Examining publication bias – A simulation-based evaluation of statistical tests on publication bias

3 Andreas Schneck[1]

4 [1] Department of Sociology, Ludwig-Maximilians-Universität München, Munich, Germany

5 Corresponding Author:

6 Andreas Schneck[1]

7 Konradstraße 6, Munich, 80801, Germany

8 Email address: andreas.schneck@lmu.de

9   **Abstract**

10   **Background**

11   Publication bias is a form of scientific misconduct. It threatens the validity of research results and

12   the credibility of science. Although several tests on publication bias exist, no in-depth evaluations

13   are available that examine which test performs best for different research settings.

14   **Methods**

15   In the study at hand four tests on publication bias, Egger's test (FAT), p-uniform, the test of

16   excess significance (TES), as well as the caliper test, were evaluated in a Monte Carlo simulation.

17   Two different types of publication bias and its degree (0%, 50%, 100%), were simulated. The

18   type of publication bias was defined either as *file-drawer*, meaning the repeated analysis of new

19   datasets, or *p-hacking*, meaning the inclusion of covariates in order to obtain a significant result.

20   In addition, the underlying effect (β = 0, 0.5, 1, 1.5), effect heterogeneity, the number of

21   observations in the simulated primary studies (N =100, 500), and the number of observations for

22   the publication bias tests (K = 100, 1000), were varied.

23   **Results**

24   All tests evaluated were able to identify publication bias both in the *file-drawer* and *p-hacking*

25   condition. The false positive rates were, with the exception of the 15%- and 20%-caliper test,

26   unbiased. The FAT had the largest statistical power in the *file-drawer* conditions, whereas under

27   *p-hacking* the TES was, except under effect heterogeneity, slightly better. The CTs were,

28   however, inferior to the other tests under effect homogeneity and had a decent statistical power

29   only in conditions with 1000 primary studies.

30   **Discussion**

31 The FAT is recommended as a test for publication bias in standard meta-analyses with no or only

32 small effect heterogeneity. If two-sided publication bias is suspected as well as under *p-hacking*

33 the TES is the first alternative to the FAT. The 5%-caliper test is recommended under conditions

34 of effect heterogeneity and a large number of primary studies, which may be found if publication

35 bias is examined in a discipline-wide setting when primary studies cover different research

36 problems.

Introduction

All scientific disciplines try to uncover truth by systematically examining their surrounding environment (Descartes 2006: 17). Natural scientists try to observe regularities in nature, whereas social scientists try to uncover patterns in the social behaviour of humans. The success, as well as the reputation, of science rests on the accuracy and unbiasedness of scientific results. Publication bias, the publication of only positive results confirming the researcher's hypothesis (cf. Dickersin & Min 1993: 135), threatens this validity. Under publication bias only results showing either statistical significance and/or the desired direction of the effects are published. The published literature in this case is merely a selective (and too optimistic) part of all existing scientific knowledge. Furthermore, science is in the case of publication bias also inefficient as studies that add substantial knowledge to the literature, but contain null-findings remain unpublished.

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

The issue of publication bias

The false positive rate of a test (commonly called *p*-value) is the probability of the estimator rejecting H_0 despite this being true. The *p*-value is therefore the probability that the observed estimate is at least as extreme given there is no effect as assumed by H_0 (Wasserstein & Lazar

61 2016). The larger the *p*-value the higher the risk of assuming an effect if none exists in the data.

62 *p*-values below a certain threshold are called statistically significant, whereas values above the

63 threshold are labelled as non-significant. In the empirical sciences the 5%-significance threshold

64 is mostly used (Cohen 1994; Labovitz 1972; Nuzzo 2014). The difference between 0.049 and

65 0.051 in the error probability is, however, marginal. Nevertheless, from the standpoint of the 5%-

66 significance threshold the first would be a significant effect, whereas the latter would be a non-

67 significant effect. In both of these two cases on average around 1 in 20 null-hypotheses of no

68 difference would be rejected, albeit true. If empirical researchers select their data/models until

69 they find, just by chance, significant evidence that seems worth publishing, publication bias is on

70 the rise, leading to inflated or even artificial effects.

71 Rosenthal (1979) constructs a worst case scenario in which only the 5% of false positive studies

72 that are "significant" solely by pure chance are published. In this case, misinterpreted results

73 shape the scientific discourse and finally result in (medical or political) interventions. Although

74 Rosenthal's example is extreme, a multitude of evidence for publication bias exists in various

75 disciplines and research fields (e.g. Doucouliagos & Stanley 2009; Jefferson et al. 2012). Godlee

76 (2012) therefore warns that scientific misconduct, under which publication bias is subsumed

77 (Chalmers 1990), may also physically harms patients.

78 In addition to the societal consequences, publication bias also has severe implications for the

79 evolution of knowledge. Under publication bias no rejection of theories (Popper 1968: 215), on

80 which all scientific progress relies, occur which leads to a state of "undead theory" (Ferguson &

81 Heene 2012: 559) where all existing theories are confirmed irrespective of their truth.

82 Motivation to commit publication bias
83 Because statistically significant results stress the originality of research findings (Merton 1957),

84 Both, authors and scientific journals (cf. Coursol & Wagner 1986; Epstein 1990, 2004; Mahoney

85 1977) have large incentives to maximise their significant results to survive in a publish or perish

86 research environment. Authors especially want to increase their publication chances, notably in

87 top-tier journals where low acceptance rates of 5%–10% are quite common (for the top

88 interdisciplinary journals Nature 2017; Science 2017; cf. for the political sciences Yoder &

89 Bramlett 2011: 266). Two distinct strategies to achieve significant results by means of publication

90 bias practices can be pointed out. Firstly, non-significant findings can be suppressed (cf. the

91 classical *file-drawer* effect described by Rosenthal 1979) and significant results are then searched

92 for in another dataset. Secondly, small bits in the data analysis can be changed (e.g. adding

93 covariates, optional stopping, exclusion of outliers, etc.) until a significant result is obtained –

94 this method is known as *p-hacking* (cf. "fishing" Gelman 2013; or "researchers degree of

95 freedom" Simmons et al. 2011: 1359). Whereas the *file-drawer* strategy can be utilised by authors

96 as well as by editors and reviewers, *p-hacking* can only be committed by authors/researchers.

97 Nonetheless, *p-hacking* strategies can be recommended by actors other than authors (e.g. editors,

98 reviewers, etc.). Especially *p-hacking* is almost without any costs, as data analysis tools/packages

99 become increasingly easy to apply (Paldam 2013).

100 Evidence on the prevalence of publication bias
101 So far, there are two strategies for identifying publication bias: the first traces studies through the

102 publication process, the second asks authors, reviewers, or editors about their publication

103 practices via surveys. In the first strategy, most of the analyses trace conference papers or ethics

104 committee decisions if those results get published or remain in the file-drawer. Overall previous

105 findings note, that studies with significant results have a substantially higher chance to get

106 published (cf. Callaham et al. 1998; Coursol & Wagner 1986; Dickersin 1990; Easterbrook et al.

107 1991). Ioannidis (1998) in addition finds that significant studies have, beside their higher

108 publication rate, also a substantially higher publication speed, meaning a shorter time between the

109 completion of the study and the final publication. This results suggest, that publication bias is a

110 beneficial strategy in order to maximize academic merits.

