

Meta-Analyzing the Multiverse: A Peek Under the Hood of Selective Reporting

Anton Olsson-Collentine¹, Robbie van Aert¹, Marjan Bakker¹, and Jelte Wicherts¹

¹Department of Methodology and Statistics, Tilburg School of Social and Behavioral Sciences, Tilburg University, the Netherlands.

Corresponding author:
Anton Olsson-Collentine¹

Email address: anton@olssoncollentine.com

ABSTRACT

Researcher degrees of freedom refer to arbitrary decisions in the execution and reporting of hypothesis-testing research that allow for many possible outcomes from a single study. Selective reporting of results (p-hacking) from this 'multiverse' of outcomes can inflate effect size estimates and false positive rates. We studied the effects of researcher degrees of freedom and selective reporting using empirical data from extensive multi-study projects in psychology (Registered Replication Reports) featuring 211 samples and 14 dependent variables. We used a counter-factual design to examine what biases could have emerged if the studies (and ensuing meta-analyses) had not been preregistered and could have been subjected to selective reporting based on significance of the outcomes in the primary studies. Our results show the substantial variability in effect sizes that researcher degrees of freedom can create in relatively standard psychological studies, and how selective reporting of outcomes can alter conclusions and introduce bias in meta-analysis. Despite the typically thousands of outcomes appearing in the multiverses of the 294 included studies, only in about 30% of studies did significant effect sizes in the hypothesized direction emerge. We also observed that the effect of a particular researcher degree of freedom was inconsistent across replication studies using the same protocol, meaning multiverse analyses often fail to replicate across samples. We recommend hypothesis-testing researchers to preregister their preferred analysis and openly report multiverse analysis. We propose a descriptive index (Underlying Multiverse Variability) that quantifies the robustness of results across alternative ways to analyze the data.

Competing interests

The authors have no financial or non-financial competing interests.

Author Note

This research was supported by a Consolidator Grant 726361 (IMPROVE) from the European Research Council (ERC, <https://erc.europa.eu>), awarded to J.M. Wicherts. The results of this study were presented as an online talk at the IOPS Winter conference (December 2020) and at the Metascience 2021 conference (September 2021), and have been made available as a preprint (<https://psyarxiv.com/43yae/>). All code and data for this project are available on the Open Science Framework (OSF) at osf.io/j8yg2/ and permanently archived at Zenodo (doi.org/10.5281/zenodo.7341292). Preregistrations of code and coding protocol is available at osf.io/h397y/. Correspondence concerning this article should be addressed to Anton Olsson-Collentine, Warandelaan 2, 5037 AB Tilburg. E-mail: anton@olssoncollentine.com

23 Researcher Degrees of Freedom (DFs, Simmons et al., 2011) refer to the many arbitrary decisions
24 that need to be made in designing, collecting, analyzing, and reporting research. In the analysis of
25 hypothesis-testing research, the focus of this article, researcher DFs involve decisions such as choosing
26 between different approaches for dealing with missing observations, excluding participants from the
27 analysis depending on different criteria, and a range of other data processing and modelling decisions
28 (for more examples, see Wicherts et al., 2016). Researcher DFs allow for many possible outcomes in
29 a single study where the reported result depends on the specific combination of decisions made. This
30 was illustrated recently by Silberzahn et al. (2018): When 29 independent teams examined the same
31 data with the same research question, the teams' estimated effect sizes (measured as odds ratios) varied
32 from 0.89 to 2.93, with 20 teams finding a statistically significant effect in the expected direction. That
33 different independent teams of researchers reached different estimates shows that there often is no clearly
34 preferable analysis in hypothesis-testing research (see also Botvinik-Nezer et al., 2020; Breznau et al.,
35 2021; Huntington-Klein et al., 2021; E.-J. Wagenmakers et al., 2022).

36 The many possible statistical results that are enabled by researcher DFs have been referred to as a
37 'multiverse' of statistical results (Steege et al., 2016), 'vibration of effects' (Patel et al., 2015), or a
38 'specification curve analysis' (Simonsohn et al., 2020). These multiverse style analyses entail (sensitivity)
39 analyses of the robustness of results to researcher DFs and offer insights into potential biases that might
40 emerge if researchers selectively report outcomes from them to present more convincing evidence in favor
41 of a hypothesized effect. The proponents of multiverse style analyses are careful to define 'reasonable' or
42 'arbitrary' decisions in light of substantive, statistical, methodological, and psychometric grounds (Del
43 Giudice & Gangestad, 2021; Patel et al., 2015; Simonsohn et al., 2020; Steege et al., 2016). In the
44 current study, we use multiverse analysis based on arbitrary choices to demonstrate the potential impact
45 of selective reporting on study-level effects and subsequent meta-analyses of resulting effect sizes. To
46 ensure arbitrariness in our researcher DFs, we consider the effect size computations to be given. That is,
47 we apply researcher DFs that we consider to not change the independent or dependent variables, and we
48 do not add covariates or change the statistical model or constructs of interest.

49 From the perspective of the broader literature the principal concern with researcher DFs is not that
50 they allow multiple statistical results to be computed, but rather that they allow for selective reporting of
51 possibly desirable outcomes. Throughout this article we use 'selective reporting' to refer to cases where
52 multiple statistical results are examined in a study, but some go unreported (Page et al., 2020). We do not
53 include in this definition the special case where no study results are reported and hence do not focus on
54 publication bias of entire studies. Selective reporting is often focused on the significance of outcomes
55 and can be intentional ('*p*-hacking') or happen unintentionally due to hindsight and confirmation biases

56 (Nickerson, 1998; Roese & Vohs, 2012). Selective reporting from the multiverse of statistical results
57 is problematic as it can allow researchers to present statistical evidence even for incredible phenomena
58 (Simmons et al., 2011). Numerous formal approaches and simulation studies have been used to show that
59 selective reporting leads to an overrepresentation of false positive findings (Ioannidis, 2005) and inflated
60 effect size estimates (Ioannidis, 2008) in the literature.

61 Unfortunately, selective reporting appears common amongst researchers. In psychology, about 50-60%
62 of researchers admit to not reporting all dependent measures in a study (Agnoli et al., 2017; John et al.,
63 2012), and in a study registry comparison 70% of studies did not report all outcome variables (Franco et
64 al., 2016). Moreover, there is an extensive literature on selective reporting in the fields of biomedicine,
65 with evidence from, for example, neurology (Fusar-Poli et al., 2014), hematology (Wayant et al., 2017),
66 pediatrics (Rosati et al., 2016), orthopedics (Rongen & Hannink, 2016), obesity (Rankin et al., 2017),
67 and cancer research (Kyzas et al., 2005). A recent study examining results in 67 trials published between
68 October and November 2015 in 5 top journals from general medicine found that 42% of pre-specified
69 outcomes went unreported (Goldacre et al., 2019). Further evidence from the fields of education (Pigott
70 et al., 2013) and studies on partner violence (Madden et al., 2019) suggests the problem of selective
71 reporting is widespread indeed.

