Testing for funnel plot asymmetry of standardized mean differences

James E. Pustejovsky¹ & Melissa Rodgers¹

¹ The University of Texas at Austin

Forthcoming in *Research Synthesis Methods*. This paper is not the version of record and may not exactly replicate the final, published version of the article. The final article will be available, upon publication, via its DOI: 10.1002/jrsm.1332

Abstract

Publication bias and other forms of outcome reporting bias are critical threats to the validity of findings from research syntheses. A variety of methods have been proposed for detecting selective outcome reporting in a collection of effect size estimates, including several methods based on assessment of asymmetry of funnel plots, such as Egger's regression test, the rank correlation test, and the Trim-and-Fill test. Previous research has demonstrated that Egger's regression test is mis-calibrated when applied to log-odds ratio effect size estimates, due to artifactual correlation between the effect size estimate and its standard error. This study examines similar problems that occur in meta-analyses of the standardized mean difference, a ubiquitous effect size measure in educational and psychological research. In a simulation study of standardized mean difference effect sizes, we assess the Type I error rates of conventional tests of funnel plot asymmetry, as well as the likelihood ratio test from a three-parameter selection model. Results demonstrate that the conventional tests have inflated Type I error due to correlation between the effect size estimate and its standard error, while tests based on either a simple modification to the conventional standard error formula or a variance-stabilizing transformation both maintain close-to-nominal Type I error.

Keywords: meta-analysis; standardized mean difference; outcome reporting bias; publication bias

Introduction

Publication bias and other forms of outcome reporting bias are critical threats to the validity of findings from research syntheses.¹ These biases arise when researchers selectively report results or when journals selectively publish studies based on the statistical significance or size of effect estimates. Evidence indicates that selective outcome reporting and selective publication practices are at work across a range of disciplines, including education,² psychology,^{3,4} and clinical medicine.⁵ One important method for mitigating the threat of publication bias is to search for and include unpublished or “grey” literature in research syntheses. On a broader scale, systemic reform efforts such as study pre-registration^{6,7} and registered reports^{8,9} aim to re-align research incentives so as to prevent selective publication and outcome reporting. However, until such practices are more widely and rigorously adopted, there will remain a need to examine potential outcome reporting biases during the analysis phase of most research synthesis projects.¹⁰

A variety of methods have been proposed for detecting and, in some cases, adjusting for selective outcome reporting in a collection of effect size estimates. One of the oldest detection techniques is to visually inspect a funnel plot, which is simply a scatter plot of the effect size estimates versus a measure of their precision, such as sample size, standard error, variance, or the inverse of any of these.^{11,12} An asymmetric funnel plot, with an absence of points in the corner of the funnel closer to null, suggests that the effects may have been selected for publication based on statistical significance. The rank correlation test,¹³ Egger’s regression test,¹⁴ and one variant of the Trim-and-Fill algorithm¹⁵ all provide formal statistical tests for funnel plot asymmetry. The Trim-and-Fill technique also generates an adjusted estimate of the overall average effect size. More recently, Stanley and Doucouliagos¹⁶ have proposed using versions of Egger’s regression to estimate the overall average effect size while correcting for funnel plot asymmetry. Other methods of testing and adjusting for publication bias include selection models,^{17,18} the p-curve¹⁹ and p-uniform²⁰ techniques; and the test of excess significance.⁵

A limitation of both visual inspection and formal asymmetry tests is that publication bias is not the only possible explanation for an asymmetric funnel plot. A variety of other mechanisms, such as researchers’ choices about inclusion criteria or the difficulty of implementing interventions with fidelity at large scale, can lead smaller studies to have relatively larger effects, thus inducing asymmetry.^{21–23} Asymmetric funnel plots might also be a product of choosing sample sizes based on accurate priors about effect magnitude²⁴ or choosing measurement procedures to compensate for sample size constraints.²⁵ Still, funnel plot asymmetry tests remain accessible and widely used tools for diagnosing the *possibility*

An earlier version of this paper was presented at the 2018 annual convention of the American Educational Research Association in New York, NY. The authors gratefully acknowledge critical feedback from Dr. Betsy Becker. Supplementary materials for this article are available on the Open Science Framework at <https://osf.io/4ntcy/>.

Correspondence concerning this article should be addressed to James E. Pustejovsky, The University of Texas at Austin, Educational Psychology Department, 1912 Speedway, MS D5800, Austin, TX 78712. E-mail: pusto@austin.utexas.edu

of selective publication, signaling the need for further probing and investigation.

Egger's regression test is one of the most commonly used methods for funnel plot asymmetry. This method entails estimating a meta-regression that includes the standard errors of the effect size estimates as the predictor, weighted by the inverse sampling variances.¹⁴ Statistical significance of the slope coefficient indicates funnel plot asymmetry, potentially due to selective outcome reporting. Subsequent methodological research has demonstrated that the original form of Egger's regression test does not work properly for certain effect size measures. For log odds ratios, the effect size estimate and its standard error are naturally correlated, even in the absence of selective outcome reporting, which causes the test to have inflated Type I error.^{26–28} Macaskill et al²⁶ and Peters et al²⁸ developed modified tests that replace the standard error with the inverse of the study's sample size. Modified versions of Egger's regression test have recently been proposed for raw proportions^{29*} and hazard ratio effect size measures.³¹

Recently, Zwetsloot and colleagues³² highlighted that Egger's regression test suffers from similar problems when applied to standardized mean difference effect size estimates. In a small simulation study, they showed that the test has inflated Type I error due to correlation between the effect size estimate and its standard error. Zwetsloot et al³² therefore recommended a) using a ratio-based effect size measure rather than the standardized mean difference or b) using sample size (rather than the standard error) as the measure of precision in Egger's regression test. However, their simulations focused on conditions observed in pre-clinical animal research studies, which typically have very small samples and may have very large true effects. Thus, the extent to which these the findings generalize to other research areas remains to be investigated.