111 The second approach asks directly about the publication practices of the involved actors. In a

112 survey of psychologists that used a sensitive question technique up to 50% of the respondents

113 claimed that they exercised publication bias (John et al. 2012: 525). Franco et al. (2014) also note

114 that most non-significant findings go to the file-drawer right after the analysis and are not even

115 written up. Also, other forms of misbehaviour, like optional stopping (stopping data collection

116 when significance is reached) or erroneous rounding of p-values to reach significant results, are

117 alarmingly widespread (prevalence rate arround 22.5% John et al. 2012: 525). These results are in

118 line with the survey of Ulrich & Miller (2017: 9), who report that researchers in the field of

119 psychology prefer significant results over non-significant results, and, furthermore, attribute more

120 value to results with smaller p-values. These estimates may even be conservative because it is

121 known from the survey literature that sensitive behaviours like scientific misconduct may be

122 underreported (Kreuter et al. 2008: 848). According to the presented research results *file-drawer*

123 and *p-hacking* behaviour is therefore quite widespread.

Methods

Publication bias tests in comparison

126 So far, the presented detection strategies ask either directly for publication preferences or

127 examine the publication fate of conference papers. Both approaches have the weakness that they

128 either rely on the potentially biased answers of the actors involved or require an immense effort

129 to follow the publication process, while publication bias may have happened before the paper is

130 submitted to a conference. Statistical tests on publication bias circumvent this problem by relying

131 only on the published literature. In the paper at hand the regression-based FAT (Egger et al. 1997;

132 Stanley & Doucouliagos 2014), PU (van Aert et al. 2016; van Assen et al. 2015), an extended

133 version of p-curve (Simonsohn et al. 2014a, b; Simonsohn et al. 2015), the TES (Ioannidis &

134 Trikalinos 2007), and the CT (Gerber & Malhotra 2008a, b) were evaluated (see online appendix

135 for an in-depth discussion of the tests).[1]

136 In order to compare the different publication bias tests, four different criteria have to be

137 established: the assumptions of the test, the measurement level, the sample used, the test method,

138 and its according limitations (see Table 1).

139 The FAT tests basically the relationship between study's precision and its effect size with all

140 available effect sizes from primary studies. If larger effects are observed for studies with low

141 precision (and low N) publication bias is suspected. Nonetheless if alternative reasons may lead

142 to this result: small studies examine specific high risk populations in which treatments may be

143 more effective (Sterne et al. 2011), this effect heterogeneity may lead to the diagnosis of

144 publication bias where none exists (Schwarzer et al. 2002).[2] The FAT has furthermore the

145 disadvantage that only one-sided publication bias either in favour of a positive or negative

146 significant effect can be tested. Alinaghi & Reed (2016: 10) show that if significant results of

147 either sign are searched for, the FAT suffers from massively inflated false positive rates. In the

148 Monte Carlo simulation at hand only the FAT is used because of its better statistical power as

149 shown in prior simulations compared to the similar rank correlation test of Begg & Mazumdar

150 (1994)[3] and the trim and fill technique (Duval & Tweedie 2000)[4].

---

1 1 Because for Fail-save-N (Rosenthal 1979) only rules of thumbs, instead of a formal statistical test, exist it was not
2 included in the simulation at hand. Although it is still widely applied (Banks et al. 2012: 183; Ferguson & Brannick
3 2012: 4), this benchmark is not recommended in the *Cochrane Handbook*, a guideline for conducting meta-analyses
4 (Higgins & Green 2008: 321f.).

5 2 For a similar result see Terrin et al. (2003) for the related Trim and Fill technique (Duval & Tweedie 2000).

6 3 For simulations see: Hayashino et al. (2005); Kicinski (2014); Macaskill et al. (2001); Sterne et al. (2000).

7 4 For simulations see: Bürkner & Doebler (2014); Kicinski (2014); Moreno et al. (2009); Renkewitz & Keiner
8 (2016).

151   PU has the assumption that every left skewness in the distribution of *p*-values smaller than the

152   significance threshold (e.g. $p < 0.05$) and conditioned on the underlying observed mean effect

153   (*pp*-value) distribution is caused by publication bias. This assumption is, however, grounded

154   mainly on the fixed-effect estimate of the mean effect, which is very sensitive to effect

155   heterogeneity. PU furthermore limits its test-value only on significant estimates in the direction

156   where publication bias is suspected (van Aert et al. 2016: 727). Therefore PU is, as the FAT, only

157   able to identify one-sided publication bias.

158   The TES (Ioannidis & Trikalinos 2007; also called ic-index see Schimmack 2012) in contrast

159   relies only on a dichotomous classifier, testing if the number of expected significant results and

160   the empirically observed number of significant effects differ. Because the TES relies, as PU on

161   the fixed-effect estimate of the mean effect of all included studies it is sensible to effect

162   heterogeneity. A large controversy in the literature is not about the TES itself, but on its

163   application. Francis (2012a, b, c, d, e, 2013) used the TES to identify singular articles in order to

164   test if they suffer from publication bias. This may invalidate the assumption of independence

165   (Morey 2013: 181) as well as inflate the false positive rate in a similar manner than in primary

166   research (cp. HARKing Kerr 1998) if the TES is used in such an exploratory manner (Simonsohn

167   2013: 175). Ioannidis however responds that if the TES is applied on prespecified research

168   questions with a large and independent number of effect sizes, the TES is even a conservative test

169   on publication bias (Ioannidis 2013: 185).

170   The CT uses the most limited sample of the included tests that includes only estimates slightly

171   over and under the chosen significance level in a distribution of *z*-values. In case of publication

172   bias the assumption of a continuous distribution that results in an approximately even distribution

173   in a narrow interval (caliper) is violated by an overrepresentation of just significant results. The

174   broader the interval is set the more it may deviate from the assumed even distribution caused by

175   the true underlying effect. This restrictive sample has the downside that the exclusion of most

176   available values may drastically reduce the statistical power of the test.

177   In contrast to the FAT, the other tests are only able to test for publication bias on pre-specified

178   levels (e.g. 0.05). Because the TES and the CT focus only on dichotomous classifiers (significant

179   or not in the case of the TES, slightly over or under the threshold for the CT) also tests on two-

180   sided publication bias are possible.

181   In previous simulation studies with a low number of included studies as well as observations PU

182   was superior to the TES (van Assen et al. 2015: 303) and the FAT (Renkewitz & Keiner 2016).

183   However no evaluations exists based on a larger number of primary studies. In particular, the

184   newer publication bias tests like PU, the TES, and the CT, are in need of an evaluation under

185   different conditions. For the CT also no studies exist regarding the best caliper width to use.

186   Despite the existence of some simulation studies on publication bias tests, so far no direct

187   comparison exists that evaluates the performance of all four publication bias tests, especially

188   under effect heterogeneity.

189   Simulation setup
190   In order to examine the performance of the four publication bias tests, a Monte Carlo simulation

191   approach is used. For the simulation two different processes have to be distinguished: firstly, the

192   data generation process (DGP), and, secondly, the meta-analytical estimation method (EM). The

193   DGP provides the ground for the hypothetical data used by the simulated actors, as well as the

194   results they report, whereas the EM applies the tests on publication bias reported in the previous

195   section. The central advantage of using Monte Carlo simulations is that controlling the DGP

196   allows to identify which simulated studies suffer from publication bias and which do not. Similar

197   to the case in experiments, different conditions can be defined to ensure a controlled setting. The

198   performance of the estimators can then be examined under the different conditions.