72 The biases created by selective reporting in primary studies are inherited by meta-analyses that seek
73 to quantitatively review effects or associations across many studies. Each of the studies included in a
74 meta-analysis have their own multiverse. Since the results used for meta-analysis are subsets from these
75 multiverses, meta-analytic result(s) also represent a subset from the multiverse of possible meta-analyses.
76 To avoid that this subset is biased, meta-analytic reporting guidelines such as PRISMA (Moher et al.,
77 2009) and MARS (Appelbaum et al., 2018) recommend meta-analysts to evaluate primary studies for
78 selective reporting. We do not consider arbitrary choices made in the context of meta-analyses themselves
79 (i.e., multiverse meta-analysis: Palpacuer et al., 2019; Voracek et al., 2019), but rather vary the analyses
80 in the primary studies while keeping the meta-analytic inclusion criteria and analysis constant (i.e., we
81 meta-analyze multiverses) to study the biasing effects of selective reporting based on researcher DFs in
82 primary studies on meta-analytic outcomes.

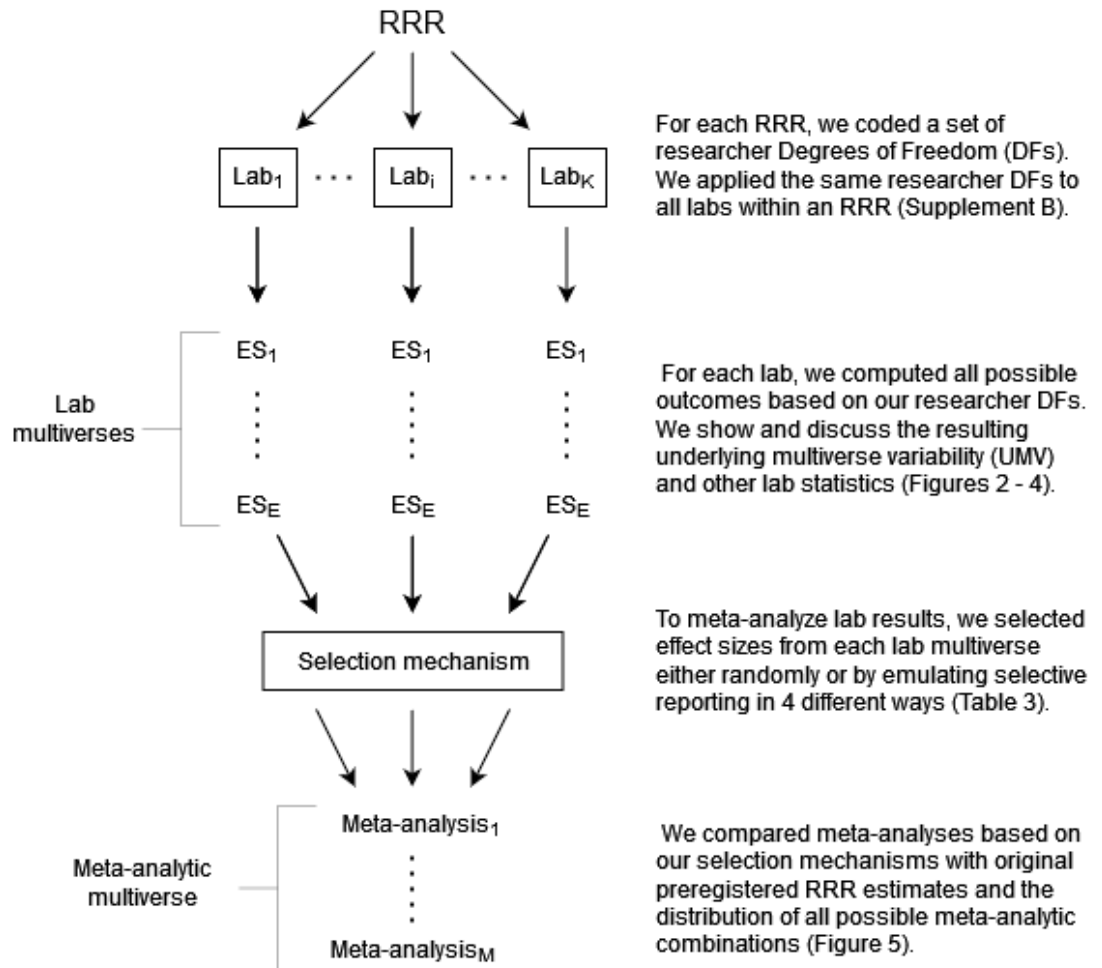
83 Such biasing effects have been studied in simulated data for meta-analysis (e.g., Botella et al., 2021;
84 Carter et al., 2019; Friese & Frankenbach, 2020), but simulated data from known distributions may not
85 be representative of actual psychological data that feature unknown (distributional) complexities. Also,
86 the effects of researcher DFs have been studied in observed data of individual studies (Botvinik-Nezer et
87 al., 2020; Breznau et al., 2021; Huntington-Klein et al., 2021; e.g., Silberzahn et al., 2018), but not in
88 meta-analytic context to inform how they might affect cumulative knowledge. We combine these streams

89 of research and study the effects of researcher DFs and selective reporting in observed meta-analytic
90 data, taking advantage of the unique opportunity offered by the open data of ten recent multi-lab direct
91 replication projects in psychology (Registered Replication Reports) that featured a total of 211 samples
92 studying 14 different outcome variables.

93 Registered Replication Reports (RRRs) each consist of a set of studies (labs) that collected data
94 on an effect in psychology using the same pre-specified research design, decision plan, and materials,
95 collectively known as a ‘preregistration.’ Each RRR can be seen as making up one (or more) meta-analysis
96 of direct (also called ‘exact’) replications, where the only difference between included studies is where
97 they collected their data. Even though the preregistrations used in the actual RRRs limited the effect of
98 researcher DFs in the original analyses, the open data from these extensive studies enable us to use a
99 counter-factual design to see what biases could have emerged if the studies (and ensuing meta-analyses)
100 had not been preregistered and could have been subjected to selective reporting based on significance
101 of the outcomes in the primary studies. In doing so, we demonstrate the variability in results that may
102 arise in meta-analytic data in the absence of preregistration, the limitations of multiverse analysis when
103 applied to a single study, and illustrate the entire process of selective reporting, from the researcher DFs
104 in primary studies that enable the practice to the consequences for meta-analysis.