The standardized mean difference, or d statistic, is a ubiquitous effect size measure in educational and psychological research, where it is used to summarize the difference between two groups on a continuous outcome measure. In this study, we investigate the performance of several tests for funnel plot asymmetry and selective publication in a comprehensive simulation, focusing on conditions that are plausible for psychology and education research. The study builds upon previous findings³² in several respects. First, in addition to demonstrating that conventional funnel plot asymmetry tests have inflated Type I error, we evaluate the performance of two possible modifications to the tests, including use of a modified measure of precision and use of a variance-stabilizing transformation for the d statistic. We describe both of these modifications in fully general terms, so that they can be applied to d -type effect sizes from a range of research designs. Second, we also evaluate the performance of a test based on the three-parameter selection model (3PSM),³³ a direct model for selective publication that does not hinge on funnel plot asymmetry. Several recent simulation studies have indicated that the 3PSM may have performance advantages relative to other methods for investigating selective publication,^{34,35} suggesting that both wider use and further study of the 3PSM is warranted. Third and finally, we use parameters for the simulation study that are informed by empirical research and chosen to emulate real

*Raw proportions refer to measures of prevalence or risk collected from observational studies. For example, Ismail and colleagues synthesized data on the prevalence of depression among patients with mild cognitive impairment, using raw proportions as the effect size metric.³⁰

meta-analyses in psychology and education. Thus, findings from our simulations are relevant to fields where standardized mean differences are commonly used.

In what follows, we explain the problem with the conventional formula for the standard error of the SMD and describe two possible modifications. We then use simulation methods to assess Type I error rates of a selection of funnel plot asymmetry tests, as well as the 3PSM, based on the conventional and modified approaches; we also briefly examine the relative power of the modified tests. Finally, we discuss implications of these findings for how to assess publication bias when meta-analyzing standardized mean differences.

Distribution Theory and Methods

To illustrate the problems with funnel plot asymmetry tests for standardized mean differences, we first review the distribution theory for Cohen's d . We present this using a general formulation that applies to d statistics estimated from a wide variety of different research designs, including not only basic between-subjects experiments but also designs that involve analysis of covariance or change-score adjustment,³⁶ longitudinal designs,³⁷ and cluster-randomized designs.³⁸

For a given study i , the standardized mean difference parameter has the form $\delta_i = \Delta_i/\sigma_i$, where Δ_i is the mean difference between two groups in the units of the outcome scale and σ_i is the standard deviation of the outcome (in one or both of the groups). Let $\hat{\Delta}_i$ and $\hat{\sigma}_i^2$ denote approximately unbiased sample estimates of Δ_i and σ_i^2 , respectively, for $i = 1, \dots, k$. Let f_i denote the degrees of freedom corresponding to the estimated standard deviation. Cohen's d estimator of the standardized mean difference is $d_i = \hat{\Delta}_i/\hat{\sigma}_i$. A delta method approximation for the sampling variance of d_i is³⁸

$$\text{Var}(d_i) \approx \frac{\text{Var}(\hat{\Delta}_i)}{\sigma_i^2} + \frac{\delta_i^2}{2f_i}. \quad (1)$$

The first term in (1) captures the contribution of the variance of $\hat{\Delta}_i$, while the second term captures the contribution of the variance of $\hat{\sigma}_i$ —that is, how precisely estimated is the *scale* of the outcome. For some study designs, the first term is known, as are the degrees of freedom, while in other designs, the first term and the degrees of freedom involve further parameters that must be estimated or imputed.* However, across study designs, the second term in (1) must generally be estimated because it involves the unknown effect size δ_i . The sampling variance of d_i is typically estimated as

$$V_i^d = W_i + \frac{d_i^2}{2f_i}, \quad (2)$$

*For example, Hedges³⁸ provides formulas for $\text{Var}(\hat{\Delta}_i)/\sigma_i^2$ and the degrees of freedom in a two-level cluster-randomized trial with a continuous outcome. Both quantities depend on the intra-class correlation of the outcome. The degrees of freedom are based on a Satterthwaite-type approximation, in which $f_i = 2\sigma_i^4/\text{Var}(\hat{\sigma}_i^2)$.

where $W_i = \text{Var}(\hat{\Delta}_i)/\sigma_i^2$ or an estimate thereof.* The standard error of d_i is taken as the square root of the sampling variance.

Using the usual estimator of the sampling variance of d_i for purposes of diagnosing funnel plot asymmetry is problematic because asymmetry tests are predicated on the null hypothesis of no dependence between d_i and $\sqrt{V_i^d}$. Because V_i^d is calculated from the effect size estimate itself, the variance will generally be correlated with d_i even in the absence of selective publication processes or any other mechanisms that could produce asymmetry. This will lead to rejection of the null more often than the nominal Type I error rate.³²

It is unclear whether this issue might also carry over to selection models such as the 3PSM. Current implementations of the 3PSM assume that the effect size estimates are normally distributed with known sampling variances, but in practice the variances must be estimated. Hedges and Vevea¹⁸ argued that sampling error in the variance estimate should not generally affect the performance of selection models. However, the robustness of the 3PSM to the assumption of known variances has not been thoroughly studied, and it is possible that its performance could be improved through modification.

We shall consider two potential solutions to the problem of dependence between d_i and V_i^d . One possible solution is simply to substitute $\sqrt{W_i}$ in place of $\sqrt{V_i^d}$ when conducting tests of funnel plot asymmetry or selective publication.[†] This will remove the artifactual correlation between the effect size estimate and the measure of precision. A further justification for using $\sqrt{W_i}$ is that statistical significance tests are based on the t statistic for the *unstandardized* mean difference Δ_i , $t_i = \hat{\Delta}_i/\sqrt{\text{Var}(\hat{\Delta}_i)} = d_i/\sqrt{W_i}$, which is invariant to the scale of the outcome. The statistical significance of the effect size estimate (and the consequent probability that the result is published or otherwise reported) is a therefore function of the $\sqrt{W_i}$ term rather than of $\sqrt{V_i^d}$. It thus seems reasonable to use $\sqrt{W_i}$ as the predictor in regression tests and other procedures for detecting funnel plot asymmetry.