199   Data setup of the primary studies and meta-analyses

200   The first step of the DGP defines different effect size conditions that underlie the analyses of the

201   simulated actors (see Table 2). As a first condition the underlying true effect was specified by a

202   linear relationship with $\beta = 0, 0.5, 1.0, 1.5$. Analogous to a linear regression model this means for

203   $\beta = 0.5$, that an increase of one unit of the independent variable x increases the dependent

204   variable *y* by 0.5. The specified linear relationship between the dependent variable *y* and the

205   independent variable *x* had a normally distributed regression error term of $\varepsilon = N(0,10)$, while the

206   variation of the independent variable was defined as $\sigma_x = 2$ (for a similar setup see Alinaghi &

207   Reed 2016; Paldam 2015). The regression coefficients can also be transformed in the Pearson

208   correlation coefficient yielding approximately $r = 0, 0.1, 0.2, 0.3$. This results are equivalent to

209   low or medium effect sizes in terms of Cohen (1992) and cover about 75% of the empirical

210   observed effects in psychology (Bosco et al. 2015: 436).[5] In addition to the homogenous

211   conditions with a common effect size, a heterogeneous condition was added that assumes no

212   fixed distribution of an underlying effect but a uniform mixture of all four effect sizes, as defined

213   above, plus an additional effect of $\beta = 2.0$ $(r = 0.4)$ in order to ensure enough variation.

214   As the FAT is based on study precision, which is mainly driven by the number of observations

215   (*N*) of the primary studies, *N* was computed as a second condition by an absolute normal

216   distribution with a mean of 100 (small *N*) or 500 (large *N*) and a standard deviation of 150. In

217   order to ensure an adequate statistical analysis for the primary studies, *N*s equal to or smaller than

218   30 were excluded. This procedure resulted in a right skewed distribution with a mean *N* of

219   roughly 500 for the large *N*, and 165 for the small *N* condition. The small *N* condition reflects the

220   observed number of observations in leading economics journals (mean: 152, own computations

221   from the publicly available dataset of Brodeur et al. 2016) as well as of typical trials included in

9   5 16.6% of the simulated studies were adequately powered with at least 80% power (Cohen 1988: 56). The setting
10   produced by the DGP also reflects the results of (Ioannidis et al. 2017: 245), who report that only 10% of the studies
11   in Economics are adequately powered.

222   Cochrane reviews (mean: 118, Mallett & Clarke 2002: 822). Because both studies refer mainly to

223   an experimental literature, the large *N* condition reflects the more common number of

224   observations especially in ex-post-facto designs (e.g. survey studies).

225   The heterogeneity of effects in each of the meta-analyses was measured by I^2, the share of

226   systematic variation in respect to the overall variation consisting of the systematic and random

227   variation (Higgins & Thompson 2002). In case of the small *N* 68.62% and for a large *N* 86.86%

228   of the variation was systematic in the heterogeneous effect condition. In terms of Higgins &

229   Thompson (2002: 1553), an I^2 larger than 50% has to be modelled explicitly in meta-analyses and

230   cannot be ignored.

231   In addition to the number of observations in the primary studies (*N)* the number of primary

232   studies that were included in the meta-analysis and form the basis of the publication bias tests (*K*)

233   was varied in the third condition. A setting with 100 studies was used as a lower condition,

234   whereas 1000 studies were set as an upper condition. Although on average the number of trials in

235   a meta-analysis is usually much lower than 100 studies (median 28 studies in the meta-meta-

236   analysis by Elia et al. 2016: 5). 100 studies were chosen because in this setting every publication

237   bias test evaluated is at least partially applicable. In other research areas like Economics, where

238   meta-regression models are more widely used to model effect heterogeneity, also higher numbers

239   of included trials estimates are quite common (e.g. 1474 effect estimates in Doucouliagos &

240   Stanley 2009).

241   Behavioural setup of publication bias
242   Building on this data setup stage of the DGP the behavioural setup adds publication bias to the

243   simulation in a fourth step (see Table 2). Publication bias was defined as the willingness to collect

244   new data or run additional analyses if statistical significance failed ($p \geq 0.05$) or a negative effect

245   occurred. In the simulation only one-sided publication bias was modelled because both, the FAT

246   and PU are not able to model two-sided publication bias that focuses only on significant results

247   irrespective of its sign. It is important to note, that only the intent to commit publication bias was

248   varied in the simulation setup. The actual publication bias depends on the data setup itself: how

249   large is the true effect size (β) and the number of observations (N) in the primary studies? Or, in

250   short: is there already a significant positive result which does not need a publication bias

251   treatment?

252   Five different publication bias conditions have to be distinguished. Firstly, the condition without

253   publication bias: in this ideal case all estimates (βx) are estimated by a bivariate ordinary least

254   squares (OLS) model and afterwards published. Publishing in terms of the simulation model

255   means that all estimates enter the final meta-analysis. Therefore, in the condition without

256   publication bias either 100 or 1000 regression results were estimated and enter the meta-analysis.

257   In the second and third conditions publication bias was present with a 50% probability. That

258   means that 50% of the actors were willing to run additional analyses in order to obtain significant

259   results. These conditions seem closest to the behavioural benchmark of the empirical studies

260   presented.

261   If a non-significant result was obtained, actors operating under the second condition chose to

262   collect new data in order to obtain significant results that can be published. This second condition

263   therefore modelled publication bias under the *file-drawer* scenario, because the datasets not used

264   remained unpublished. An actor tried to run analyses on the basis of up to nine additional datasets

265   and only stopped earlier if a significant result with a positive sign was obtained. If none of the 10

266   datasets yielded a significant relationship with a positive sign, the estimate which was closest to

267   the significance threshold has been published. This rule served two purposes: firstly, it seemed

268   plausible that an actor who has tried that many analyses wants to get the results published in the

269   end to compensate for the invested effort and to avoid sunk costs (Thaler 1980). Secondly, from a

270   technical point of view, this allowed to keep the number of observations in a meta-analysis K

271   constant across all simulation conditions.

272   In the third condition an actor did not try to achieve significant results by running the same

273   bivariate analysis on different samples, but rather tried to run different model specifications on

274   the same data by including control variables (z_j) to achieve statistical significance of the

275   coefficient of interest (βx). The third condition therefore modelled publication bias as a form of

276   *p-hacking*, because the existing dataset was optimised to receive a significant *p*-value. The actor

277   was able to add three different control variables to the model. The control variables were defined

278   as collider variables that are both an effect of *x* as well as *y*, which biases the effect of interest

279   (Cole et al. 2009; Greenland et al. 1999). The effect of *x* and *y* on z_j was, however, only small ($\gamma =$

280   0.5). The error term of the equation defining *z* was normally distributed $\mathsf{N}(0,10)$. With three

281   available control variables z_j an actor had seven different combinations to improve the research

282   results in order to obtain a significant effect of *x* on *y*.

283   In contrast to the second and third conditions, where 50% of the actors had the intention to

284   commit publication bias, in the fourth and fifth conditions all actors had the intention to engage in

285   publication bias practices, once again either through *file-drawer* (fourth condition) or *p-hacking*

286   behaviour (fifth condition). Part from the higher degree of intention to engage in publication bias

287   practices the settings remained the same. Although the two conditions where all actors had the

288   intent to engage in publication bias are far too pessimistic, they allow to evaluate the performance

289   of the tests in the most extreme publication bias environment. Tests that are not able to detect

290   publication bias even under such extreme conditions are of low utility to the research community.

291   The resulting design matrix had 100 different combinations resulting from 20 data setup

292   conditions multiplied by the five publication bias conditions. In order to obtain reliable estimates

293   similarly to in an experiment (Carsey & Harden 2013: 4f.), every single cell of the design matrix

294   had to be replicated multiple times.[6]

295   The aim of the simulation study at hand was to compare the performance of the four tests in

296   respect of: A) their capability to detect publication bias if present (true positive, statistical power),

297   as well as B) consistent false positive classification (α-error). Because the conditions with and

298   without publication bias are known in a simulation study, the power of the tests and the false

299   positive rate is computable (Mooney 1997: 77-79). In a first step, a dummy variable (s) was

300   constructed, with the value 1 for a significant test result below the significance threshold (5%

301   significance level ; $s = 1 \: if \: p < 0.05$). The statistical power, was than defined as the proportion

302   of significant results s in respect to all *runs* with publication bias ($\sum_{i=1}^{r} s_i / run \: if \: PB > 0$). The

303   false positive rate was computed equivalently but in conditions without publication bias (

304    $\sum_{i=1}^{r} s_i / run \: if \: PB = 0$).