105 **METHODS**

106 Figure 1 summarizes the design of this study. We identify decision points in each RRR where reasonable
107 alternative decisions could have been made (absent any preregistration) and compute all resulting outcomes
108 (create a multiverse) for each included lab. We then combine effect sizes from the lab multiverses in
109 meta-analysis within each RRR. This design allows us to explore the effects of researcher DFs on research
110 output by 1) examining the underlying multiverse variability in effect size estimates at the primary study
111 level, 2) examining the resulting multiverse variability at the meta-analytic level, and 3) examining
112 different mechanisms for selecting effect sizes from primary study multiverses for inclusion in the meta-
113 analysis. We refer to the variability due to researcher DFs as the Underlying Multiverse Variability (UMV,
114 statistically defined in ‘the multiverses’ section).



115

116 *Figure 1.* Summary of the study design. For each lab in a Registered Replication Report (RRR)
 117 multiverses were computed, analyzed and used for meta-analysis. ES = Effect size. Each RRR consists
 118 of K labs, lab $i = 1, 2, 3, \dots, K$. Each lab has E effect size estimates in its multiverse. There are M
 119 possible combinations of the E effect sizes across labs, resulting in a meta-analytic multiverse of size
 120 M . We approximate the meta-analytic multiverse by randomly sampling 10^5 meta-analyses from the
 121 meta-analytic multiverse. For details on how we selected researcher DFs, see Methods under the header
 122 ‘Selection and justification of researcher DFs.’ This figure was created using the website draw.io.

123 **Transparency and openness**

124 All our code and data for this project are available on the Open Science Framework (OSF) at osf.io/j8yg2/,
 125 and permanently archived at Zenodo (doi.org/10.5281/zenodo.7341292). We refer directly to relevant
 126 files on the OSF using brackets and links in the sections below. We registered data cleaning code and the
 127 researcher DFs available in each RRR before proceeding to analysis (osf.io/h397y/). We only made minor
 128 code corrections and clarifications of researcher DFs after registration, fully detailed in Supplement A
 129 (osf.io/xem2y/). We handled all data in R version 4.0.2 (R Core Team, 2020), and cite used packages in

130 the reference list.

131 **Data Collection**

132 We included all ten RRRs available at the time of data collection (i.e., published up until 2019-05-15)
133 available in the journals ‘Perspectives on Psychological Science’ and ‘Advances in Methods and Practices
134 in Psychological Science’ (see also Olsson-Collentine et al., 2020). Three RRRs (RRR3, RRR5, RRR9)
135 had multiple primary outcome variables (as explicitly identified in the accompanying publications). In
136 total, we included 10 projects containing 14 primary outcome variables that could be meta-analyzed,
137 consisting of 211 unique lab samples and 34,357 participants (Table 1). These values correspond to the
138 sum of the labs and participants of RRR1 - RRR9, as RRR9 and RRR10 used the same sample.

139 We use the RRRs because the meta-analyses they offer allow us to consider the effects of (selective
140 reporting from) multiverse analyses in relation to a benchmark based on meticulously collected data
141 from multiple labs (and different samples) using the same protocol. We selected the RRRs for our study
142 over other multilab replication initiatives for two reasons. First, we wished to examine researcher DFs
143 within a meta-analytic structure, which is what the RRRs nicely provide. The RRRs have the additional
144 advantage that most of them report average results not significantly different from zero, allowing us to
145 examine the bias from selective reporting under the most problematic circumstances (i.e., when there is
146 no genuine effect) and the percentage of significant outcomes appearing across multiverses (Type I error
147 rate). Second, we wished to allow for researcher DFs to depend on study design. The Many Labs series of
148 replication projects (which we have worked with previously in Olsson-Collentine et al., 2020) consists of
149 many effects studied at the same time in the same samples, meaning (almost) all researcher DFs will be
150 identical across all studied effects. Hence the RRRs allow us to delve deep into the generalizability of the
151 multiverse variability across labs and effects.

152

Table 1.

153

Preregistered Multi-Lab Replication Projects

RP	Paper	Countries	Labs	Effects	N	Sample and Settings	Description of Effects
RRR1	Alogna et al. (2014)	10	31	1	4832	31/32 samples were undergraduate students aged 18-25, 1 general population which was also the only online sample.	Verbal overshadowing 1; Independent two-group experiment. Participants either described a robber after watching a video or listed countries/capitals and after a filler task attempted to identify the robber in a lineup.
RRR2	Alogna et al. (2014)	8	26	1	2932	22/23 samples were undergraduate students aged 18-25, 1 general population which was also the only online sample.	Verbal overshadowing 2; Different from 1 only in that the filler task took place before the descriptive task instead of after.
RRR3	Eerland et al. (2016)	2	10	3	1210	11/12 samples were undergraduate students mostly aged 18-25, one of which was online. 1 sample was a broader online sample.	Grammar's effect on interpretation; Independent two-group vignette experiment with three outcome variables. Participants read about actions either described in imperfect or perfect tense and then rated protagonist's intentions (intentionality/intention attribution/detailed processing).
RRR4	Hagger et al. (2016)	10	24	1	3127	All samples consisted of in-lab undergraduate students	Ego depletion; Independent two-group experiment. Participants either assigned to a cognitively demanding or a neutral task, and performance was then measured in a subsequent cognitive task.
RRR5	Cheung et al. (2016)	5	16	2	2279	All samples consisted of in-lab undergraduate students aged 18-25	Commitment on neglect/exit; Independent two-group experiment with two outcome variables. Participants either primed to think about commitment to or independence from partner.
RRR6	Wagenmakers et al. (2016)	8	17	1	2542	All but one sample explicitly consisted of students and all took place in-lab. The last sample was recruited at university grounds.	Facial feedback hypothesis; Independent two-group experiment. Participants either induced to 'smile' or 'pout' by holding a pen in their mouth differently and simultaneously rated funniness of cartoons.
RRR7	Bouwmeester et al. (2017)	12	21	1	3669	All samples consisted of in-lab undergraduate students aged 18-34.	Intuitive cooperation; Independent two-group experiment. Economic game with money contribution to a common pool either under time pressure or time delay.
RRR8	O'Donnell et al. (2017)	13	40	1	7041	All samples consisted of in-lab undergraduate students aged 18-25	Professor priming; Independent two-group experiment. Participants primed with either a 'professor' or 'hooligan' stimuli. Outcome was percentage correct trivia answers.
RRR9	McCarthy et al. (2018)	13	26	2	6720	All samples consisted of in-lab students aged 18-25	Hostility priming; Independent two-group experiment with two outcome variables. Participants descrambled sentences, either 20% or 80% were hostile, then rated an individual and a list of ambiguous behaviors on perceived hostility.
RRR10	Verschuere et al. (2018)	12	25	1	3245	All samples consisted of in-lab students aged 18-25	Moral reminder; Independent two-group experiment. Participants either recalled the Ten Commandments or books they had read. Outcome was degree of cheating when reporting results.