A second possible solution is to apply a transformation to d , such that the variance of the transformed effect estimate does not involve the effect size parameter itself. This approach is similar to use of the Fisher- z transformation for the correlation coefficient, where the transformed effect estimate has sampling variance that is a simple function of sample size. The variance-stabilizing transformation for the standardized mean difference is given by

$$h_i = h(d_i) = \sqrt{2} \times \text{sgn}(d_i) \left[\log \left(|d_i| + \sqrt{d_i^2 + a_i^2} \right) - \log(a_i) \right], \quad (3)$$

where $a_i = \sqrt{2W_i f_i}$.³⁹ The transformed effect size has approximate sampling variance

*If Δ_i is a simple difference in means between two groups with sample sizes n_{0i} and n_{1i} and common variance σ_i , then $W_i = (n_{0i} + n_{1i})/(n_{0i}n_{1i})$ and $f_i = n_{0i} + n_{1i} - 2$.³⁶ In more complex designs, W_i can be estimated as $V_i^\Delta/\hat{\sigma}_i^2$, where V_i^Δ is the estimated sampling variance of $\hat{\Delta}_i$.

[†]In a simple two-group experiment with equally sized groups, $\sqrt{W_i} = 2/\sqrt{n_i}$, where n_i is the per-group sample size. However, in unbalanced designs or designs with more complicated features, W_i will depend on other features of the design beyond just n_i . Thus, funnel plot asymmetry tests based on $\sqrt{W_i}$ will not necessarily be equivalent to tests based on $1/\sqrt{n_i}$.

$V_i^h = 1/f_i$, which does not depend on δ_i .

It is not clear on the basis of theory which, if either, of these approaches maintains correct Type I error and which yields more powerful tests for funnel plot asymmetry and selective publication. To investigate these questions and provide guidance about what methods should be used in practice, we turned to Monte Carlo simulations.

Simulations

We conducted Monte Carlo simulations to investigate the performance of conventional tests for funnel plot asymmetry and selective outcome reporting, as well as modified versions of the same tests based on either a) using W_i in place of V_i^d or b) using the transformed effect size and sampling variance, h_i and V_i^h . Our primary focus was on the Type I error rates of the tests when effect size estimates are generated without publication bias. In addition, for the subset of tests that maintain nominal Type I error rates, we examined power under varying degrees of funnel plot asymmetry induced by selection of statistically significant results.

Methods

The simulations used a random effects data-generating process, in which the effect size parameters from each primary study are sampled from a distribution with mean μ and standard deviation τ . Given a true effect size δ_i , we simulated standardized mean difference estimates based on a two-group experiment with equal sample sizes, where the sample sizes are chosen independently of the true effect sizes. Effect size estimates and sampling variances were calculated using Hedges' g correction term to adjust for small-sample bias.⁴⁰ For each simulated dataset, we conducted tests for funnel plot asymmetry and selective outcome reporting using five different conventional procedures:

1. the rank correlation test;
2. Egger's regression test, using the standard error $\sqrt{V_i^d}$ as the predictor;
3. Egger's regression test, using the estimated sampling variance V_i^d as the predictor;
4. the R0 variant of the Trim-and-Fill algorithm; and
5. the likelihood ratio test based on the 3PSM.¹⁸

We then repeated the same tests a) using W_i in place of V_i^d and b) using h_i and V_i^h in place of d_i and V_i^d .

For the regression tests, several different techniques have been proposed for estimating standard errors and conducting hypothesis tests. We examined three different variants, all of which involve estimating the regression using weights proportional to the inverse sampling variance. Egger and colleagues¹⁴ originally proposed calculating standard errors based on the assumption that random effects are proportional to the sampling variances (i.e., a multiplicative dispersion model). Alternately, one could assume that random effects are additive (as in the conventional random effects model) and estimate their variance using the DerSimonian and Laird⁴¹ moment estimator. This approach can also be combined with

the small-sample adjustments proposed by Knapp and Hartung.⁴² Finally, one could use robust standard errors that are valid under arbitrary forms of heteroskedasticity.⁴³

Simulations were carried out using the R (Version 3.4.0) statistical computing environment.⁴⁴ The rank correlation test and Trim-and-Fill test were calculated using the `metafor` package;⁴⁵ the likelihood ratio test was calculated using the `weightr` package;⁴⁶ the variants of Egger’s regression tests were calculated using custom-written R code. The supplementary materials provide further details about the implementation of the tests.⁴⁷

The simulations varied five factors: the number of studies included in the meta-analysis (k), the distribution of sample sizes in the primary studies, the magnitude of the average effect (μ), the amount of heterogeneity (τ), and the strength of the selective publication process. For each of the factors, we used parameter values similar to those found in meta-analyses of psychological and educational research, where standardized mean difference effect sizes are prevalent. Simulation results are summarized across 8000 replications. The supplementary materials⁴⁷ include R code for implementing all of the simulations and complete numerical results.

Number of primary studies per meta-analysis. In a recent review of meta-analyses of educational and psychological research, Polanin et al² found that meta-analyses included an average of 79 primary studies; a similar review by Ahn et al⁴⁸ reported an average of 55 primary studies per meta-analysis. Meta-analyses of as few as 10 studies were observed in both reviews, but meta-analyses with so few studies generally lack power to make assessments of selective outcome reporting. We therefore examined sample sizes of $k = 20, 50, 80$ studies per meta-analysis, consistent with the range of the studies included in recent syntheses in psychology and education.

Primary study sample sizes. In meta-analyses of psychological research, study-level sample sizes are often small, on average, and positively skewed.⁴⁹ To simulate meta-analyses that have realistic sample size distributions, we drew sample sizes at random from a recent synthesis of the literature on ego depletion effects,^{50,51} which included 118 unique effect sizes. This review is a particularly useful reference point because nearly half of the effect sizes ($n = 48$) were from unpublished experiments and half ($n = 59$) were not statistically significant. Total sample sizes ranged from 20 to 251, with a median of 48 (inter-quartile range: 31-68). To explore the consequences of different study-level sample size distributions, we multiplied all sample sizes from the Carter et al⁵⁰ review by a factor of 1, 2, or 3, yielding median sample sizes of 48, 96, or 144.

