sensitivity analysis should only include methods that can be expected to perform reasonably well. Put differently, if a method is known to perform poorly under the conditions that apply to a meta-analysis at hand, it should not be included in a sensitivity analysis, or it should at least be treated with skepticism and given less weight when evaluating the results.

One of the goals of this paper is to encourage and facilitate sensitivity analysis in meta-analysis by suggesting a two-step procedure: First, evaluate which bias-correcting methods can be a priori expected to perform reasonably well in research conditions that are plausible for the meta-analysis at hand (*method performance check*). Second, compute and present meta-analytic estimates from all included methods side by side in order to evaluate the variability (or robustness) of conclusions (*sensitivity analysis*). This evaluation should respect the results from the method performance check and weigh methods accordingly.

For doing a sensible sensitivity analysis, we recommend that meta-analysts and consumers of meta-analyses focus on the question, “Do my conclusions depend on a meta-analytic method that performs poorly in plausible conditions?” If the answer is “yes,” then those conclusions should clearly be revisited. To help ask this question, we provide an interactive app (<http://www.shinyapps.org/apps/metaExplorer/>) that, for a given method and a given definition of “performs poorly,” identifies all of the conditions in our simulation for which the answer to the question is “yes.” In the following, we provide an illustration of how one might perform this form of method performance check, and how this guides the subsequent sensitivity analysis.

#### **A real world example: Two data sets on ego depletion.**

We use data from studies on the topic of ego depletion for this example because it is relatively easy to understand and there are data from both meta-analyses of the literature and a large-scale pre-registered replication.

Briefly, the limited strength model of self-control holds that any act of self-control will result in subsequent acts of self-control being less likely to succeed—a state referred to as ego depletion (Muraven, Tice, & Baumeister, 1998). Typically experiments aimed at examining this hypothesis involve participants completing a sequence of two tasks—a manipulation task and an outcome task. Prior to the outcome task, participants in the “depletion” condition are given a version of the manipulation task that is designed to require more self-control than the version given to participants in the control condition. Support for ego depletion is claimed when participants in the depletion condition perform worse on the sub-

sequent outcome task than participants in the control condition. Following convention, we represent this effect as a standardized mean difference ( $d$ ) where the depletion effect is represented by *higher* values. In the following, we analyze both a pre-existing meta-analytic sample of  $k = 116$  studies (Carter, Kofler, Forster, & McCullough, 2015) and a data set of  $k = 23$  studies from a pre-registered replication report (Hagger et al., 2016). We apply each of our methods to these two datasets—see Table 2 and Table 3.

Imagine that researchers agree to the logic detailed by Carter and McCullough (2018) that a depletion effect of  $\delta \leq 0.15$  should be considered practically equivalent to zero. Unfortunately, in this case different meta-analytic methods would come to different conclusions (see Table 2). On the basis of the results from the random-effects model, WAAP-WLS, the trim-and-fill method,  $p$ -curve,  $p$ -uniform, and the selection model, a group of researchers could conclude that the depletion effect is practically significant with  $\delta > 0.15$ . In contrast, on the basis of results from PET, PEESE, or PET-PEESE, a separate group of researchers could conclude that the depletion effect is practically non-significant (i.e.,  $\delta \leq 0.15$ ). Hence, a “naive” sensitivity analysis would be inconclusive, as the variability in results is so large that either conclusion can be drawn. To help overcome this inconclusiveness, we want researchers to ask “Do my conclusions depend on a meta-analytic method that performs poorly in plausible conditions?”. Such a method performance check gives guidance which results should be given more weight and credibility.