305   Results

---

12   6 In order to specify the number of replications that are necessary to achieve a sufficient statistical power of at least
13   80% (Cohen 1988: 56) a power analysis was conducted for the statistical power, as well as the false positive rate
14   estimates. For the false positive rate a small deviation of 1 percentage point from the set 5%-false positive rate has to
15   be correctly identified with at least an 80% chance. To achieve this goal, every condition without publication bias
16   had to be supported with 3729 runs. As deviations in power are, though important, not as essential as the false
17   positive rate (Cohen 1988: 56) a difference of 3 percentage points is set as acceptable. In order to identify a 3
18   percentage point deviation from the target power of 80% each of the 80 conditions with existing publication bias
19   needed 1545 runs. In total, 198080 runs were necessary, resulting in nearly 109 million primary studies that in the
20   case of publication bias contained up to 10 different regression models.

306 Prevalence of publication bias

307 Because publication bias in the experimental setup was implemented as the intent to commit

308 publication bias, three variables are useful to address the actual publication bias and its impact on

309 the overall bias. Firstly, the share of actual studies per meta-analysis that suffer from publication

310 bias (if $p < 0.05$ or negative result are obtained as a first result), secondly the share of studies that

311 achieve their goal of a significant positive result by publication bias, and thirdly the impact of

312 publication bias on the p-value of a fixed-effect meta-analysis (deflation factor of the p-value).

313 Because the heterogeneous effect condition of the simulation does not allow an absolute bias

314 measure the p-value deflation factor was used for all conditions.

315 In a first step the focus is on how the opportunity structures of the simulation conditions shape

316 the committed publication bias. In the first two columns of Table 3 the actual committed

317 publication bias is shown dependent on study characteristics like the mean number of

318 observations in the primary studies (N) and the underlying effect (β), including effect

319 heterogeneity. As expected around 50% respectively 100% of the studies committed publication

320 bias in case of an underlying null-effect, because only 2.5% of the results had the right positive

321 sign and were significant just by chance. In case of an underlying effect the share of committed

322 publication bias decreased because an already significant finding made a publication bias

323 treatment unnecessary. For $\beta = 0.5$ in the 50% publication bias condition only 35%; for $\beta = 1$,

324 15% for $\beta = 1.5$ only 9%; and in the heterogeneous condition 22% of the studies employed

325 publication bias practices. The 100% publication bias condition approximately doubled the

326 prevalence rates of the 50% condition as expected.

327 Besides the necessity of publication bias to achieve significant results also the success probability

328 in respect to committed publication bias depended on the conditions of the primary studies as

329 shown by the third column of Table 3. For small studies ($N = 100$) with an underlying null-effect

330   ($\beta = 0$) the success-probability in respect to the committed publication bias was about 31.3%.

331   Publication bias got more effective if larger studies ($N = 500$) provide the primary study with

332   more statistical power. The success probability of publication bias rose dramatically around 50

333   percentage points if a specific underlying empirical effect existed. Also in case of effect

334   heterogeneity the success probability increased about 35.8 percentage points. Slight differences

335   could be observed in the effectivity of the publication bias mechanism, as *p-hacking* was with 10

336   percentage points less effective than the *file-drawer* condition to achieve significant results.

337   As publication bias deflates *p*-values and therefore biases meta-analytical effect estimates the

338   impact of the actual observed publication bias (the share of committed and successful publication

339   bias) on the meta-analytical *p*-value is presented. The fourth column of Table 3 shows that with

340   an average proportion of publication bias committed (32.6%) as well successfully implemented

341   (18.8%) in a meta-analysis with 100 studies ($K = 100$) the *p*-value of the meta-analysis more than

342   quartered. This is further aggravated if the share of committed as well as successful publication

343   bias rose by 10 percentage points. The actual impact of successful publication bias deflated the *p*-

344   values by 7.7 percentage points and was more pronounced than the deflation caused by non-

345   successfully committed publication bias (deflation by 1.7 percentage points). The deflation also

346   was less severe if *p-hacking* procedures, as implemented in the simulation at hand, were used.

347   Nonetheless, the meta-analytical *p*-value in case of *p-hacking* is still less than half the size

348   (42.2%) of the unbiased estimate.

349   False positive rate of publication bias tests
350   For the evaluated tests on publication bias consistent false positive rates are most important. In

351   the simulation none of the tests should exceed the prespecified 5% error probability in any

352   condition. The false positive rate of the test was fixed in the simulation setting to 0.05, so all false

353 positive rates should be equal to, or even smaller than, 0.05. Positive deviations from 0.05 point

354 to inflated false positive rates, which lead to more false conclusions than expected.

355 Table 4 shows the false positive rate in dependence of the number of studies included ($K = 100$,

356 1000) and effect heterogeneity measured by I^2. In the constant condition of a meta-analysis with

357 $K = 100$ and no effect heterogeneity none of the tests had larger false positive rates than the

358 expected 0.05. In particular, the TES, the 3% and 5% CTs were very conservative. A larger meta-

359 analytical sample increased the false positive rates for the TES and the CTs. The broadest 15%

360 CT missed the expected significance threshold of 5%, with 7.8% clearly. The false positive rates

361 for PU in contrast were slightly lower. Increasing effect heterogeneity resulted in more

362 conservative false positive rates for PU, the 15% CT, and to a smaller extent also for the FAT, the

363 TES and the 10% CT. The narrower 3% and 5% CTs were unaffected by effect heterogeneity.

364 The overall influence of the varied conditions on the false positive rate was small, as can be seen

365 by the small share of explained variance ($R^2 < 1.7\%$). Looking at the false positive rates by each

366 condition (Table A1 in the online appendix) only the 10%- and 15%-caliper showed increased

367 false positive rates because the underlying true effect rather than publication bias elicited an

368 overrepresentation of just significant values. Note however, that the 3% and 5% CTs showed now

369 increased false positive rates.

370 Statistical power of publication bias tests
371 The following regression model (Table 5) addresses the statistical power conditional on the type

372 of publication bias and its occurrence (committed as well as successful publication bias). Starting

373 from the baseline condition of a meta-analysis with $K = 100$, a mean share of publication bias

374 committed (32.6%), as well as successfully applied (18.8%) via a *file-drawer* procedure and no

375 effect heterogeneity, the FAT had a superior power of 56.9%, followed by the TES (51.5%) and

376 the PU (48.3%). The CTs performed worst and yielded only a power of 0.0%–38.6%.