Recent meta-analytic reviews of studies from educational and psychological journals found sample size distributions similar to the range we explored. In a review of studies published in top journals of social-personality psychology research, Fraley and Vazire⁴⁹ reported an overall median sample size of 104. Sample sizes also varied depending on the sub-field, with medians of 79 for attitudes and social cognition, 95 for interpersonal relations and group processes, and 122 for personality processes and individual differences. In education research, a review of articles published from 2010 to 2015 in *Review of Educational Research* found a median study-level sample size of 171 with an interquartile range of 66 to 702.⁵²

Average effect size and between-study heterogeneity. Mean effect sizes were varied across a wide range of values, including $\mu = 0.0, 0.1, 0.2, \dots, 1.0$. This broad range of average effects allows for an assessment of trends in the performance of the procedures tested. For heterogeneity, we examined standard deviations of $\tau = 0.0$, corresponding to a constant-effect model, as well as three degrees of effect heterogeneity ($\tau = 0.1, 0.2, 0.4$). These standard deviations of the true effects were determined by first considering cut-points suggested by Pigott⁴⁰ for small (25%), medium (50%), and large (75%) values of I^2 . Specifically, combining the average variance in the Carter et al⁵⁰ dataset ($\bar{V} = 0.09$) and the three selected values of τ yields I^2 values that correspond closely to these cut-points. In a recent review of 200 psychology meta-analyses, Stanley et al⁵³ estimated a median I^2 of 74%, which in their sample translated to $\tau = 0.35$ among the true effects. In addition, a review of 705 *Psychological Bulletin* meta-analyses published between 1990 and 2013 van Erp et al⁵⁴ reported an interquartile range in I^2 estimates of 33% to 87%—quite similar to the Pigott et al⁴⁰ cut-points. This range of values is also consistent with previous simulation studies of publication bias methods (e.g., Carter et al³⁵ used $\tau = 0, 0.2, 0.4$).

Selective publication. We examined data-generating processes that do not involve selective publication as well as conditions that involve selection on the basis of the direction and statistical significance of study results. Following methods similar to recent simulation studies,^{27,35,51} we used a simple selection model to induce varying degrees of reporting bias and funnel plot asymmetry. In this model, statistical significance is operationalized as a one-sided p-value of less than .025; all statistically significant effect sizes are observed; and effect sizes that are not statistically significant are observed with probability π . Effect size estimates are simulated until k studies are observed. We examined conditions with no selective publication ($\pi = 1$) and varying degrees of selection, ranging from very mild selection ($\pi = 0.8$) to absolute selection ($\pi = 0.0$), in which only statistically significant results are observed. For conditions with selective publication ($\pi < 1$), we limited the levels of the mean effect to $\mu = 0.0, 0.4$, or 0.8 .

Convergence of the 3PSM. Maximum likelihood estimates of the 3PSM were obtained using iterative numerical algorithms that did not always converge, which precluded use of the likelihood ratio test. Non-convergence occurred under three distinct conditions: 1) when heterogeneity estimates were near zero, 2) when none or very few of the effect estimates were statistically significant, or 3) when all or nearly all of the estimates were statistically significant. Under condition 1), we re-estimated the 3PSM and conducted the likelihood ratio test under the constraint that $\tau = 0$. Under condition 2), when very few effect estimates were statistically significant and the 3PSM did not converge, we assumed that a researcher’s practical conclusion would be to fail to reject the null. Unfortunately, no clear analogue was available under condition 3), where all or nearly all estimates were statistically significant. We therefore calculated rejection rates of the 3PSM likelihood ratio test after excluding such replications. In the results that follow, rejection rates for the 3PSM are only depicted for conditions where no more than 25% of replications were excluded due to non-convergence under condition 3). The supplementary materials provide further information about convergence rates.⁴⁷

Results: Type I error rates

Type I error rates of the conventional and modified funnel plot asymmetry tests were calculated for the conditions without selective publication, using the conventional level of $\alpha = .05$. Figure 1 depicts the rejection rates for five tests of funnel plot asymmetry or selective publication, for a sample size of $k = 50$ studies, and for varying levels of average effect size (on the horizontal axis), heterogeneity (rows), and study sample size (columns). Both versions of Egger's regression test use the additive random effects model and Knapp-Hartung adjustment for calculating standard errors.

At the smallest level of within-study sample size (left column), the error rates of all of the tests tend to be above the nominal level whenever the mean effect is greater than zero. Comparing the top and bottom rows of results, it can be seen that the extent of inflation is more severe when the between-study heterogeneity is larger. As the distribution of primary study sample sizes increases, the rejection rates of all of the procedures decrease and grow closer to the nominal α . However, even with the largest sample size distribution, the funnel plot asymmetry tests still reject at above-nominal levels when the effect sizes are heterogeneous. Notably, the rejection rate of the 3PSM grows closer to nominal more quickly than all of the other tests and performs well regardless of between-study heterogeneity.

The supplementary materials include graphs comparing the same conditions and publication bias detection methods for meta-analyses of $k = 20$ or 80 studies.⁴⁷ Across conditions, the performance of the tests when $k = 80$ is similar to when $k = 50$. For meta-analyses with 20 studies, rejection rates are closer to or below $\alpha = .05$ when heterogeneity is low ($\tau \leq 0.1$), regardless of the sample size distribution. When heterogeneity increases to $\tau = 0.4$, Type I error rates with $k = 20$ are similar to those with $k = 50$ and no heterogeneity. Overall, the Type-I error calibration of the funnel plot asymmetry tests degrades as the number of studies increases, the primary study sample size distribution decreases, and heterogeneity increases. The 3PSM maintains rejection rates that are closer to nominal regardless of the number of studies included, although rates still exceed nominal under some conditions.