Note:

All RRRs published up until 2019-05-15 in the journals 'Perspectives on Psychological Science' and 'Advances in Methods and Practices in Psychological Science'. RP = Replication Project, Countries = number of lab country locations, Effects = number of primary effects studied, N = participants before exclusions, RRR = Registered Replication Report. Table adapted with permission from Olsson-Collentine et al. (2020). Code to reproduce table: osf.io/jehpy/

154

We downloaded individual level data for all RRRs in Table 1. Summary data of all RRRs were available on the OSF. When the raw lab data were not publicly available via the OSF we contacted authors by email to request them. Only for one lab in RRR1 and 2 and two labs in RRR3 were we unable to acquire their individual level raw data.

158

For each RRR, we standardized data formatting across labs, fixed minor mistakes (e.g., mislabeled

159 columns in RRR8), and prepared the datasets for multiverse analysis (osf.io/cf86y/). We prepared the data
160 in the same way as the original RRRs to the extent possible. However, we largely had to write our own
161 code, because the alternative decisions needed to create our multiverses (e.g., exclusion criteria) could
162 not be taken in the code by the original authors. In preparing the datasets for analysis, we only excluded
163 participants due to reported experimenter error or when participants were reported to not have followed
164 instructions or completed the experiment. Note that exclusions based on ‘not following instructions’ are
165 usually *ad hoc*, and hence are distinct from formalized exclusions based on attention checks.

166 **The Multiverses**

167 Although multiverse type analyses have been suggested by multiple authors under somewhat different
168 names (Patel et al., 2015; Simonsohn et al., 2020; Steegen et al., 2016) all multiverse analyses consist of
169 identifying points in the research process where multiple reasonable decisions could have been made,
170 identifying what these decisions might be, and examining the impact of these decisions on the study results.
171 A core assumption of multiverse type analysis is that the alternative decisions are all (approximately)
172 equally reasonable (Simonsohn et al., 2020; Steegen et al., 2016; see also Del Giudice & Gangestad,
173 2021).

174 It is important that researcher DFs are defined such that these choices are indeed ‘reasonable’ or
175 ‘arbitrary’ on a priori substantive and methodological grounds (Patel et al., 2015; Simonsohn et al.,
176 2020; Steegen et al., 2016). Del Giudice & Gangestad (2021) rightly pointed out that many decisions
177 implemented in published multiverse analyses were not truly arbitrary (are ‘non-equivalent’) because they
178 can a priori be expected to result in different a) measurement reliability/validity, b) studied psychological
179 effects (e.g., when including a covariate that changes the prediction) or c) power/precision. Only decisions
180 for which this does not hold (i.e., we are either confident they are equivalent or uncertain) should according
181 to Del Giudice & Gangestad (2021) be included in a multiverse analysis.

182 We agree with Del Giudice & Gangestad (2021) that such substantive and methodological issues
183 should be considered when performing multiverse analysis as a sensitivity analysis. Hence, we carefully
184 selected our researcher DFs to reflect choices for which we saw no prior substantive or methodological
185 grounds to expect them to affect the true effect sizes tapped by the different labs. However, we could
186 imagine that others might object to some of those choices and hence we offer a range of supplementary
187 results to assess how alternative choices in designing the multiverse affect our results. We also make our
188 data and code available for re-analysis using alternative multiverse setups.

189 ***Selection and justification of researcher DFs***

190 We selected our researcher DFs to correspond to normative researcher behavior that is at risk of selective
191 reporting in the fields of the RRRs that make up our data. These RRRs belong to the fields of social

192 and cognitive psychology (Olsson-Collentine et al., 2020). Although there are many researcher DFs
193 before analyzing the data (Wicherts et al., 2016), due to using already collected data we were only able
194 to vary post data collection decisions. Moreover, because our focus was on researcher DFs in primary
195 studies and their consequences for downstream meta-analysis, we only varied decisions in data processing
196 (the data multiverse, Steegen et al., 2016) and not the statistical models used in data analysis (the model
197 multiverse). Consequently, several chosen researcher DFs concerned using different exclusion criteria
198 (which we prepend with ‘E’), although we also varied how the composite score was computed from
199 multiple indicators (researcher DFs prepended with ‘S’).

200 When creating our list of researcher DFs, we proceeded in two steps: We 1) set up a list of ‘common’
201 researcher DFs, and then 2) set up a list of researcher DFs unique to each RRR. These were then combined
202 to create our final list of researcher DFs for each RRR, which we registered before analyzing any data
203 (Supplement B). Because all labs in an RRR used the same design, it was only necessary to identify
204 decisions and create associated options once for each RRR and not for each lab/study separately. Our
205 coded researcher DFs each consisted of a decision that needs to be made and several associated potential
206 options for that decision. When defining the researcher DFs, we explored the data in the sense of
207 examining whether potential DFs could be applied (i.e., whether the variables existed and how they were
208 defined) but did not examine what effect applying them would have.

209 We created our list of common researcher DFs based on recommendations in statistical textbooks,
210 common decisions by applied researchers as reported in research literature, data analytic decisions made
211 by the included RRRs, and our own experience of decisions encountered in the literature. Table 2 provides
212 an overview of all common researcher DFs. We considered option a) across decisions to be the default
213 option, corresponding to no scale adjustments or participant exclusions (although for some researcher
214 DFs, an active decision must be made; S2, E1 Table 2). In Supplement C we detail how we selected
215 each researcher DF and its options. We acknowledge that many additional multiverses could be run in
216 these and other studies, but we consider our setup typical of researcher DFs that could be used in practice
217 across a range of psychological studies, and as such useful to study the influence of selective reporting.
218 Researchers instead interested in using multiverse analysis as a sensitivity analysis for a particular effect
219 should carefully consider the advice of Del Giudice & Gangestad (2021) on equivalent pathways before
220 applying any of the researcher DFs in Table 2.