377   The underperformance of the CTs is largely explained by the small number of studies in the meta-

378   analyses. With $K = 100$ hardly any study falls within the small caliper around the significance

379   threshold. This limitation on just significant or non-significant effects also led to missing values,

380   because without observations in the caliper no CT could be performed. The underperformance of

381   the CT changed if 1000 studies were included, which improved the estimated power substantially,

382   by 30.7–57.3 percentage points, while smaller calipers profited most. The FAT, the TES, and the

383   PU, profited moderately from an increased number of studies, by 24.4, 23.8 and 16.5 percentage

384   points, respectively. When focussing on the influence of heterogeneity in the meta-analyses the

385   PU and the TES showed a drastic drop in power, by 6.5 and 6.4 percentage points, if the

386   heterogeneity measured by I^2 rose by 10 percentage points. This decrease in power shows that

387   neither PU nor TES were able to cope with heterogeneity. In contrast, the FAT and the CTs

388   actually showed a slight increased statistical power. Varying the publication bias procedure from

389   *file-drawer* to *p-hacking,* which is less related to the standard error of the effect estimates,

390   increased the power of PU, TES, and the CTs. The CTs profited most, increasing the statistical

391   power by around 18 percentage points. The TES and PU showed a smaller increase of power, by

392   7.5 and 4.8 percentage points. The FAT, in contrast lost about 11 percentage points of its power

393   under *p-hacking* compared to the *file-drawer* condition. Although the differences in power are

394   dependent especially on the operationalization of the *p-hacking* condition in the simulation, this

395   result points on a weakness of the FAT under non *file-drawer* conditions that are less related to

396   the standard error of the estimate but are still detectable in the distribution of *z-* or *p*-values.

397   The structural difference between tests based on a continuous effect distribution (FAT, PU) and

398   tests that focus only on a dichotomous classification (TES, CTs)[7] becomes evident looking at the

399   effect of the proportion of studies that underwent a publication bias treatment in the simulation

---

21   7 Significant or not (TES) over- or under-caliper (CTs).

400   and the proportion of studies that had a successful outcome after publication bias. Increasing the

401   share of studies under publication bias lifted the power by 3.0 (FAT) and 5.1 (PU) percentage

402   points. A 10 percentage point increase in studies successfully applying publication bias increases

403   the power by 9.9 (FAT) and 10.3 percentage points (PU). The TES and the CTs, however, were

404   only able to detect successful publication bias. An increase only in studies committing

405   publication bias (whether successful or not) in contrast reduced the statistical power. Both tests

406   were therefore not able to detect all possible outcomes of publication bias. This is especially

407   problematic as non-successful publication bias may also inflate the overall estimated effect in

408   meta-analyses. All effects presented are statistically significant ($p < 0.05$).

409   In contrast to the influence of the varied conditions on the false positive rate, the influence on

410   statistical power was substantial, varying from 30.6% in the case of the FAT to 57.2% for the PU.

411   This finding underlines the fact that all publication bias tests have their strengths and weaknesses

412   in specific conditions.

413   Discussion & Conclusion

414   In the simulation at hand, the performance of four different tests (PU, FAT, TES, CTs) were

415   evaluated in a Monte Carlo simulation. Different conditions were varied: the underlying true

416   effect size, including effect heterogeneity, the number of observations in the primary studies, the

417   number of studies in the meta-analyses, the degree of publication bias and its form as either *file-*

418   *drawer* or *p-hacking*.

419   Limitations

420   In order to compare the tests in a realistic setting that is nonetheless at least from the assumptions

421   of all four tests applicable, four central limitations have to be pointed out:

422   Firstly the simulation and its according publication bias procedures rest on the assumptions that

423   all the correlation between the study's precision and its effect size is caused by publication bias.

424   In case that studies with larger effects are, for example after a pre-study power analysis (Lau et

425   al. 2006) conducted with a lower number of observation especially the FAT may yield increased

426   false positive rates (Schwarzer et al. 2002).

427   Secondly the number of observations included in the meta-analyses either set to $K = 100$ or 1000

428   is large compared to the average meta-analysis (Elia et al. 2016). The results however showed

429   that even in such large meta-analyses and especially in the more realistic condition in which 50%

430   of the actors are willing to commit publication bias, the tests hardly yielded an adequate

431   statistical power under most conditions. Increasing the number of included studies is therefore

432   important to assure an adequately powered test on publication bias.

433   Thirdly the analysis focused only on one specific form of *p-hacking* that could occur in both

434   small ($N = 100$) or large studies ($N = 500$). Especially for studies where N is small, other

435   strategies like optional stopping may also be applied. Further research on publication bias should

436   therefore focus on the different impact of other *p-hacking* practices.

437   As a fourth limitation only one-sided publication bias against insignificant or negative results was

438   simulated. By assumption especially PU limits only on the negative or positive signed studies that

439   were supposed to be affected by publication bias. Also the FAT is not able to detect two-sided

440   publication bias because the funnel in this situation may still be perfectly symmetric. The

441   suggestions for applications if two-sided publication bias is suspected are therefore limited to the

442   TES and the CTs only.

443   Conclusion
444   The following five conclusions can be derived from the results: Firstly, for homogenous research

445   settings and with publication bias favouring only effects in one direction (one-sided publication

446   bias) the FAT is recommended due to its most consistent false positive rate as well as its superior

447   statistical power. Secondly, if there are concerns whether there are any correlations between the

448   precision of the study and its effect size for other reasons than publication bias (see first

449   limitation) and if *p-hacking* is suspected, the TES should be preferred to the FAT under effect

450   homogeneity. As the 5% CT offers more relaxed assumptions it is therefore the first alternative

451   for the FAT under effect heterogeneity if a large number of studies is included in the meta-

452   analysis.

453   Despite the analysis focussed only on one-sided publication bias, also two-sided publication bias,

454   favouring significant results with either sign may also be present. As PU and the FAT are not able

455   to identify one-sided publication bias only the TES and the CTs remain for two-sided publication

456   bias. Therefore, fourthly the TES is recommended under effect homogeneity because of its larger

457   statistical power compared to the CTs. Fifthly, in the case of heterogeneous effect sizes and a

458   sufficient number of observations in the meta-analysis the 5% CT provides the best trade-off

459   between a conservative false positive rate and a decent statistical power.

460   The 5% CT is therefore best used to identify publication bias in an effect heterogeneous

461   discipline-wide setting which relies per definition on completely different underlying effects but

462   offers enough studies to compensate for the low statistical power. Because the wider 10% and

463   15% CTs yield inflated false positive rates, at least in some conditions, they are not recommended

464   to identify publication bias.

465   Identifying publication bias in substantial meta-analyses as well as focussing on publication as a

466   general problem within the scientific domain is necessary in order to establish and retain trust in

467   scientific results. Further research, however, should not only focus on the diagnosis of publication

468   bias just stating a problem that is well known (Morey 2013). Beyond the nonetheless important

469   diagnosis of the scientific "disease" a further examination of the risk factors, either on the side of

470   the involved actors or with regard to the incentive structure within the discipline (see for example

471   Auspurg & Hinz 2011) seem essential. This includes also the evaluation of possible interventions

472   (e.g. an open data policy). Research on publication bias is inevitable to maintain trust in scientific

473   results and avoid wasted research funds that also limit the efficiency of science at a whole.

474   Beside the diagnosis of publication bias and its risk factors, also estimators of the unbiased effect,

475   that are beyond the scope of this paper, like the effect estimates provided by PU and the PET (for

476   example the PET/PEESE procedure of Stanley & Doucouliagos 2014) should be evaluated

477   comparatively. This is at most important for meta-analyses with a heterogeneous effect that try to

478   uncover the underlying true effect rather than test for publication bias alone.

479   ## Acknowledgements

480   I thank Katrin Auspurg and the two reviewers, Joseph Hilgard, Bob Reed and the responsible

481   editor Robert Winkler for their valuable comments on this article.