Figure 2 depicts the same conditions as Figure 1 (again for samples of $k = 50$ studies), but for the modified tests based on $\sqrt{W_i}$ rather than $\sqrt{V_i^d}$. The regression tests again use the additive random effects and Knapp-Hartung adjustment for calculating standard errors. When there is no heterogeneity (top row), most of the funnel plot asymmetry tests reject at or below the nominal level. The exception is Trim-and-Fill, which tends to over-reject slightly when the average effect size is large ($\mu \geq 0.75$). As heterogeneity increases, the Trim-and-Fill method slightly exceeds the nominal level when based on the $\sqrt{W_i}$ method, but not to the extent of the conventional $\sqrt{V_i^d}$ method. In the presence of any degree of heterogeneity ($SD = 0.1, 0.2, 0.4$), the rejection rates of the rank correlation test are consistently below the nominal α . The rejection rates of the likelihood ratio test based on the 3PSM are consistently at or below the nominal level, even under conditions of high heterogeneity. However, the 3PSM frequently failed to converge under conditions with large average effect size and low heterogeneity.* Across tests, results are consistent when based on

*The 3PSM has convergence problems when the sample is comprised of studies that are all or nearly all

Figure 1. Type I error rates for five tests of funnel plot asymmetry or selective publication based on conventional standard error $\sqrt{V_i^d}$, for samples of $k = 50$ studies. Results for the 3PSM are omitted for conditions where less than 75% of replications converged.

Figure 2. Type I error rates for five modified tests of funnel plot asymmetry or selective publication based on $\sqrt{W_i}$, for samples of $k = 50$ studies. Results for the 3PSM are omitted for conditions where less than 75% of replications converged.

$k = 20$ or $k = 80$ studies included in the meta-analysis; the supplementary materials provide further details.⁴⁷

Figure 3 depicts the same conditions as the previous figures, but for the modified tests based on the transformed effect size h_i and standard error $\sqrt{V_i^h}$. Generally, the results are very similar to those that use $\sqrt{W_i}$: Rejection rates of the regression tests are very close to nominal; those of the rank correlation test are consistently below nominal levels; and those of the Trim-and-Fill test can be in excess of nominal when within-study sample sizes are small and the mean effect size is large. Rejection rates for the 3PSM are also at or below nominal levels, although the model frequently failed to converge under conditions with large average effects and low heterogeneity. Convergence rates for the 3PSM test based on h_i and $\sqrt{V_i^h}$ were consistently worse than those based on d_i and $\sqrt{W_i}$.

As noted in the previous section, there are several alternative techniques for calculating standard errors for the regression tests. Figure 4 depicts the Type I error rates for the modified versions of the regression tests based on $\sqrt{W_i}$, with standard errors calculated based on a multiplicative dispersion model, an additive random effects model, or heteroskedasticity-robust variance estimation.* In the case of no heterogeneity ($SD = 0$), all methods have rejection rates at or below the nominal level. In the presence of heterogeneity, the additive standard errors have rejection rates closest to $\alpha = 0.05$, with both predictors performing similarly. The error rates of the tests based on robust standard errors are very close to nominal for low or moderate degrees of heterogeneity, but tend to slightly exceed the nominal level for high levels of heterogeneity. The multiplicative method has the worst performance with regard to Type I error rates, with above-nominal rejection rates at higher levels of between-study heterogeneity. The same performance trends are evident with lower or higher numbers of studies in the meta-analysis and with other within-study sample size distributions.

Results: Power of modified tests

Although not our primary focus, we briefly summarize results comparing the relative power of the modified funnel plot asymmetry tests and 3PSM likelihood ratio tests. Our discussion is limited to the modified tests because the conventional tests do not adequately control Type I error.

Figure 5 presents the power levels (vertical axis) of all five modified tests based on $\sqrt{W_i}$ as a function of the degree of selective publication (horizontal axis, running from 1 to 0) for a sample size of $k = 50$ studies, study sample size of $n \times 1$, and for varying levels of average effect size and heterogeneity. As the probability of publishing a nonsignificant

statistically significant. This is more likely to occur under conditions where the true effects are large and the degree of heterogeneity is small or where there is strong selective publication.

*Results for the modified tests based on h_i and V_i^h are included in the supplementary materials.⁴⁷ The pattern of results is very similar to those depicted in Figure 4. The supplementary materials also include results for regression tests based on the modified measure of precision that use W_i as the predictor variable. The pattern of results is again very similar to those for the tests that use $\sqrt{W_i}$ as the predictor, but the degree of Type I error inflation is less pronounced.

Figure 3. Type I error rates for five modified tests of funnel plot asymmetry or selective publication based on h_i and $\sqrt{V_i^h}$, for samples of $k = 50$ studies. Results for the 3PSM are omitted for conditions where less than 75% of replications converged.

Figure 4. Type I error rates for modified regression test based on $\sqrt{W_i}$, using different methods for calculating standard errors, for samples of $k = 50$ studies.

Figure 5. Power to detect publication bias for five tests of funnel plot asymmetry or selective publication based on modified standard error $\sqrt{W_i}$, for samples of $k = 50$ studies with $n \times 1$.

result (π) decreases towards zero, the degree of selective publication increases and power rates of all of the tests improve. Across conditions and degrees of selective publication, the 3PSM has substantially higher power than the other tests. The funnel plot asymmetry tests have little to no power when the true effect size is $\mu = 0$. When $\mu = 0.4$, the power of the regression test and rank correlation test reaches adequate levels of 80% only under very strong selective publication (i.e., $\pi \leq .1$), while the Trim-and-Fill test has very little power except when selection is very strong. As between-study heterogeneity increases, power decreases for all tests except the 3PSM. For average effect sizes greater than zero, the power of all of the funnel plot asymmetry tests increases with the number of studies included in the meta-analysis; see the supplementary materials for further details.⁴⁷ The same trends are evident in the power of the modified tests based on h_i and $\sqrt{V_i^h}$.*

Given that both the tests based on the modified measure of precision and the tests based on the variance-stabilizing transformation adequately control Type I error in the absence of publication bias, it is reasonable to compare the relative power of the two approaches. Figure 6 depicts the difference in power between the two approaches for each of the funnel plot asymmetry and publication bias tests; positive values correspond to higher power for tests based on the variance-stabilizing transformation. The variance-stabilizing approach has a consistent, albeit relatively small, advantage over the approach based on the modified measure of precision. The advantage is stronger when overall average effects are larger—exactly the conditions where there is a higher degree of dependence between the effect size estimate and its sampling variance. The advantage is also more pronounced for the regression tests than for the 3PSM likelihood ratio test. For a given test, the power advantage of the variance-stabilizing transformation is reduced when within-study sample sizes are larger (i.e., $n \times 2$ or $n \times 3$; see supplementary materials for details).