**What are “plausible conditions”?** For a method performance check, we need “plausible conditions”. Comparable to an a-priori power analysis, these considerations should relate to the specific research environment of the meta-analysis at hand. Only if no specific prior knowledge is available, general knowledge about the field can be used as an approximation. For example, some degree of bias seems possible: In the field of psychology/psychiatry, more than 90% of all published hypothesis tests are significant (Fanelli, 2011)—despite the average power being estimated as around 35% (Bakker, van Dijk, & Wicherts, 2012)—and, whereas reported effects tend to be statistically significant, unreported effects tend not to be (Franco, Malhotra, & Simonovits, 2016). Moreover, there is both direct (Franco et al., 2016; LeBel, McCarthy, Earp, Elson, et al., 2017) and self-report (John et al., 2012) evidence of the use of QRPs, and several studies have found evidence of small-study effects consistent with the presence of publication bias (Bakker et al., 2012; Kühberger, Fritz, & Scherndl, 2014; Fanelli et al., 2017).<sup>8</sup> In addition to bias, a degree of heterogeneity

<sup>8</sup>It should be noted that, in contrast to the above cited work, two “meta-meta-analyses” suggest that the influence of bias in psychology is relatively small (T. Stanley, Carter, & Doucouliagos, *in press*; R. C. M. van Aert, 2018). Because of the specifics of each of these

seems very likely when diverse experimental paradigms are summarized in a meta-analysis (e.g., not in multilab registered reports, van Erp et al., 2017; Tackett, McShane, Bockenholt, & Gelman, 2017; T. Stanley et al., *in press*). Finally, it seems that the typical true effect in psychology research is rather small: The median published effect size is around  $d = 0.3$  to  $0.4$  (Richard, Bond Jr, & Stokes-Zoota, 2003; Bosco, Aguinis, Singh, Field, & Pierce, 2015), and as this estimate is not corrected for publication bias, the typical true effect is likely smaller. This general observation, of course, does not preclude that some effects in psychology indeed are large. However, as the (uncorrected) meta-analysis from the (Hagger, Wood, Stiff, & Chatzisarantis, 2010) meta-analysis shows an overall effect of  $d = 0.62$ , which most likely is inflated, we define for our specific example that plausible conditions for ego depletion research have a true effect size of  $\delta \leq 0.5$ . Therefore, of the conditions we simulated, we defined that the most plausible conditions for a meta-analysis of the set of ego-depletion studies published in (Carter et al., 2015) will have medium to high publication bias and QRPs,  $\tau \geq 0.2$ , and effect sizes under  $H_1$  of  $\delta \leq 0.5$ . In our interactive app, we evaluate the performance of all included estimators under these plausible conditions at  $k = 100$ , which is the closest simulated value to the observed  $k = 116$ .

**What is “poor performance” of a meta-analytic method?** Given the above definition of plausible conditions, the next step for a method performance check is to identify defensible choices for the definition of “poor performance.” For simplicity, in this example we consider only one metric, mean error (ME)<sup>9</sup>, and ask whether the method is likely to be biased enough that a true null effect is mistaken for a practically significant effect, or conversely, a true effect is mistaken for a practically null effect. For each possible true effect, mean error large enough to make this mistake is considered “poor” performance. Where  $\delta = 0$ , a maximum of  $ME \geq 0.15$  is poor, since this would mistake a true null effect for a practically significant effect, on average. Where  $\delta = 0.2$ , a maximum downward bias of  $ME \leq 0.05$  is poor, since a practically significant true effect would be underestimated as practically non-significant. Where  $\delta = 0.5$ , a maximum downward bias of  $ME \leq 0.35$  would cause the same mistake.

**Considering method performance in a sensitivity analysis.** With these definitions for poor performance and plausible conditions in hand, we can use the interactive app (<http://www.shinyapps.org/apps/metaExplorer/>) to judge whether meta-analytic conclusions rely on methods that perform poorly in plausible conditions. As a result, we see that the conclusion that the depletion effect is practically significant based on the results from the random-effects model, the trim-and-fill, WAAP-WLS, 3PSM,  $p$ -curve, and  $p$ -uniform is indefensible, since each of these methods perform poorly in at least one of the plausible conditions we defined (primarily

when  $\delta = 0$ ; Table 2). In contrast, the conclusion that the depletion effect is practically non-significant—which is based on results from PET, PEESE, and PET-PEESE—appears to be reasonable on the basis that PET does not perform poorly in any of the plausible conditions we examined. To summarize, if one gives more weight to methods that a priori perform well under the hypothesized conditions, one would lean towards the conclusion that the depletion effect is practically non-significant.