221
222

Table 2.
List of common researcher degrees of freedom applied to Registered Replication Reports

Decision	Options	Short Explanation
S1. Post hoc scale length	a) No adjustment b) Drop the item with the lowest item-rest correlation c) Drop the two items with the lowest item-rest correlations	It is unclear how common it is to post hoc drop items 'that don't work' from a scale, but dropping more than a few seems unlikely. In the research we are looking at (experimental) there are rarely long scales. Excluding 1-item scales, the average scale length in a large sample of psychological research in 2014 was 6.87 (SD = 7.18) (Flake et al., 2017)
S2. Composite score	a) Unweighted average score b) Sum score c) PCA score: Varimax rotation, force two components and pick the first, requires at least 3 items.	For Likert-type scales with multiple items. Other DVs e.g., reaction time variability (RRR3), dichotomous correct/incorrect (RRR1&2), continuous measures (RRR7), single item DVs (RRR10) may need more unique choice options. We chose Varimax rotation to maximize variance between outcomes.
E1. Missingness DV	a) Any missing items -> list-wise deletion b) If ≤25% items missing then pair-wise deletion of missing items. Otherwise list-wise.	List-wise deletion appears to be by far the most common approach to missing data. In van Ginkel et al.'s (2010) review of personality psychology 97% used list-wise deletion for missing data and several reviews in medicine have also found it to be an extremely common method (Eekhout et al., 2012; Rombach et al., 2016; Burton & Altman, 2004). Nonetheless, we see that for example RRR6 used pair-wise deletion (option b) which may seem reasonable to some researchers, in particular with a longer scale.
E2. Missingness E3-E4 variables	a) No exclusion b) Missing on any variable -> list-wise deletion	RRRs that excluded data based on a E3-E4 variable (e.g., age) did list-wise deletion when data was missing. For other non-DV variables we make no exclusions based on missingness, unless this was explicitly done by the project (e.g., 'task completion' RRR10).
E3. Age	a) No exclusion b) Not 18-24 c) Not 18-23 d) Not 18-22 e) Not 18-21	Used by 9/10 RRRs for exclusions. Across 25 cohorts of Dutch bachelor psychology students 96.7% of students were below 25, 92.7% below 24, 86.6% below 23 and 77.8% below 22 (Wicherts et al, 2012). The oldest students in this dataset were 25. We choose a set of age ranges based on these data that we believed might go unremarked if used as exclusion criteria in psychological literature with an accompanying motivation such as 'we only included young adults'.
E4. Language / Student / Ethnicity	a) No exclusion b) Exclude participants not belonging to the dominant category	Used by 3/10 RRRs. Demographic variables which are sometimes used for exclusions. Language includes variables such as 'native speaker' which may have a yes/no response. Ethnicity includes similar variables such as 'country of birth' or 'race'. If multiple of these demographic variables are available they are treated as separate exclusion criteria.
E5. Attention check	a) No exclusion b) Exclude if failed >50% of attention check items (i.e., with two items, must fail both, e.g., RRR7) c) Exclude if failed any attention check item	Attention checks are common in psychology, as evidenced by the more than 1500 citations of Oppenheimer et al. (2009) who introduced 'instructional manipulation checks'. Curran (2016) suggests 'conservative' exclusions based on 50% failed attention checks when multiple items are used. This category does not include manipulation checks which vary more in format.
E6. Univariate outliers	a) No exclusion b) DV score > 2 SD from mean c) > 3 SD from mean d) > 1.5 times the interquartile range	Used by 1/10 RRRs. Commonly recommended cutoffs (Bakker & Wicherts, 2014). Test for outliers across groups.
E7. Multivariate outliers	a) No exclusion b) Mahalanobi's distance with $p < .001$	If the outcome variable is a correlation. Tabachnik, Fidell and Ullman (2007) recommend using Mahalanobi's distance with a cutoff of $p < .001$ for detecting multivariate outliers. Outliers tested within groups as recommended by Tabachnik et al.

Note:

S. = Degree of Freedom (DF) affecting Scale composition, E. = Exclusion DF, DV = Dependent Variable, SD = Standard Deviation, PCA = Principal Component Analysis, RRR = Registered Replication Report. Code to reproduce table: osf.io/jehpy/.

223
224
225
226
227
228

In addition to the list of common researcher DFs, which we applied to all RRRs, each RRR had several unique researcher DFs. These arise from the uniqueness of each research topic and design and consisted of different exclusion criteria. We coded between 2 (RRR3) and 10 (RRR7) unique researcher DFs for each RRR, each decision with 2-6 associated options. Due to the large number of unique researcher DFs we do not describe them all in detail here, but provide only a broad overview and refer interested readers to Supplement B.

229
230

We can separate between two types of 'unique' researcher DFs: either 1) the RRRs excluded participants based on some variable that was not defined in our list of common decisions or 2) an RRR measured

231 variables (not in our list of common decisions) that they could have used for exclusions. As an example
232 of the first case, in RRR4 (ego-depletion) participants with less than 80% correct on the main task were
233 excluded. However, 80% is a largely arbitrary number, and someone might also consider values such
234 as 75%, 85%, 90%, or many others, in addition to no exclusions. In cases like these, when there are an
235 infinite number of possible values to choose from, we have elected only a maximum of six possible values
236 that we believe an applied researcher would reasonably pick.

237 As an example of where an RRR measured variables they could have used for exclusions: RRR5
238 (commitment to romantic partner), amongst other things, asked participants whether they lived within 60
239 miles of their partner (yes/no) but did not use this variable in their analysis. However, another researcher
240 might have found it relevant to only consider participants (not) living close to each other and used this
241 variable for exclusions. Collecting data on a variable with no clear purpose thus adds researcher DFs and
242 increases the risk of selective reporting, which we in this case used to create our multiverses.

243 ***Applying researcher DFs to the RRRs***

244 After registering the coding protocol for ‘common’ researcher DFs, we coded the applicability of each
245 common researcher DFs to each RRR, which differed depending on, for example, how the outcome
246 variable was measured (binary vs. continuous, one item vs. a scale) and how projects coded their data.
247 Because some labs within RRRs pre-screened their participants for the original RRR exclusion criteria, it
248 was not always possible to apply all exclusion criteria to all labs in an RRR. Nonetheless, we still included
249 such labs, prioritizing the inclusion of more labs over the possibility of less multiverse variation. The
250 coded common and unique researcher DFs for all RRRs are available in Supplement B (osf.io/wj38n/).