Discussion

This study has demonstrated that conventional funnel plot asymmetry tests do not adequately control Type I error rates when used with standardized mean difference effect sizes and conventional variance estimates, particularly when within-study sample sizes are relatively small or between-study heterogeneity is high. However, using either a modified measure of precision ($\sqrt{W_i}$ or W_i) or a variance-stabilizing transformation removes the artifactual correlation with the effect size estimates and results in tests that have the correct Type I error rates. For Egger’s regression tests, the modified approaches maintain correct Type I error rates when inference is conducted using an additive random effects model and Knapp-Hartung correction, though not when using the multiplicative dispersion model as originally proposed.

The logic of tests based on the modified measure of precision aligns closely with alternative regression tests for funnel plot asymmetry that have been developed for odds ratio, proportion, and hazard ratio effect size measures.^{26–29,31} Although less familiar and less

*The supplementary materials⁴⁷ include graphs depicting the power of the modified tests based on h_i and $\sqrt{V_i^h}$.

Figure 6. Differences in power between tests based on transformed effect size h_i and $\sqrt{V_i^h}$ versus tests based on modified standard error $\sqrt{W_i}$, for samples of $k = 50$ studies with $n \times 1$.

simple to apply, funnel plot asymmetry tests based on the variance-stabilizing transformation have a modest power advantage over tests based on the modified measure of precision.

The simulations also examined the performance of a likelihood ratio test based on a 3PSM, a direct test for selective publication, which posits a model for the probability that an effect size based on a non-significant result is reported in the literature. The 3PSM test out-performs all of the other methods we examined, in terms of both controlling Type I error rates in the absence of selective publication and achieving power in the presence of varying degrees of selection. Type I error rates of the 3PSM test based on the conventional variance estimator sometimes exceeded nominal levels, although the excess was less severe than for funnel plot asymmetry tests and diminished more quickly with increased within-study sample size. In practice, we recommend conducting 3PSM tests based on the modified measure of precision because this approach adequately controlled Type I error rates and had higher convergence rates than when based on the variance-stabilizing transformation.

The results of the present study align with the findings of other recent simulations,^{34,35} which evaluated the 3PSM with the conventional variance estimator and under other data-generating process. Although the 3PSM still remains less widely known than Eggers' regression and other tests of funnel plot asymmetry, recently developed, easy-to-use software tools⁴⁶ make the technique much more accessible to applied researchers than in the past. The basic 3PSM has been extended in several ways, including incorporation of covariates to explain variation in average effect sizes,⁵⁵ and further extensions continue to be developed and studied.⁵⁶⁻⁵⁹

Findings from the simulation study have several limitations that should be noted. Although we have described the modified funnel plot asymmetry tests in general terms that could be applied to effect size estimates from a range of different study designs, the simulations were limited to SMD effect size estimates generated from simple, two-group between-subjects designs, as might be found in brief psychological experiments conducted in laboratory settings. Research syntheses in other contexts might require calculating effect sizes from studies that have considerably more complex experimental designs (e.g., longitudinal designs, cluster-randomized designs). We believe that it is reasonable to expect that the main findings of the study will generalize to these more complex scenarios, although the exact levels of Type I error and power reported here are likely to vary.

The effective performance of the 3PSM is likely due to the fact that the estimation model is exactly aligned with the data-generating process of the simulation. Under more complex forms of selection bias, the properties and relative advantages of the 3PSM and the funnel plot asymmetry tests might differ.³⁴ Similarly, the findings might differ if the data-generating process involves p-hacking or other questionable research practices.³⁵ The 3PSM also bears an important limitation in that it can only be calculated when the model estimates converge, which makes it impossible to fit based when the sample includes *only* statistically significant results. This situation occurs more frequently when study sample sizes are large and true effects are large and consistent (with low heterogeneity). Future work should examine alternative procedures, such as score tests, based upon the 3PSM.

A final limitation is that we have only simulated meta-analytic samples where the effect

size estimates are all independent. In practice, meta-analyses in psychology, education, and other disciplines very often include studies that contribute multiple, statistically dependent effect size estimates.⁶⁰ A diverse array of methods for meta-analysis and meta-regression of dependent effect sizes are now available.^{61–64} However, there remains a dearth of methods for testing or adjusting for selective outcome reporting under these conditions. Of the tests that we have considered, only Egger’s regression test can readily be applied to meta-analytic samples that include dependent effect size estimates, through the use of robust variance estimation methods.⁶⁵ To our knowledge, the operating characteristics of this method have not been evaluated. A critical avenue for future research is thus to develop and rigorously evaluate methods for investigating publication bias in meta-analytic samples that include dependent effects.