482   ## References

483   Alinaghi N, and Reed WR. 2016. Meta-Analysis and Publication Bias: How Well Does the FAT-
484   PET-PEESE Procedure Work? https://ideas.repec.org/p/cbt/econwp/16-26.html.
485   Auspurg K, and Hinz T. 2011. What Fuels Publication Bias? Theoretical and Empirical Analyses
486   of Risk Factors Using the Caliper Test. *Journal of Economics and Statistics* 231:636-660.
487   Banks GC, Kepes S, and McDaniel MA. 2012. Publication Bias: A Call for Improved Meta-
488   Analytic Practice in the Organizational Sciences. *International Journal of Selection and*
489   *Assessment* 20:182-196. doi:10.1111/j.1468-2389.2012.00591.x.
490   Begg CB, and Mazumdar M. 1994. Operating Characteristics of a Bank Correlation Test for
491   Publication Bias. *Biometrics* 50:1088-1101. doi:10.2307/2533446.
492   Bosco FA, Aguinis H, Singh K, Field JG, and Pierce CA. 2015. Correlational Effect Size
493   Benchmarks. *Journal of Applied Psychology* 100:431-449. doi:10.1037/a0038047.
494   Brodeur A, Le M, Sangnier M, and Zylberberg Y. 2016. Star Wars: The Empirics Strike Back.
495   *American Economic Journal-Applied Economics* 8:1-32. doi:10.1257/app.20150044.
496   Bürkner P-C, and Doebler P. 2014. Testing for Publication Bias in Diagnostic Meta-Analysis: A
497   Simulation Study. *Statistics in Medicine* 33:3061-3077. doi:10.1002/sim.6177.

498 Callaham ML, Wears RL, Weber EJ, Barton C, and Young G. 1998. Positive-Outcome Bias and
499   Other Limitations in the Outcome of Research Abstracts Submitted to a Scientific
500   Meeting. *JAMA* 280:254-257.
501 Carsey TM, and Harden JJ. 2013. *Monte Carlo Simulation and Resampling Methods for Social*
502   *Science*: Sage Publications.
503 Chalmers l. 1990. Underreporting Research is Scientific Misconduct. *JAMA* 263:1405-1408.
504   doi:10.1001/jama.1990.03440100121018.
505 Cohen J. 1988. *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, N.J.: L.
506   Erlbaum Associates.
507 Cohen J. 1992. A Power Primer. *Psychological Bulletin* 112:155-159.
508 Cohen J. 1994. The Earth Is Round (p<.05). *American Psychologist* 49:997-1003.
509   doi:10.1037/0003-066x.49.12.997.
510 Cole SR, Platt RW, Schisterman EF, Chu H, Westreich D, Richardson D, and Poole C. 2009.
511   Illustrating Bias Due to Conditioning on a Collider. *International Journal of*
512   *Epidemiology* 39:417-420. doi:10.1093/ije/dyp334.
513 Coursol A, and Wagner EE. 1986. Effect of Positive Findings on Submission and Acceptance: A
514   Note of Meta-Analysis Bias. *Professional Psychology: Research and Practice* 17.
515   doi:10.1037/0735-7028.17.2.136.
516 Descartes R. 2006. *A Discourse on the Method of Correctly Conducting One's Reason and*
517   *Seeking Truth in the Sciences*. Oxford ; New York: Oxford University Press.
518 Dickersin K. 1990. The Existence of Publication Bias and Risk Factors for Its Occurrence. *JAMA*
519   263:1385-1389. doi:10.1001/jama.1990.03440100097014.
520 Dickersin K, and Min Y-I. 1993. Publication Bias - the Problem that Won't Go Away. *Annals of*
521   *the New York Academy of Sciences* 703:135-148. doi:10.1111/j.1749-
522   6632.1993.tb26343.x.
523 Doucouliagos H, and Stanley TD. 2009. Publication Selection Bias in Minimum-Wage Research?
524   A Meta-Regression Analysis. *British Journal of Industrial Relations* 47:406-428.
525   doi:10.1111/j.1467-8543.2009.00723.x.
526 Duval S, and Tweedie R. 2000. Trim and Fill: A Simple Funnel-Plot–Based Method of Testing
527   and Adjusting for Publication Bias in Meta-Analysis. *Biometrics* 56:455-463.
528   doi:10.1111/j.0006-341X.2000.00455.x.
529 Easterbrook PJ, Berlin JA, Gopalan R, and Matthews DR. 1991. Publication Bias in Clinical
530   Research. *Lancet* 337:867-872.
531 Egger M, Smith GD, Schneider M, and Minder C. 1997. Bias in Meta-Analysis Detected by a
532   Simple, Graphical Test. *British Medical Journal* 315:629-634.
533 Elia N, von Elm E, Chatagner A, Pöpping DM, and Tramèr MR. 2016. How Do Authors of
534   Systematic Reviews Deal with Research Malpractice and Misconduct in Original Studies?
535   a Cross-Sectional Analysis of Systematic Reviews and Survey of Their Authors. *BMJ*
536   *open* 6. doi:10.1136/bmjopen-2015-010442.
537 Epstein WM. 1990. Confirmational Response Bias among Social-Work Journals. *Science*
538   *Technology & Human Values* 15:9-38. doi:10.1177/016224399001500102.
539 Epstein WM. 2004. Confirmational Response Bias and the Quality of the Editorial Processes
540   among American Social Work Journals. *Research on Social Work Practice* 14:450-458.
541   doi:10.1177/1049731504265838.
542 Ferguson CJ, and Brannick MT. 2012. Publication Bias in Psychological Science: Prevalence,
543   Methods for Identifying and Controlling, and Implications for the Use of Meta-Analyses.
544   *Psychological Methods* 17:120-128. doi:10.1037/a0024445.

545 Ferguson CJ, and Heene M. 2012. A Vast Graveyard of Undead Theories: Publication Bias and
546   Psychological Science's Aversion to the Null. *Perspectives on Psychological Science*
547   7:555-561. doi:10.1177/1745691612459059.
548 Francis G. 2012a. Evidence That Publication Bias Contaminated Studies Relating Social Class
549   and Unethical Behavior. *Proceedings of the National Academy of Sciences of the United*
550   *States of America* 109:E1587-E1587. doi:10.1073/pnas.1203591109.
551 Francis G. 2012b. The Psychology of Replication and Replication in Psychology. *Perspectives on*
552   *Psychological Science* 7:585-594. doi:10.1177/1745691612459520.
553 Francis G. 2012c. Publication Bias and the Failure of Replication in Experimental Psychology.
554   *Psychonomic Bulletin & Review* 19:975-991. doi:10.3758/s13423-012-0322-y.
555 Francis G. 2012d. The Same Old New Look: Publication Bias in a Study of Wishful Seeing. *i-*
556   *Perception* 3:176-178. doi:10.1068/i0519ic.
557 Francis G. 2012e. Too Good to Be True: Publication Bias in Two Prominent Studies from
558   Experimental Psychology. *Psychonomic Bulletin & Review* 19:151-156.
559   doi:10.3758/s13423-012-0227-9.
560 Francis G. 2013. Publication Bias in 'Red, Rank, and Romance in Women Viewing Men,' by
561   Elliot et al. (2010). *Journal of Experimental Psychology: General* 142:292-296.
562   doi:10.1037/a0027923.
563 Franco A, Malhotra N, and Simonovits G. 2014. Publication Bias in the Social Sciences:
564   Unlocking the File Drawer. *Science*. doi:10.1126/science.1255484.
565 Gelman A. 2013. Too Good to Be True. *Available at*
566   *http://www.slate.com/articles/health_and_science/science/2013/07/statistics_and_psychol*
567   *ogy_multiple_comparisons_give_spurious_results.html* (accessed 6.1.2017).
568 Gerber AS, and Malhotra N. 2008a. Do Statistical Reporting Standards Affect What Is Published?
569   Publication Bias in Two Leading Political Science Journals. *Quarterly Journal of*
570   *Political Science* 3:313-326. doi:10.1561/100.00008024.
571 Gerber AS, and Malhotra N. 2008b. Publication Bias in Empirical Sociological Research.
572   *Sociological Methods & Research* 37:3-30. doi:10.1177/0049124108318973.
573 Godlee F. 2012. Research Misconduct is Widespread and Harms Patients. *BMJ* 344:e14.
574   doi:10.1136/bmj.e14.
575 Greenland S, Pearl J, and Robins JM. 1999. Causal Diagrams for Epidemiologic Research.
576   *Epidemiology* 10:37-48. doi:10.1097/00001648-199901000-00008.
577 Hayashino Y, Noguchi Y, and Fukui T. 2005. Systematic Evaluation and Comparison of
578   Statistical Tests for Publication Bias. *Journal of Epidemiology* 15:235-243.
579   doi:10.2188/jea.15.235.
580 Higgins JPT, and Green S. 2008. *Cochrane Handbook for Systematic Reviews of Interventions*:
581   Wiley Online Library.
582 Higgins JPT, and Thompson SG. 2002. Quantifying Heterogeneity in a Meta-Analysis. *Statistics*
583   *in Medicine* 21:1539-1558. doi:10.1002/sim.1186.
584 Ioannidis JP. 1998. Effect of the Statistical Significance of Results on the Time to Completion
585   and Publication of Randomized Efficacy Trials. *JAMA* 279:281-286.
586   doi:10.1001/jama.279.4.281.
587 Ioannidis JPA. 2013. Clarifications on the Application and Interpretation of the Test for Excess
588   Significance and Its Extensions. *Journal of Mathematical Psychology* 57:184-187.
589   doi:10.1016/j.jmp.2013.03.002.
590 Ioannidis JPA, Stanley TD, and Doucouliagos H. 2017. The Power of Bias in Economics
591   Research. *The Economic Journal* 127:F236-F265. doi:10.1111/ecoj.12461.
592 Ioannidis JPA, and Trikalinos TA. 2007. An Exploratory Test for an Excess of Significant
593   Findings. *Clinical Trials* 4:245-253. doi:10.1177/1740774507079441.