251 We computed resulting effect sizes from all possible combinations of decisions for each lab in an
252 RRR (osf.io/zhdxr/). Incompatible decision combinations were not applied. For example, if we wished
253 to drop two items from a scale (Table 2; S1c) but required at least three items in the scale for Principal
254 Component Analysis (Table 2; S2c), this decision combination was inapplicable to scales with fewer
255 than 5 items. We standardized mean differences (Cohen’s *d*, p. 226, Borenstein, 2009) and computed
256 log odds ratios for RRR1 and RRR2. Effect sizes were originally analyzed unstandardized in all RRRs
257 except for RRR4, and if certain researcher DF lower the within sample variance, as is highly likely, then
258 standardized effect sizes will appear larger. However, because certain of our researcher DFs change the
259 dependent variable, and we wanted to draw conclusions across RRRs, it was necessary to standardize
260 effect sizes. As most meta-analyses use standardized effect sizes and we are interested here in the biasing
261 effects of selective reporting on typical unregistered meta-analyses, we do not consider standardization of
262 effect sizes a major concern for our analysis.

263 To prevent including lab multiverses with an unrealistically small number of participants, we only

264 included lab multiverses with at least 24 participants per experimental group, the median sample size in
265 psychology (Bakker et al., 2012), in our primary analyses. Three labs in RRR2 (L09, L17, L26) and one
266 lab in RRR8 (L24) had smaller sample sizes than required in all conditions and were excluded from these
267 analyses. We present the results of our analyses also without this sample size restriction in Supplement D.

268 **Analysis**

269 A consequence of assuming that the alternative decisions in the multiverse are equally reasonable is that
270 under the null hypothesis that no researcher DF has a systematic effect (i.e., is an actual moderator of
271 the effect) we can consider the distribution of effect sizes in the multiverse as random variability around
272 a true score. We refer to the variability underlying a given set of researcher DFs as the Underlying
273 Multiverse Variability (UMV) and define it as the standard deviation (SD) in effect size estimates that are
274 in the multiverse of the same study. A different set of researcher DFs will reveal different UMV. Other
275 researchers have focused on the distribution of *p*-values (Simonsohn et al., 2020; Steegen et al., 2016) or
276 on the range of effect sizes in the multiverse (e.g., Patel et al., 2015), but we consider it more useful to
277 treat multiverse variability in terms of the standard deviation of effect size estimates, in line with how
278 sampling error is defined. The UMV should be seen as a descriptive tool that highlights some degree of
279 variability that might have a relation with bias due to selective reporting over and beyond sampling error,
280 rather than a well-defined statistic.

281 To demonstrate the effects of researcher DFs on research output, we 1) examined the variance in
282 effect size estimates at the lab level (lab multiverses) and 2) compared meta-analytic average effect size
283 estimates based on how lab outcomes were selected from their multiverses. To examine the variance in
284 effect size estimates at the lab level, we created funnel plots, computed UMV, and standard deviations
285 in effect size resulting from variation across the options within a single researcher DF. For the funnel
286 plots, we plotted all effect size estimates at the lab level using either the standard error (for log odds
287 ratios; RRR1/RRR2) or sample size as the y-axis (for standardized mean difference effect sizes; RRR3 –
288 RRR10). We used total sample size (N) on the y-axis for all standardized mean difference (SMD) effect
289 types since most of our coded researcher DFs affected sample size.

290 To examine how large the effects of applying a single researcher DF can be and the relative impact of
291 our different researcher DFs, we computed the standard deviation in a lab's estimated effect size across
292 the options associated with each decision. For each researcher DF, we computed the standard deviation in
293 effect size when all other researcher DFs were set to their default value (corresponding to option 'a' for
294 each researcher DF, see Table 2 and Supplement B). In addition to examining standard deviations for labs
295 within RRRs, we also disaggregated these lab estimates across RRRs and then aggregated them across
296 common and unique researcher DF categories. In doing so, we treated all unique researcher DFs as one

297 category.

298 Due to computational limitations, and because it is often the case that some researcher DF must be
299 applied before another (e.g., outliers cannot be removed before the composite score has been computed),
300 we only applied the researcher DFs in a single fixed order. That is, if we have three researcher DFs (1, 2,
301 3) then we always applied them in the order 1, 2, 3 regardless of chosen option, rather than also varying
302 the order (e.g., 2, 1, 3). This fixed order may affect results when removing items with the lowest item-rest
303 correlation from a scale, or excluding participants based on outlier criteria, although we see no reason to
304 expect a systematic interaction between these two and any other researcher DFs. The fixed order also
305 makes it impossible to compute the impact of a single researcher DF across all possible researcher DF
306 combinations, although it remains possible to compute its impact when not applying any other researcher
307 DFs (see previous paragraph).

308 When comparing meta-analytic average estimates, we compared a) the original (preregistered) RRR
309 estimates, with b) an estimate of the distribution of all possible meta-analytic combinations, c) randomly
310 selected lab effect sizes, and d) lab effect sizes selected by one of four biased selection mechanisms
311 (see below). We ran all meta-analyses as random-effects models with the restricted maximum likelihood
312 estimator for estimating the between-study variance using the R-package ‘metafor’ (Viechtbauer, 2010).

313 The huge number of possible effect size combinations across labs for each RRR, the smallest consisting
314 of 697×10^{33} possible meta-analyses, made it impossible to compute the full distributions of possible
315 meta-analytic outcomes. Instead, we drew large random samples to approximate the distributions. For
316 each RRR (or outcome variable when an RRR contained multiple primary outcomes), we proceeded as
317 follows: We drew one random effect size from all possible effect sizes from lab 1, one random effect
318 size from all possible effect sizes from lab 2, one from lab 3, one from lab 4, and so on until we had
319 drawn one effect size from all labs in an RRR. The drawn effect sizes across labs were then combined
320 using a meta-analysis. We repeated this procedure, sampling with replacement from each lab’s multiverse
321 of effect sizes, until we had sampled 100,000 effect sizes from each lab, and consequently computed
322 100,000 meta-analyses for each RRR. These samples of meta-analyses constituted our approximation of
323 the distribution of possible meta-analyses for each RRR (or outcome variable when an RRR contained
324 multiple primary outcomes). The means of these distributions (and the means of the estimated lower/upper
325 95% Confidence Intervals; CIs) constituted our random sample of estimates.