References

1. Rothstein HR, Sutton AJ, Borenstein M. *Publication Bias in Meta-Analysis: Prevention, Assessment and Adjustments*. John Wiley & Sons; 2006.
2. Polanin JR, Tanner-Smith EE, Hennessy EA. Estimating the difference between published and unpublished effect sizes a meta-review. *Review of Educational Research*. 2016;86(1):207-236. <http://rer.sagepub.com/content/86/1/207.short>.
3. Franco A, Malhotra N, Simonovits G. Underreporting in psychology experiments: Evidence from a study registry. *Social Psychological and Personality Science*. 2016;7(1):8-12. doi:10.1177/1948550615598377
4. Ferguson CJ, Heene M. A vast graveyard of undead theories: Publication bias and psychological science's aversion to the null. *Perspectives on Psychological Science*. 2012;7(6):555-561. doi:10.1177/1745691612459059
5. Ioannidis JP, Trikalinos TA. An exploratory test for an excess of significant findings. *Clinical Trials*. 2007;4(3):245-253. doi:10.1177/1740774507079441
6. AllTrials. All trials registered. All results reported. AllTrials. <http://www.alltrials.net/>. Published 2014. Accessed January 19, 2018.
7. ClinicalTrials. ClinicalTrials.Gov. <https://clinicaltrials.gov/>. Published 2018. Accessed February 2, 2018.
8. Chambers CD. Registered Reports: A new publishing initiative at Cortex. *Cortex*. 2013;49(3):609-610. doi:10.1016/j.cortex.2012.12.016
9. Nosek BA, Lakens D. Registered reports: A method to increase the credibility of published results. *Social Psychology*. 2014;45(3):137-141. doi:10.1027/1864-9335/a000192
10. Higgins JPT, Lasserson T, Chandler J, Tovey D, Churchill R. *Methodological Expectations of Cochrane Intervention Reviews (MECIR): Standards for the conduct and reporting of new Cochrane Intervention Reviews, reporting of protocols and the planning, conduct and reporting of updates*. London; 2016.
11. Light RJ, Pillemer DB. *Summing up*. Harvard University Press; 1984.
12. Sterne JA, Egger M. Funnel plots for detecting bias in meta-analysis: Guidelines on choice of axis. *Journal of Clinical Epidemiology*. 2001;54(10):1046-1055.
13. Begg CB, Mazumdar M. Operating characteristics of a rank correlation test for publication bias. *Biometrics*. 1994;50(4):1088. doi:10.2307/2533446
14. Egger M, Smith GD, Schneider M, Minder C. Bias in meta-analysis detected by a simple, graphical test. *BMJ*. 1997;315(7109):629-634.
15. Duval S, Tweedie R. A nonparametric "trim and fill" method of accounting for publication bias in meta-analysis. *Journal of the American Statistical Association*. 2000;95(449):89-98. doi:10.2307/2669529

16. Stanley TD, Doucouliagos H. *Meta-Regression Analysis in Economics and Business*. Vol 5. Routledge; 2012.
17. Copas J, Shi JQ. A sensitivity analysis for publication bias in systematic reviews. *Statistical Methods in Medical Research*. 2001;10:251-265.
18. Hedges L, Vevea J. Selection method approaches. In: *Publication Bias in Meta-Analysis: Prevention, Assessment, and Adjustments*. Chichester, England: John Wiley & Sons; 2005:145-174.
19. Simonsohn U, Nelson LD, Simmons JP. P-curve and effect size: Correcting for publication bias using only significant results. *Perspectives on Psychological Science*. 2014;9(6):666-681. doi:10.1177/1745691614553988
20. Assen MALM van, Van Aert RCM, Wicherts JM. Meta-analysis using effect size distributions of only statistically significant studies. *Psychological Methods*. 2015;20(3):293-309. doi:http://dx.doi.org/10.1037/met0000025
21. Sterne JAC, Sutton AJ, Ioannidis JPA, et al. Recommendations for examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *BMJ*. 2011;343:d4002. doi:10.1136/bmj.d4002
22. Stuck AE, Rubenstein LZ, Wieland D. Bias in meta-analysis detected by a simple, graphical test. Asymmetry detected in funnel plot was probably due to true heterogeneity. *BMJ: British Medical Journal*. 1998;316(7129):469.
23. Tang J-L, Liu JL. Misleading funnel plot for detection of bias in meta-analysis. *Journal of Clinical Epidemiology*. 2000;53(5):477-484. doi:10.1016/S0895-4356(99)00204-8
24. Simonsohn U. The funnel plot is invalid because of this crazy assumption $r(n,d)=0$. <http://datacolada.org/58>. Published March 21, 2017. Accessed April 1, 2018.
25. Morey R. Asymmetric funnel plots without publication bias. <http://bayesfactor.blogspot.com/2016/01/asymmetric-funnel-plots-without.html>. Published January 9, 2016. Accessed April 1, 2018.
26. Macaskill P, Walter SD, Irwig L. A comparison of methods to detect publication bias in meta-analysis. *Stat Med*. 2001;20(4):641-654. doi:10.1002/sim.698
27. Moreno SG, Sutton AJ, Ades A, et al. Assessment of regression-based methods to adjust for publication bias through a comprehensive simulation study. *BMC Medical Research Methodology*. 2009;9(1). doi:10.1186/1471-2288-9-2
28. Peters JL, Sutton AJ, Jones DR, Abrams KR, Rushton L. Comparison of two methods to detect publication bias in meta-analysis. *Journal of the American Medical Association*. 2006;295(6):676-680.
29. Hunter JP, Saratzis A, Sutton AJ, Boucher RH, Sayers RD, Bown MJ. In meta-analyses of proportion studies, funnel plots were found to be an inaccurate method of assessing publication bias. *Journal of Clinical Epidemiology*. 2014;67(8):897-903. doi:10.1016/j.jclinepi.2014.03.003