Using the data from Hagger et al. (2016) and the same definitions of poor performance as above, we apply our analysis using a very different set of plausible conditions. Because the Hagger et al. (2016) data come from a pre-registered replication report, there is no reason to think that the results are influenced by publication bias or QRPs to any substantial degree. Furthermore, one can expect significantly less heterogeneity in these data as compared to those from Carter et al. (2015) because data collection at each location was conducted using identical study designs executed with a pre-registered, automated script. Furthermore, there is evidence that heterogeneity in these kinds of multilab registered reports is generally low (Klein et al., 2018). Thus, of the conditions we simulated, we can imagine that researchers would view the most plausible conditions as having no publication bias, no QRPs,  $\tau = 0$ , and effect sizes under  $H_1$  of  $\delta \leq 0.5$ . In our interactive table, we evaluate each method’s performance using the same criteria of “poor performance” as above and at  $k = 30$  (which is the closest simulated value to the observed  $k$  of 23 studies).

The results from applying the various meta-analytic methods to the Hagger et al. (2016) data set are shown in Table 3. Unlike the results from the Carter et al. (2015) data (Table 2), we suspect that the vast majority of researchers would conclude that the results from the Hagger et al. (2016) data uniformly suggest that the depletion effect is practically non-significant. Again, we suggest to use a method’s performance to evaluate whether such a conclusion is based on results from methods that perform poorly in plausible con-

studies, it is difficult to reconcile their general conclusions. For our analysis here, we decided to take the conservative route and err on the side of assuming the existence of bias, but we recommend that, when applying our approach, meta-analysts consider these issues and explicitly describe their reasoning.

<sup>9</sup>Note that by using this metric, we are focused on the point estimate provided by each method, not the upper or lower bounds of the confidence interval. Furthermore, such an evaluation of methods should also consider RMSE, error rates, and other performance metrics, which all are provided in the online app. Furthermore, we want to note that these considerations only take *statistical* properties of the estimators into account. It has been argued that other dimensions of quality, such as the presence or absence of proper randomization, should also go into the weighting of primary studies in a meta-analysis (Detsky, Naylor, O’Rourke, McGeer, & L’Abbé, 1992). But this is a topic for another article.

Table 2  
 Method performance checks for plausible conditions of the Carter, Kofler, Forster, and McCullough (2015) data

Method	Estimate	Poor ( $\delta = 0$ )	Poor ( $\delta = 0.20$ )	Poor ( $\delta = 0.50$ )
Random-effects	0.43 [0.34, 0.52]	Yes	No	No
Trim-and-fill	0.24 [0.13, 0.34]	Yes	No	No
WAAP-WLS	0.35 [0.26, 0.43]	Yes	No	No
<i>p</i> -curve	0.55 [NA, NA]	Yes	No	No
<i>p</i> -uniform	0.55 [0.33, 0.71]	Yes	No	No
PET	-0.27 [-0.52, 0.00]	No	No	No
PEESE	0.00 [-0.14, 0.15]	Yes	No	No
PET-PEESE	-0.27 [-0.52, 0.00]	Yes	No	No
3PSM	0.33 [0.19, 0.47]	No	Yes	No

ditions. By using our interactive table, it appears that, on the basis of the definitions above, this conclusion is perfectly defensible: None of the methods—except *p*-curve and *p*-uniform—perform poorly in any of the plausible conditions we defined (see Table 3).

To summarize, we performed two method performance checks in two different plausible research environments. In both meta-analyses, those methods which a priori could be expected to show a reasonably good performance suggest that the true effect is not practically and significantly different from zero.

**Further considerations for method performance checks and sensitivity analyses.** There are several important points that should be considered for our proposed approach. First, this form of method performance check depends on the specifics of our simulation design. Because our simulation covers only a limited set of possible data generation processes, it is possible that our approach does not generalize to real world situations that meta-analysts and consumers of meta-analysis will encounter. Indeed, because of the infinite possible processes that might generate real world data, generalizability will always be a concern.

Second, we suspect that some readers will worry that the method performance checks we describe here are subject to a kind of “assumption hacking” whereby a researcher who is partial to a certain view can pick and choose the definitions of plausible conditions and poor performance that provide the result they want. This is technically correct, but the key strength of our approach is that it is explicit and transparent.