594   Jefferson T, Jones MA, Doshi P, Del Mar CB, Heneghan CJ, Hama R, and Thompson MJ. 2012.
595   Neuraminidase Inhibitors for Preventing and Treating Influenza in Healthy Adults.
596   *Cochrane database of systematic reviews*. doi:10.1002/14651858.CD008965.pub3.
597   John LK, Loewenstein G, and Prelec D. 2012. Measuring the Prevalence of Questionable
598   Research Practices With Incentives for Truth Telling. *Psychological Science* 23:524-532.
599   doi:10.1177/0956797611430953.
600   Kerr NL. 1998. HARKing: Hypothesizing After the Results are Known. *Personality and Social*
601   *Psychology Review* 2:196-217. doi:10.1207/s15327957pspr0203_4.
602   Kicinski M. 2014. How Does Under-Reporting of Negative and Inconclusive Results Affect the
603   False-Positive Rate in Meta-Analysis? A Simulation Study. *BMJ open* 4:1-8.
604   doi:10.1136/bmjopen-2014-004831.
605   Kreuter F, Presser S, and Tourangeau R. 2008. Social Desirability Bias in CATI, IVR, and Web
606   Surveys: The Effects of Mode and Question Sensitivity. *Public Opinion Quarterly*
607   72:847-865. doi:10.1093/poq/nfn063.
608   Labovitz S. 1972. Statistical Usage In Sociology: Sacred Cows and Ritual. *Sociological Methods*
609   *& Research* 1:13-37. doi:10.1177/004912417200100102.
610   Lau J, Ioannidis JPA, Terrin N, Schmid CH, and Olkin I. 2006. The Case of the Misleading
611   Funnel Plot. *BMJ* 333:597-600. doi:10.1136/bmj.333.7568.597.
612   Macaskill P, Walter SD, and Irwig L. 2001. A Comparison of Methods to Detect Publication Bias
613   in Meta-Analysis. *Statistics in Medicine* 20:641-654. doi:10.1002/sim.698.
614   Mahoney MJ. 1977. Publication Prejudices: An Experimental Study of Confirmatory Bias in the
615   Peer Review System. *Cognitive Therapy and Research* 1:161-175.
616   doi:10.1007/BF01173636.
617   Mallett S, and Clarke M. 2002. The Typical Cochrane Review. How Many Trials? How Many
618   Participants? *International Journal of Technology Assessment in Health Care* 18:820-823.
619   Merton RK. 1957. Priorities in Scientific Discovery: A Chapter in the Sociology of Science.
620   *American sociological review* 22:635-659. doi:10.2307/2089193.
621   Mooney CZ. 1997. *Monte Carlo Simulation*. Thousand Oakes: Sage Publications.
622   Moreno SG, Sutton AJ, Ades AE, Stanley TD, Abrams KR, Peters JL, and Cooper NJ. 2009.
623   Assessment of Regression-based Methods to Adjust for Publication Bias through a
624   Comprehensive Simulation Study. *Bmc Medical Research Methodology* 9.
625   doi:10.1186/1471-2288-9-2.
626   Morey RD. 2013. The Consistency Test Does Not – and Cannot – Deliver What Is Advertised: a
627   Comment on Francis (2013). *Journal of Mathematical Psychology* 57:180-183.
628   doi:10.1016/j.jmp.2013.03.004.
629   Nature. 2017. Getting Published in Nature: The Editorial Process. *Available at*
630   *http://www.nature.com/nature/authors/get_published/* (accessed 05.01.2017).
631   Nuzzo R. 2014. Scientific Method: Statistical Errors. *Nature* 506:150-152. doi:10.1038/506150a.
632   Paldam M. 2013. Regression Costs Fall, Mining Ratios Rise, Publication Bias Looms, and
633   Techniques Get Fancier: Reflections on Some Trends in Empirical Macroeconomics.
634   *Econ Journal Watch* 10:136-156.
635   Paldam M. 2015. Simulating an Empirical Paper by the Rational Economist. *Empirical*
636   *Economics*:1-25. doi:10.1007/s00181-015-0971-6.
637   Popper KR. 1968. *Conjectures and Refutations: The Growth of Scientific Knowledge*. New York:
638   Harper & Row.
639   Renkewitz F, and Keiner M. 2016. How to Detect Publication Biases from Published Data? A
640   Monte Carlo Simulation of Different Methods. 50 Kongress der Deutschen Gesellschaft
641   für Psychologie. Leipzig.

Rosenthal R. 1979. The File Drawer Problem and Tolerance for Null Results. *Psychological Bulletin* 86:638-641. doi:10.1037/0033-2909.86.3.638.

Schimmack U. 2012. The Ironic Effect of Significant Results on the Credibility of Multiple-Study Articles. *Psychological Methods* 17:551-566. doi:10.1037/a0029487.

Schwarzer G, Antes G, and Schumacher M. 2002. Nflation of Type I Error Rate in Two Statistical Tests for the Detection of Publication Bias in Meta-Analyses with Binary Outcomes. *Statistics in Medicine* 21:2465-2477. doi:10.1002/sim.1224.

Science. 2017. The Science Contributors FAQ. *Available at* http://www.sciencemag.org/site/feature/contribinfo/faq/#pct_faq (accessed 05.01.2017).

Simmons JP, Nelson LD, and Simonsohn U. 2011. False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science* 22:1359-1366. doi:10.1177/0956797611417632.

Simonsohn U. 2013. It Really Just Does Not Follow, Comments on Francis (2013). *Journal of Mathematical Psychology* 57:174-176. doi:10.1016/j.jmp.2013.03.006.