326 When selectively reporting results, researchers may exhibit different behavior. We included four
327 types of biased selection mechanisms (Table 3: ‘Most significant,’ ‘Below alpha,’ ‘Random significant,’
328 ‘Bounded significant’) with different motivations. All selection mechanisms focused on statistical signif-
329 icance, and we used a two-tailed test with $\alpha = .05$ for hypothesis testing. First, we selected the effect

330 size with the lowest p -value in each lab. This allowed us to examine the most extreme selection of results
331 possible due to p -hacking ('most significant'). We included this scenario as a worst-case scenario. Second,
332 selective reporting may sometimes result in a 'bump' just below $p = .05$ when aggregating p -values across
333 selectively reported studies ('below alpha'). This is most likely in the case of incremental p -hacking
334 approaches such as optional stopping (e.g., Hartgerink, 2017). To compare what a meta-analysis of such
335 data might look like, in the 'below alpha' condition, we selected, for each lab in an RRR, the outcome
336 with a p -value closest below .05 (or, if there were no p -values below .05, the lowest value). These two
337 approaches ('most significant,' 'below alpha') attempted to select a single result from the multiverse, but
338 it may be that several effect sizes have equivalent p -values due to being based on exactly the same sample.
339 If so, we picked the effect size with the fewest researcher DFs deviating from their default option a.

340 Third, we represent a p -hacker who is satisfied with any significant effect size they encounter (in
341 the expected direction), by picking a random effect size out of those that were statistically significant
342 ('random significant'). If no effect sizes were significant, the effect size with the lowest p -value was
343 picked. Fourth and finally, when a p -hacking researcher tries multiple analyses, they might choose to
344 report the analysis that resulted in the smallest p -value. However, selecting the result with the smallest
345 p -value across the full multiverse suggests that the p -hacker systematically explored the full multiverse to
346 find the strongest possible effect, whereas reality probably consists of a more ad hoc and limited search.
347 Hence, we represent a 'bounded' search by 1) randomly drawing 100 possible outcomes and 2) out of
348 these 100 outcomes selecting the one with the smallest p -value.

349 With all our biased selection mechanisms (i.e., excluding the random draw and original meta-analytic
350 results, see Table 3), we applied a "hypothesized direction filter." That is, when selecting an effect size at
351 the lab level, we excluded all effect sizes that were in the opposite direction of the originally predicted
352 effect (osf.io/r2dum). If there were no effect sizes in the predicted direction, we excluded all significant
353 effect sizes in the 'wrong' direction and selected outcomes from the remainder. We added this filter
354 because we believe researchers who apply selective reporting in reality are unlikely to be agnostic about
355 the direction of their effect, and our focus in this study is on selective reporting, not Hypothesizing After
356 Results are Known (HARKing, Kerr, 1998)

357
358

Table 3.
Summary of outcome selection mechanisms

Selection mechanism	Hypothesized Direction Filter	Single outcome	Description
Pre-registered	No	Yes	The original RRR meta-analytic average effect size with pre-registered decisions.
Random draw	No	No	The average point estimate and upper/lower 95% CI from 10^5 meta-analyses randomly sampled from all possible meta-analyses.
Most significant	Yes	Yes	Select the effect size in the multiverse with the smallest p -value.
Below alpha	Yes	Yes	Select the effect size in the multiverse with a p -value closest below $p = 0.05$. If no p -value below the cutoff, pick the smallest.
Random significant	Yes	No	Identical to the random draw, but with effect sizes first limited to only significant effect sizes.
Bounded significant	Yes	No	We drew 100 effect sizes from a lab's multiverse, and selected the effect size with the lowest p -value. This was repeated 10^5 times, resulting in 10^5 values per lab. These were then meta-analyzed and summarized as above for the random draw.

Note:

Description of different implemented selection mechanisms for selecting effect sizes at the lab-level to meta-analyze. 'Hypothesized Direction Filter' = exclude effect sizes not in the predicted direction (yes/no), 'Single outcome' = selection mechanism resulting in a single meta-analytic result (yes/no). Code to reproduce table: osf.io/jehpy/.

359
360
361
362
363
364

Publication bias, the complete suppression of a study being published, and selective reporting (selection of reported results amongst multiple possibilities) are closely related, and it is intuitively appealing to believe correcting for publication bias may be sufficient for generally removing biases in the meta-analytic data (e.g., Kvarven et al., 2019). We applied three publication bias correction methods [PET-PEESE, 3PSM, and p -uniform*; Stanley & Doucouliagos (2014); Vevea & Hedges (1995); van Aert & van Assen (2020)] to examine their applicability to selective reporting in the absence of publication bias.

365 RESULTS

366
367
368
369
370
371
372
373

After excluding conditions that resulted in fewer than 24 participants per experimental group, 8/14 RRR multiverses decreased in size, as can be seen in Table 4. The absolute decrease was largest for the largest multiverses (RRR05 and RRR07), with RRR07 showing the largest absolute decrease and decreasing from 2,621,440 to 525,680 (an 80% decrease) potential outcomes. However, the proportionally largest decrease was seen in RRR08, which decreased from 115,200 to 19,200 (83% decrease) potential outcomes. More importantly, as evidenced by the median number of multiverses per lab, even in the RRRs with relatively few researcher DFs, the researcher DFs jointly created thousands of alternative outcomes per lab. Nonetheless, many labs found zero significant (at $p = .05$) effect sizes within their multiverses. Across all

374 studies (counting labs with multiple DVs as separate studies), 205 / 294 (70%) encountered no significant
375 effect sizes in the hypothesized direction in their multiverses and 134 / 294 (46%) found no significant
376 effect sizes in any direction. The size of the multiverse was strongly correlated with the number of studies
377 that encountered significant effect sizes in the hypothesized direction within their multiverses (Pearson's r
378 = 0.63, 95% CI [0.14, 0.87]).

379
380

Table 4.
Multiverse sizes before and after filtering out outcomes with < 24 participants per experimental group