Researchers will need to clearly state their assumptions to run this method performance check, and consumers of the results can assess whether such assumptions appear reasonable to them. If not, our interactive table can easily be used to run an alternative method performance check, thereby preventing the possibility of effective assumption hacking. Furthermore, we suggest that analysts preregister a method performance check prior to data collection and define which methods will be given the biggest weight in case they disagree. Finally, we encourage researchers to report results from all meta-analytic methods that reasonably can be considered, even if they did not pass some of the method performance checks, because other researchers might want to apply a different emphasis in their subjective evaluation.

### Ways forward

Based on our results, we stress that meta-analysis in psychology is *difficult*. Observable issues such as small samples—in both the primary literature and at the level of the meta-analysis—interact with heterogeneity and bias, both of which are unknowable in terms of severity and functional form (e.g., do the true effects follow a normal distribution? Is publication bias applied as we have modeled it?). Thus, it is hard to interpret the results of a meta-analysis in psychology, just as it is difficult to interpret the results of any single replication study (Braver, Thoemmes, & Rosenthal, 2014; D. J. Stanley & Spence, 2014; Fabrigar & Wegener, 2016).

Meta-analysts might have the hope that different bias-correcting methods all converge on a true value. However,

Table 3  
*Meta-analytic estimates of the Hagger et al., 2016 data*

Method	Estimate	Poor ( $\delta = 0$ )	Poor ( $\delta = 0.20$ )	Poor ( $\delta = 0.50$ )
Random-effects	0.04 [-0.06, 0.14]	No	No	No
Trim-and-fill	0.04 [-0.06, 0.14]	No	No	No
WAAP-WLS	0.04 [-0.08, 0.15]	No	No	No
<i>p</i> -curve	0.00 [NA, NA]	Yes	No	No
<i>p</i> -uniform	0.00 [NA, NA]	Yes	No	No
PET	-0.03 [-0.80, 0.74]	No	No	No
PEESE	0.00 [-0.41, 0.40]	No	No	No
PET-PEESE	-0.03 [-0.80, 0.74]	No	No	No
3PSM	0.01 [-0.10, 0.12]	No	No	No

our simulations show that different methods often do *not* converge. For example, in the case of a null effect and strong publication bias, PET and trim-and-fill will virtually never give the same answer because trim-and-fill performs so poorly. For this reason we caution against ideas of “triangulation” or majority votes of multiple methods. One should instead think carefully about which method(s) can be expected to perform well. We think a good approach is the combination of a method performance check with a subsequent sensitivity analysis, either as we have defined it or as put forward by others (e.g., J. Copas & Shi, 2000; Vevea & Woods, 2005; J. B. Copas, 2013; Kim, Bangdiwala, Thaler, & Gartlehner, 2014).

Furthermore, and along the same lines of what has been argued for replication results (e.g., D. J. Stanley & Spence, 2014), we conclude that the field should modify its expectations about meta-analysis. Researchers in psychology should not expect to produce a conclusive, debate-ending result by conducting a meta-analysis on an existing literature. Instead, we imagine meta-analyses may serve best to draw attention to the existing strengths and/or weaknesses in a literature (e.g., Carter et al., 2015; van Elk et al., 2015; Hilgard, Engelhardt, & Rouder, 2017). Those results can then inspire a careful re-examination of methodology and theory, perhaps followed by large-scale, preregistered replication efforts (e.g., Hagger et al., 2016). Such registered replications reports can then be summarized with random-effects meta-analytic methods, which show the best performance in the absence of bias (Figures 3 and 4).

## Conclusion

In simulations we compared bias-correcting meta-analysis methods using effect sizes, sample sizes, questionable research practices, and degrees of publication bias that plausibly represent real data in psychology. We found that each of the seven methods we implemented show unacceptable performance in at least some conditions. This is not an entirely surprising result given previous simulation studies (e.g., Hedges & Vevea, 1996; Moreno et al., 2009; Rucker et al., 2011; T. Stanley & Doucouliagos, 2014; Simonsohn et al., 2014; R. C. van Aert et al., 2016; McShane et al., 2016; T. Stanley, 2017). However, it highlights an important conclusion that we believe needs to be more widely acknowledged: Meta-analysts and consumers of meta-analyses in psychology should not expect to come to definitive conclusions. Instead, we believe that the most productive outcomes will be generated by method performance checks, sensitivity analyses, and a willingness to carefully design and conduct pre-registered replications.