Simonsohn U, Nelson LD, and Simmons JP. 2014a. P-Curve and Effect Size: Correcting for Publication Bias Using Only Significant Results. *Perspectives on Psychological Science* 9:666-681. doi:10.1177/1745691614553988.

Simonsohn U, Nelson LD, and Simmons JP. 2014b. P-curve: A Key to the File Drawer. *Journal of Experimental Psychology-General* 143:534-547. doi:10.1037/a0033242.

Simonsohn U, Simmons JP, and Nelson LD. 2015. Better P-Curves: Making P-Curve Analysis More Robust To Errors, Fraud, and Ambitious P-Hacking, A Reply To Ulrich and Miller (2015). *Journal of Experimental Psychology-General* 144:1146-1152. doi:10.1037/xge0000104.

Stanley TD, and Doucouliagos H. 2014. Meta-Regression Approximations to Reduce Publication Selection Bias. *Research Synthesis Methods* 5:60-78. doi:10.1002/jrsm.1095.

Sterne JAC, Gavaghan D, and Egger M. 2000. Publication and Related Bias in Meta-Analysis: Power of Statistical Tests and Prevalence in the Literature. *Journal of Clinical Epidemiology* 53:1119-1129. doi:10.1016/S0895-4356(00)00242-0.

Sterne JAC, Sutton AJ, Ioannidis JPA, Terrin N, Jones DR, Lau J, Carpenter J, Rücker G, Harbord RM, Schmid CH, Tetzlaff J, Deeks JJ, Peters J, Macaskill P, Schwarzer G, Duval S, Altman DG, Moher D, and Higgins JPT. 2011. Recommendations for examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *BMJ* 343. doi:10.1136/bmj.d4002.

Terrin N, Schmid CH, Lau J, and Olkin I. 2003. Adjusting for Publication Bias in the Presence of Heterogeneity. *Statistics in Medicine* 22:2113-2126. doi:10.1002/sim.1461.

Thaler R. 1980. Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior & Organization* 1:39-60. doi:10.1016/0167-2681(80)90051-7.

Ulrich R, and Miller J. 2017. Some Properties of p-Curves, With an Application to Gradual Publication Bias. *Psychological Methods*. doi:10.1037/met0000125.

van Aert RCM, Wicherts JM, and van Assen MALM. 2016. Conducting Meta-Analyses Based on p Values. *Perspectives on Psychological Science* 11:713-729. doi:10.1177/1745691616650874.

van Assen MALM, van Aert RCM, and Wicherts JM. 2015. Meta-Analysis Using Effect Size Distributions of only Statistically Significant Studies. *Psychological Methods* 20:293-309. doi:10.1037/met0000025 10.1037/met0000025.supp (Supplemental).

Wasserstein RL, and Lazar NA. 2016. The ASA's Statement on p-Values: Context, Process, and Purpose. *American Statistician* 70:129-131. doi:10.1080/00031305.2016.1154108.

689   Yoder S, and Bramlett BH. 2011. What Happens at the Journal Office Stays at the Journal Office:
690   Assessing Journal Transparency and Record-Keeping Practices. *Ps-Political Science &*
691   *Politics* 44:363-373. doi:10.1017/S1049096511000217.

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

Data generating process (DGP) of Monte Carlo simulation

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

Conditions	Values	Functional form	*N* (conditions)
Data setup:			
1. True effects *β*:	$β$ = 0; 0.5; 1; 1.5; Het $ε$ = NV(0,10) $σ_x$ = 2	y = $β$x + $ε$	**5**
2. Number of observations *N*:	$μ_N$ = 100; 500	$|N(μ_N)|$ N > 30	**2**
3. Number of studies *K*:	*K* = 100; 1000		**2**
Behavioural setup:			
4. Publication bias (PB)	*PB = 0; 0.5; 1.0*	$β$ > 0 & *p* < 0.05	**1+2*2 = 5**
4.1. *File-drawer*	Draw new sample size *N*	*(max. 9 additional samples)*	
4.2. *p-hacking*	Run new analyses with same dataset	y = $β$x + $γ_j z_j$ + $ε$ z = 0.5x + 0.5y + $ε$ *(max. 3 z's = 7 combinations)*	
			5*2*2*5 = 100

Risk factors for publication and its impact on bias in the simulated data (OLS regression)

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

	Publication bias committed (50% intention)	Publication bias committed (100% intention)	Publication bias successful (in relation to committed)	Deflation of *p*-value
N = 500	-0.105***	-0.211***	0.179***	
(ref. N = 100)	(0.000)	(0.001)	(0.001)	
β = 0.5	-0.196***	-0.391***	0.471***	
(ref. β = 0)	(0.001)	(0.001)	(0.001)	
β = 1	-0.389***	-0.777***	0.513***	
	(0.001)	(0.001)	(0.001)	
β = 1.5	-0.451***	-0.899***	0.503***	
	(0.001)	(0.001)	(0.002)	
β = heterogeneous	-0.319***	-0.636***	0.358***	
	(0.001)	(0.001)	(0.001)	
p-hacking			-0.100***	0.197***
(ref. file-drawer)			(0.001)	(0.002)
Committed PB [+10ppts]				-0.018***
(ref. mean = 32.5%)				(0.001)
Successful PB [+10ppts]				-0.077***
(ref. mean = 18.8%)				(0.001)
Constant	0.541***	1.080***	0.313***	0.225***
	(0.000)	(0.001)	(0.001)	(0.002)
Observations	61,760	61,760	115,843	123,520
R²	0.939	0.958	0.648	0.168

Standard errors in parentheses; *** p<0.001, ** p<0.01, * p<0.05

Conditional false positive rates of the publication bias tests (OLS regression)

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

	PU	FAT	TES	3% CT	5% CT	10% CT	15% CT
$K = 1000$	-0.005***	0.000	0.006***	0.026***	0.023***	0.026***	0.050***
(ref. $K = 100$)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
I^2 [+10 percentage points]	-0.003***	-0.001***	-0.001***	0.000	0.000	-0.001***	-0.003***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Constant [a]	0.023***	0.049	0.010***	0.003**	0.008***	0.019***	0.028***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	73,960	74,560	74,560	62,644	66,546	69,718	70,936
R^2	0.005	0.000	0.002	0.010	0.007	0.006	0.017

[a] Test H0: constant = 0.05; Standard errors in parentheses; *** p<0.001, ** p<0.01, * p<0.05

Conditional statistical power of the publication bias tests (OLS regression)

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

	PU	FAT	TES	3% CT	5% CT	10% CT	15% CT
K = 1000	0.165***	0.244***	0.238***	0.573***	0.513***	0.382***	0.307***
(ref. K = 100)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
I^2 [+10 percentage points]	-0.065***	0.001*	-0.064***	0.006***	0.005***	0.008***	0.010***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
p-hacking	0.048***	-0.110***	0.075***	0.179***	0.187***	0.186***	0.177***
(ref. file-drawer)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Comitted PB [+10ppts]	0.051***	0.030***	-0.065***	-0.035***	-0.053***	-0.073***	-0.084***
(ref. mean = 32.6%)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Successful PB [+10ppts]	0.103***	0.099***	0.221***	0.162***	0.193***	0.224***	0.234***
(ref. mean = 18.8%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Constant [a]	0.483***	0.569***	0.515***	-0.002	0.125***	0.300***	0.386***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	123,520	123,520	123,520	107,736	111,315	115,243	117,207
R^2	0.572	0.306	0.473	0.497	0.483	0.457	0.446

[a] Test H0: constant = 0.8; Standard errors in parentheses; *** $p<0.001$, ** $p<0.01$, * $p<0.05$