30. Ismail Z, Elbayoumi H, Fischer CE, et al. Prevalence of depression in patients with mild cognitive impairment: A systematic review and meta-analysis. *JAMA Psychiatry*. 2017;74(1):58-67.
31. Debray TP, Moons KG, Riley RD. Detecting small-study effects and funnel plot asymmetry in meta-analysis of survival data: A comparison of new and existing tests. *Research Synthesis Methods*. 2018. doi:10.1002/jrsm.1266
32. Zwetsloot PP, Van Der Naald M, Sena ES, et al. Standardized mean differences cause funnel plot distortion in publication bias assessments. *eLife*. 2017;6:1-20. doi:10.7554/eLife.24260
33. Hedges LV, Vevea JL. Estimating effect size under publication bias: Small sample properties and robustness of a random effects selection model. *Journal of Educational and Behavioral Statistics*. 1996;21(4):299. doi:10.2307/1165338
34. McShane BB, Böckenholt U, Hansen KT. Adjusting for publication bias in meta-analysis an evaluation of selection methods and some cautionary notes. *Perspectives on Psychological Science*. 2016;11(5):730-749. <http://pps.sagepub.com/content/11/5/730.short>. Accessed November 1, 2016.
35. Carter E, Schönbrodt F, Gervais WM, Hilgard J. Correcting for bias in psychology: A comparison of meta-analytic methods. 2017. <https://osf.io/preprints/psyarxiv/9h3nu/download?format=pdf>.
36. Borenstein M. Effect sizes for continuous data. In: Cooper HM, Hedges LV, Valentine JC, eds. *The Handbook of Research Synthesis and Meta-Analysis*. New York, NY: Russell Sage Foundation; 2009:221-236.
37. Feingold A. Effect sizes for growth-modeling analysis for controlled clinical trials in the same metric as for classical analysis. *Psychological Methods*. 2009;14(1):43-53. doi:10.1037/a0014699
38. Hedges LV. Effect sizes in cluster-randomized designs. *Journal of Educational and Behavioral Statistics*. 2007;32(4):341-370. doi:10.3102/1076998606298043
39. Hedges LV, Olkin I. *Statistical Methods for Meta-Analysis*. Orlando, FL: Academic Press; 1985.
40. Pigott T. *Advances in Meta-Analysis*. Springer Science & Business Media; 2012.
41. DerSimonian R, Laird N. Meta-analysis in clinical trials. *Controlled Clinical Trials*. 1986;7(3):177-188.
42. Knapp G, Hartung J. Improved tests for a random effects meta-regression with a single covariate. *Statistics in Medicine*. 2003;22(17):2693-2710. doi:10.1002/sim.1482
43. Sidik K, Jonkman JN. Robust variance estimation for random effects meta-analysis. *Computational Statistics & Data Analysis*. 2006;50(12):3681-3701. doi:10.1016/j.csda.2005.07.019

44. R Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing; 2017. <https://www.R-project.org/>.
45. Viechtbauer W. Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*. 2010;36(3):1-48.
46. Coburn KM, Vevea JL. *Weightr: Estimating Weight-Function Models for Publication Bias*; 2017. <https://CRAN.R-project.org/package=weightr>.
47. Pustejovsky JE, Rodgers MA. Testing for funnel plot asymmetry of standardized mean differences: Supplementary materials. July 2018. osf.io/4ntcy.
48. Ahn S, Ames AJ, Myers ND. A review of meta-analyses in education: Methodological strengths and weaknesses. *Review of Educational Research*. 82(4):436-476. doi:10.3102/0034654312458162
49. Fraley RC, Vazire S. The n-pact factor: Evaluating the quality of empirical journals with respect to sample size and statistical power. *PLoS ONE*. 2014;9(10):e109019. doi:10.1371/journal.pone.0109019
50. Carter EC, Kofler LM, Forster DE, McCullough ME. A series of meta-analytic tests of the depletion effect: Self-control does not seem to rely on a limited resource. *Journal of Experimental Psychology: General*. 2015;144(4):796-815. doi:10.1037/xge0000083
51. Inzlicht M, Gervais W, Berkman E. *Bias-Correction Techniques Alone Cannot Determine Whether Ego Depletion Is Different from Zero: Commentary on Carter, Kofler, Forster, & McCullough, 2015*. Social Science Research Network; 2015. <https://papers.ssrn.com/abstract=2659409>.
52. Park S, Beretvas S. Using total sample size weights in meta-analysis of log-odds ratios. *Journal of Experimental Education*. 2018.
53. Stanley TD, Carter EC, Doucouliagos H. *What Meta-Analyses Reveal About the Replicability of Psychological Research (Working Paper)*. Deakin Laboratory for the Meta-Analysis of Research; 2017. http://www.deakin.edu.au/___data/assets/pdf_file/0007/1198456/WhatMeta-AnalysesReveal_WP.pdf.
54. Erp S van, Verhagen J, Grasman RPPP, Wagenmakers E-J. Estimates of between-study heterogeneity for 705 meta-analyses reported in *psychological bulletin* from 1990–2013. *Journal of Open Psychology Data*. 2017;5(1). doi:10.5334/jopd.33
55. Vevea JL, Hedges LV. A general linear model for estimating effect size in the presence of publication bias. *Psychometrika*. 1995;60(3):419-435. doi:10.1007/BF02294384
56. Coburn KM, Vevea JL. Publication bias as a function of study characteristics. *Psychological Methods*. 2015;20(3):310-330. doi:10.1037/met0000046
57. Citkovicz M, Vevea JL. A parsimonious weight function for modeling publication bias. *Psychological Methods*. 2017;22(1):28-41. doi:10.1037/met0000119
58. Rufibach K. Selection models with monotone weight functions in meta analysis. *Biometrical Journal*. 2011;53(4):689-704. doi:10.1002/bimj.201000240

59. Du H, Liu F, Wang L. A bayesian “fill-in” method for correcting for publication bias in meta-analysis. *Psychological Methods*. 2017;22(4):799.
60. Becker BJ. Multivariate Meta-analysis. In: Brown SD, Tinsley HE, eds. *Handbook of Applied Multivariate Statistics and Mathematical Modeling*. San Diego, CA: Academic Press; 2000:499-525. doi:10.1016/B978-012691360-6/50018-5
61. Hedges LV, Tipton E, Johnson MC. Robust variance estimation in meta-regression with dependent effect size estimates. *Research Synthesis Methods*. 2010;1(1):39-65. doi:10.1002/jrsm.5
62. Jackson D, Riley RD, White IR. Multivariate meta-analysis: Potential and promise. *Statistics in Medicine*. 2011. doi:10.1002/sim.4172
63. Van den Noortgate W, López-López JA, Marín-Martínez F, Sánchez-Meca J. Three-level meta-analysis of dependent effect sizes. *Behavior Research Methods*. 2013;45(2):576-594. doi:10.3758/s13428-012-0261-6
64. Chen Y, Hong C, Riley RD. An alternative pseudolikelihood method for multivariate random-effects meta-analysis. *Statistics in Medicine*. 2014. doi:10.1002/sim.6350
65. Frieze M, Frankenbach J, Job V, Loschelder DD. Does Self-Control Training Improve Self-Control? A Meta-Analysis. *Perspectives on Psychological Science*. 2017;12(6):1077-1099. doi:10.1177/1745691617697076