### Glossary

**$d_i$  and  $v_i$ .** The observed effect size ( $d$ ) and its associated variance ( $v$ ) for the  $i^{th}$  study. We calculate  $d$  as  $\frac{M_1 - M_2}{S}$ , where  $M_1$  and  $M_2$  are the means of the two groups and  $S$  is the pooled standard error of the two groups. The variance of  $d$  can be calculated as  $\frac{n_1 + n_2}{n_1 n_2} + \frac{d^2}{2(n_1 + n_2 - 2)} \cdot \frac{n_1 + n_2}{n_1 n_2 - 2}$ , where  $n_1$  and  $n_2$  are the sample sizes of the two groups.

**$\delta$ .** Under the fixed-effect model,  $\delta$  is the hypothetical true underlying effect estimated by each study. Under the random-effects model,  $\delta$  is the mean of the distributions of hypothetical true underlying effects.

**$\tau$ .** The standard deviation of the distribution of hypothetical true underlying effects assumed by the random-effects model. It is also referred to as measuring between-study heterogeneity.

**$k$ .** The number of studies in a meta-analytic sample.

**Mean error (ME).** The average of the deviations of each estimate from the true effect (often called bias). Nonzero ME indicates that the expected value of the estimate does not converge on the true value in the long run, being instead too high or too low. ME is not sensitive to variance in estimates, so it is possible that a method produces low ME by large *but equal* over- and underestimation. Such a case would yield estimates that are accurate on average, but any individual estimate could be quite far from the truth.

**Root Mean Square Error (RMSE).** RMSE incorporates information about average error as well as the variance (i.e., the efficiency) in the estimates. It is possible to observe low RMSE even when a method produces estimates that are consistently biased in one direction. Imagine a very narrow distribution of estimates that is centered a bit above the true value. On average, the estimates are too high, but the variability of these estimates will be low. Thus, a method's estimation performance must be considered in terms of both ME and RMSE. For both ME and RMSE, values as close to zero as possible are desirable.

**95% Coverage Probability.** The proportion of each method's 95% confidence intervals that included the true value of  $\delta$ . Optimally, the coverage probability is at the nominal level of 95%. Low coverage is problematic as most intervals do not contain the true value. Coverage rates higher than 95% may indicate exceedingly wide intervals.

### References

- APA Publications and Communications Board Working Group on Journal Article Reporting Standards. (2008). Reporting standards for research in psychology: Why do we need them? what might they be? *The American Psychologist*, 63(9), 839–851. doi:10.1037/0003-066X.63.9.839
- Augusteijn, H. E. et al. (2018). The effect of publication bias on the q test and assessment of heterogeneity. *Psychological methods*.
- Baker, R. D. & Jackson, D. (2013). Meta-analysis inside and outside particle physics: Two traditions that should converge? *Research Synthesis Methods*, 4(2), 109–124.
- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, 7(6), 543–554.
- Bayarri, M. J. & DeGroot, M. H. (1991). *The analysis of published significant results*. Citation Key: Bayarri:1991tj. Retrieved from <http://www.stat.purdue.edu/docs/research/tech-reports/1991/tr91-21.pdf>
- Bediou, B., Adams, D. M., Mayer, R. E., Tipton, E., Green, C. S., & Bavelier, D. (2018). Meta-analysis of action video game impact on perceptual, attentional, and cognitive skills. *Psychological bulletin*, 144(1), 77.
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2011). *Introduction to meta-analysis*. John Wiley & Sons.
- Bosco, F. A., Aguinis, H., Singh, K., Field, J. G., & Pierce, C. A. (2015). Correlational effect size benchmarks. *Journal of Applied Psychology*, 100(2), 431.
- Boulesteix, A.-L., Wilson, R., & Hapfelmeier, A. (2017). Towards evidence-based computational statistics: Lessons from clinical research on the role and design of real-data benchmark studies. *BMC Medical Research Methodology*, 17, 138. doi:10.1186/s12874-017-0417-2
- Braver, S. L., Thoenmes, F. J., & Rosenthal, R. (2014). Continuously cumulating meta-analysis and replicability. *Perspectives on Psychological Science*, 9(3), 333–342.
- Burton, A., Altman, D. G., Royston, P., & Holder, R. L. (2006). The design of simulation studies in medical statistics. *Statistics in medicine*, 25(24), 4279–4292.
- Carter, E. C., Kofler, L. M., Forster, D. E., & McCullough, M. E. (2015). A series of meta-analytic tests of the depletion effect: Self-control does not seem to rely on a limited resource. American Psychological Association.
- Carter, E. C. & McCullough, M. E. (2018). A simple, principled approach to combining evidence from meta-analysis and high-quality replications. *Advances in Methods and Practices in Psychological Science*, 1(2), 174–185. doi:10.1177/2515245918756858. eprint: <https://doi.org/10.1177/2515245918756858>
- Coburn, K. M. & Vevea, J. L. (2017). *Weightr: Estimating weight-function models for publication bias*. R package version 1.1.2. Retrieved from <https://CRAN.R-project.org/package=weightr>

- Cooper, H., Hedges, L. V., & Valentine, J. C. (2009). *The handbook of research synthesis and meta-analysis*. Russell Sage Foundation.
- Copas, J. B. (2013). A likelihood-based sensitivity analysis for publication bias in meta-analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(1), 47–66.
- Copas, J. & Shi, J. Q. (2000). Meta-analysis, funnel plots and sensitivity analysis. *Biostatistics*, 1(3), 247–262.
- Detsky, A. S., Naylor, C. D., O'Rourke, K., McGeer, A. J., & L'Abbé, K. A. (1992). Incorporating variations in the quality of individual randomized trials into meta-analysis. *Journal of Clinical Epidemiology*, 45(3), 255–265.
- Du, H., Liu, F., & Wang, L. (2017, December). A Bayesian "fill-in" method for correcting for publication bias in meta-analysis. *Psychological Methods*, 22, 799–817. doi:10.1037/met0000164. pmid: 29265851
- 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.
- Egger, M., Smith, G. D., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *BMJ*, 315(7109), 629–634.
- Fabrigar, L. R. & Wegener, D. T. (2016). Conceptualizing and evaluating the replication of research results. *Journal of Experimental Social Psychology*, 66, 68–80.
- Fanelli, D. (2011). Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3), 891–904.
- Fanelli, D., Costas, R., & Ioannidis, J. P. (2017). Meta-assessment of bias in science. *Proceedings of the National Academy of Sciences*, 201618569.
- Ferguson, C. J. & Heene, M. (2012). A vast graveyard of undead theories: Publication bias and psychological science's aversion to the null. *Perspectives on Psychological Science*, 7(6), 555–561.
- Franco, A., Malhotra, N., & Simonovits, G. (2016). Underreporting in psychology experiments: Evidence from a study registry. *Social Psychological and Personality Science*, 7(1), 8–12.
- Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, 82(1), 1.
- Guan, M. & Vandekerckhove, J. (2016). A bayesian approach to mitigation of publication bias. *Psychonomic bulletin & review*, 23(1), 74–86.
- Hagger, M. S., Wood, C., Stiff, C., & Chatzisarantis, N. L. D. (2010). Ego depletion and the strength model of self-control: A meta-analysis. *Psychological Bulletin*, 136(4), 495–525. doi:10.1037/a0019486
- Hagger, M. S., Chatzisarantis, N. L., Alberts, H., Anggono, C. O., Batailler, C., Birt, A., ... Bruyneel, S., et al. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, 11(4), 546–573.
- Hartgerink, C. H., van Aert, R. C., Nuijten, M. B., Wicherts, J. M., & Van Assen, M. A. (2016). Distributions of p-values smaller than .05 in psychology: What is going on? *PeerJ*, 4, e1935.
- Hedges, L. V. (1984). Estimation of effect size under nonrandom sampling: The effects of censoring studies yielding statistically insignificant mean differences. *Journal of Educational Statistics*, 9(1), 61–85.
- 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.
- Hilgard, J., Engelhardt, C. R., & Rouder, J. N. (2017). Overstated evidence for short-term effects of violent games on affect and behavior: A reanalysis of anderson et al. (2010). *Psychological Bulletin*, 143(7), 757–774. doi:10.1037/bul0000074
- Idris, N. & Ruzni, N. (2012). Performance of the trim and fill method in adjusting for the publication bias in meta-analysis of continuous data. *American Journal of Applied Sciences*, 9(9), 1512–1517.
- Ioannidis, J., Stanley, T., & Doucouliagos, H. (2017). The power of bias in economics research. *The Economic Journal*, 127(605).
- Iyengar, S. & Greenhouse, J. B. (1988). Selection models and the file drawer problem. *Statistical Science*, 109–117.
- 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.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 0956797611430953.
- Kim, N. Y., Bangdiwala, S. I., Thaler, K., & Gartlehner, G. (2014). Samurai: Sensitivity analysis of a meta-analysis with unpublished but registered analytical investigations (software). *Systematic reviews*, 3(1), 27.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Reginald B. Adams, J., Alper, S., ... Nosek, B. A. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443–490. doi:10.1177/2515245918810225
- Koehler, E., Brown, E., & Haneuse, S. J.-P. A. (2009). On the assessment of monte carlo error in simulation-based statistical analyses. *The American Statistician*, 63(2), 155–162. doi:10.1198/tast.2009.0030
- Kühberger, A., Fritz, A., & Scherndl, T. (2014). Publication bias in psychology: A diagnosis based on the corre-

- lation between effect size and sample size. *PLoS one*, 9(9), e105825.
- LeBel, E. P., McCarthy, R. J., Earp, B. D., Elson, M., et al. (2017). A unified framework to quantify the trustworthiness of empirical research.
- Marszalek, J. M. (2011). Sample size in psychological research over the past 30 years [data file]. Retrieved from <https://mospace.umsystem.edu/xmlui/handle/10355/62220>
- 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.
- 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.
- Moreno, S. G., Sutton, A. J., Ades, A., Stanley, T., Abrams, K. R., Peters, J. L., & Cooper, N. J. (2009). Assessment of regression-based methods to adjust for publication bias through a comprehensive simulation study. *BMC Medical Research Methodology*, 9(1), 2.
- Muraven, M., Tice, D. M., & Baumeister, R. F. (1998). Self-control as a limited resource: Regulatory depletion patterns. *Journal of personality and social psychology*, 74(3), 774.
- Peters, J. L., Sutton, A. J., Jones, D. R., Abrams, K. R., & Rushton, L. (2007). Performance of the trim and fill method in the presence of publication bias and between-study heterogeneity. *Statistics in medicine*, 26(25), 4544–4562.
- Pigott, T. (2012). *Advances in meta-analysis*. Springer Science & Business Media.
- R Core Team. (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Reed, W. R., Florax, R. J., & Poot, J. (2015). *A monte carlo analysis of alternative meta-analysis estimators in the presence of publication bias*. Economics Discussion Papers.
- Rice, K., Higgins, J., & Lumley, T. (2017). A re-evaluation of fixed effect (s) meta-analysis. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*.
- Richard, F. D., Bond Jr, C. F., & Stokes-Zoota, J. J. (2003). One hundred years of social psychology quantitatively described. Educational Publishing Foundation.
- Rothstein, H. R., Sutton, A. J., & Borenstein, M. (2006). *Publication bias in meta-analysis: Prevention, assessment and adjustments*. John Wiley & Sons.
- Rücker, G., Carpenter, J. R., & Schwarzer, G. (2011). Detecting and adjusting for small-study effects in meta-analysis. *Biometrical Journal*, 53(2), 351–368.
- Schmid, C. H. (2017). Heterogeneity: Multiplicative, additive or both? *Research synthesis methods*, 8(1), 119–120.
- 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.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2018). P-curve handles heterogeneity just fine. Retrieved from <http://datacolada.org/67>
- 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.
- Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2015a). Better p-curves: Making p-curve analysis more robust to errors, fraud, and ambitious p-hacking, a reply to ulrich and miller (2015).
- Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2015b). Specification curve: Descriptive and inferential statistics on all reasonable specifications.
- Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2017a). Trim-and-fill is full of it (bias). Retrieved from <http://datacolada.org/30>
- Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2017b). Why p-curve excludes  $ps > .05$ . Retrieved from <http://datacolada.org/61>
- Stanley, D. J. & Spence, J. R. (2014). Expectations for replications: Are yours realistic? *Perspectives on Psychological Science*, 9(3), 305–318.
- Stanley, T. (2017). Limitations of PET-PEESE and other meta-analysis methods. *Social Psychological and Personality Science*, 1948550617693062.
- Stanley, T., Carter, E. C., & Doucouliagos, H. (in press). What meta-analyses reveal about the replicability of psychological research. Retrieved from [http://www.deakin.edu.au/\\_\\_data/assets/pdf\\_file/0007/1198456/WhatMeta-AnalysesReveal\\_WP.pdf](http://www.deakin.edu.au/__data/assets/pdf_file/0007/1198456/WhatMeta-AnalysesReveal_WP.pdf)
- Stanley, T. & Doucouliagos, H. (2014). Meta-regression approximations to reduce publication selection bias. *Research Synthesis Methods*, 5(1), 60–78.
- Stanley, T. & Doucouliagos, H. (2015). Neither fixed nor random: Weighted least squares meta-analysis. *Statistics in Medicine*, 34(13), 2116–2127.
- Stanley, T. & Doucouliagos, H. (2017). Neither fixed nor random: Weighted least squares meta-regression. *Research synthesis methods*, 8(1), 19–42.

- Stanley, T., Doucouliagos, H., & Ioannidis, J. (2017). Finding the power to reduce publication bias. *Statistics in Medicine*, *36*(10), 1580–1598.
- Sterling, T. D., Rosenbaum, W. L., & Weinkam, J. J. (1995). Publication decisions revisited: The effect of the outcome of statistical tests on the decision to publish and vice versa. *The American Statistician*, *49*(1), 108–112.
- Sterne, J. A., 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.
- Tackett, J. L., McShane, B. B., Bockenholt, U., & Gelman, A. (2017). Large scale replication projects in contemporary psychological research. *arXiv:1710.06031 [stat]*. arXiv: 1710.06031. Retrieved from <http://arxiv.org/abs/1710.06031>
- 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.
- van Aert, R. C. (2017). *Puniform: Meta-analysis methods correcting for publication bias*. R package version 0.0.3.
- van Aert, R. C. M. (2018). *Meta-analysis: Shortcomings and potential* (Doctoral dissertation, Tilburg University). Retrieved from <https://doi.org/10.31222/osf.io/eqhjd>
- van Aert, R. C. M. & van Assen, M. A. L. M. (2018). Correcting for Publication Bias in a Meta-Analysis with the P-Uniform\* Method. doi:10.31222/osf.io/zqjr9
- van Aert, R. C., Wicherts, J. M., & van Assen, M. A. (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.
- van Assen, M. A., van Aert, R., & Wicherts, J. M. (2015). Meta-analysis using effect size distributions of only statistically significant studies. *Psychological Methods*, *20*(3), 293.
- van Elk, M., Matzke, D., Gronau, Q., Guan, M., Vandekerckhove, J., & Wagenmakers, E.-J. (2015). Meta-analyses are no substitute for registered replications: A skeptical perspective on religious priming. *Frontiers in Psychology*, *6*, 1365. doi:10.3389/fpsyg.2015.01365
- 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*. Retrieved from [osf.io/myu9c](https://osf.io/myu9c)
- Veroniki, A. A., Jackson, D., Viechtbauer, W., Bender, R., Bowden, J., Knapp, G., ... Salanti, G. (2016). Methods to estimate the between-study variance and its uncertainty in meta-analysis. *Research synthesis methods*, *7*(1), 55–79.
- Vevea, J. L. & Woods, C. M. (2005). Publication bias in research synthesis: Sensitivity analysis using a priori weight functions. *Psychological methods*, *10*(4), 428.
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, *36*(3), 1–48.