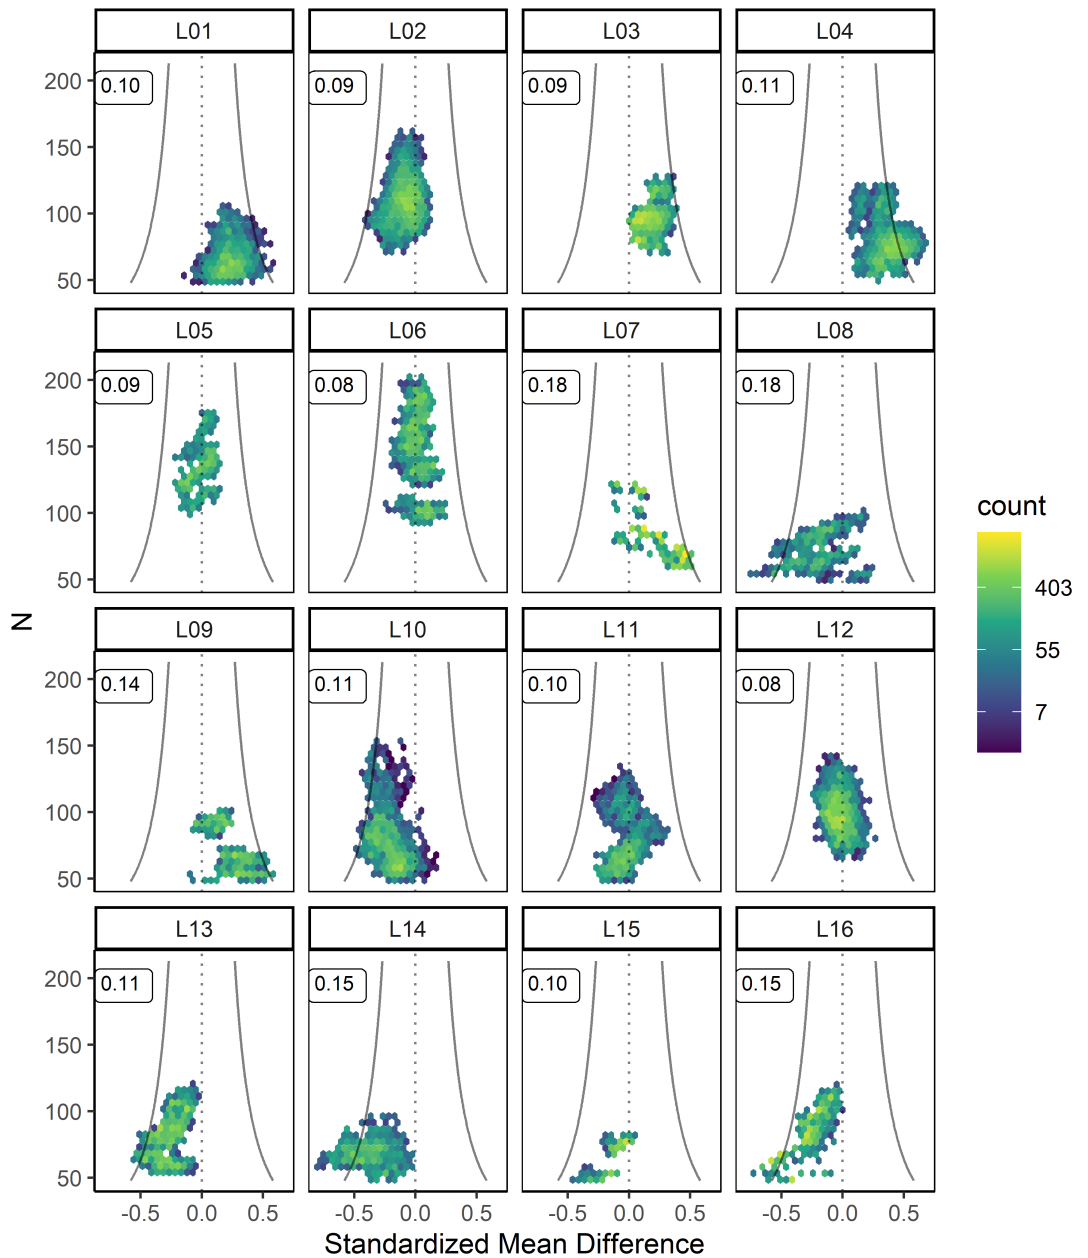
Meta-Analysis	Common DFs	Unique DFs	<i>N</i> [lq, uq]	Multiverse size before exclusion	Multiverse size after exclusion	Labs after exclusion	Labs with any sig.	Labs with hyp. sig.
RRR01	5	5	116 [107, 125]	3,840	3,840 (100%)	31	4 (13%)	4 (13%)
RRR02	5	5	88 [84, 98]	3,840	3,840 (100%)	23	10 (43%)	10 (43%)
RRR03 Attribution	6	2	84 [82, 84]	3,840	3,816 (99%)	10	5 (50%)	3 (30%)
RRR03 Intention	6	2	84 [82, 84]	3,840	3,840 (100%)	10	3 (30%)	1 (10%)
RRR03 Process	7	2	84 [83, 85]	7,680	7,680 (100%)	10	2 (20%)	0 (0%)
RRR04	5	3	76 [68, 90]	23,040	20,160 (88%)	24	14 (58%)	7 (29%)
RRR05 Exit	6	9	82 [70, 94]	2,488,320	1,503,904 (60%)	16	12 (75%)	7 (44%)
RRR05 Neglect	6	9	82 [70, 94]	2,488,320	1,540,176 (62%)	16	15 (94%)	10 (62%)
RRR06	7	5	96 [77, 111]	122,880	61,440 (50%)	17	9 (53%)	5 (29%)
RRR07	3	10	74 [63, 95]	2,621,440	525,680 (20%)	21	18 (86%)	10 (48%)
RRR08	8	4	79 [65, 102]	115,200	19,204 (17%)	39	31 (79%)	16 (41%)
RRR09 Behavior	8	4	169 [114, 218]	46,080	46,080 (100%)	26	16 (62%)	5 (19%)
RRR09 Hostility	8	4	168 [114, 218]	46,080	46,080 (100%)	26	13 (50%)	10 (38%)
RRR10	5	5	90 [77, 107]	11,520	3,808 (33%)	25	8 (32%)	1 (4%)

Note:

Meta-analytic distributions and estimates after excluding analytic choices that resulted in <24 participants per experimental group at the study-level. Three labs in RRR2 (L09, L17, L26) and one lab in RRR8 (L24) always had fewer than 24 participants and were excluded. 'Labs with any sig.' = number of labs (%) with any significant (at $p = .05$) effect size in their multiverse, 'Labs with hyp. sig.' = number of labs (%) with any significant (at $p = .05$) effect size in the hypothesized direction in their multiverse. DFs = Researcher Degrees of Freedom, 'Common DFs' = DF from a common list of potential DFs, 'Unique DFs' = study-unique DFs, M = Median study multiverse size, N [lower quartile, upper quartile] = Median study sample sizes across their multiverses. Code to reproduce table: osf.io/jehpy/.

381 **Lab multiverses**

382 There can be substantial variation in effect sizes within labs due to researcher DFs. Figure 2 shows
383 effect sizes across the multiverses for 16/24 labs in RRR04. Similar plots for all RRRs (or outcome
384 variables when an RRR contains multiple), including for all labs in RRR04, can be found in Supplement
385 E (osf.io/2htc6/).



386

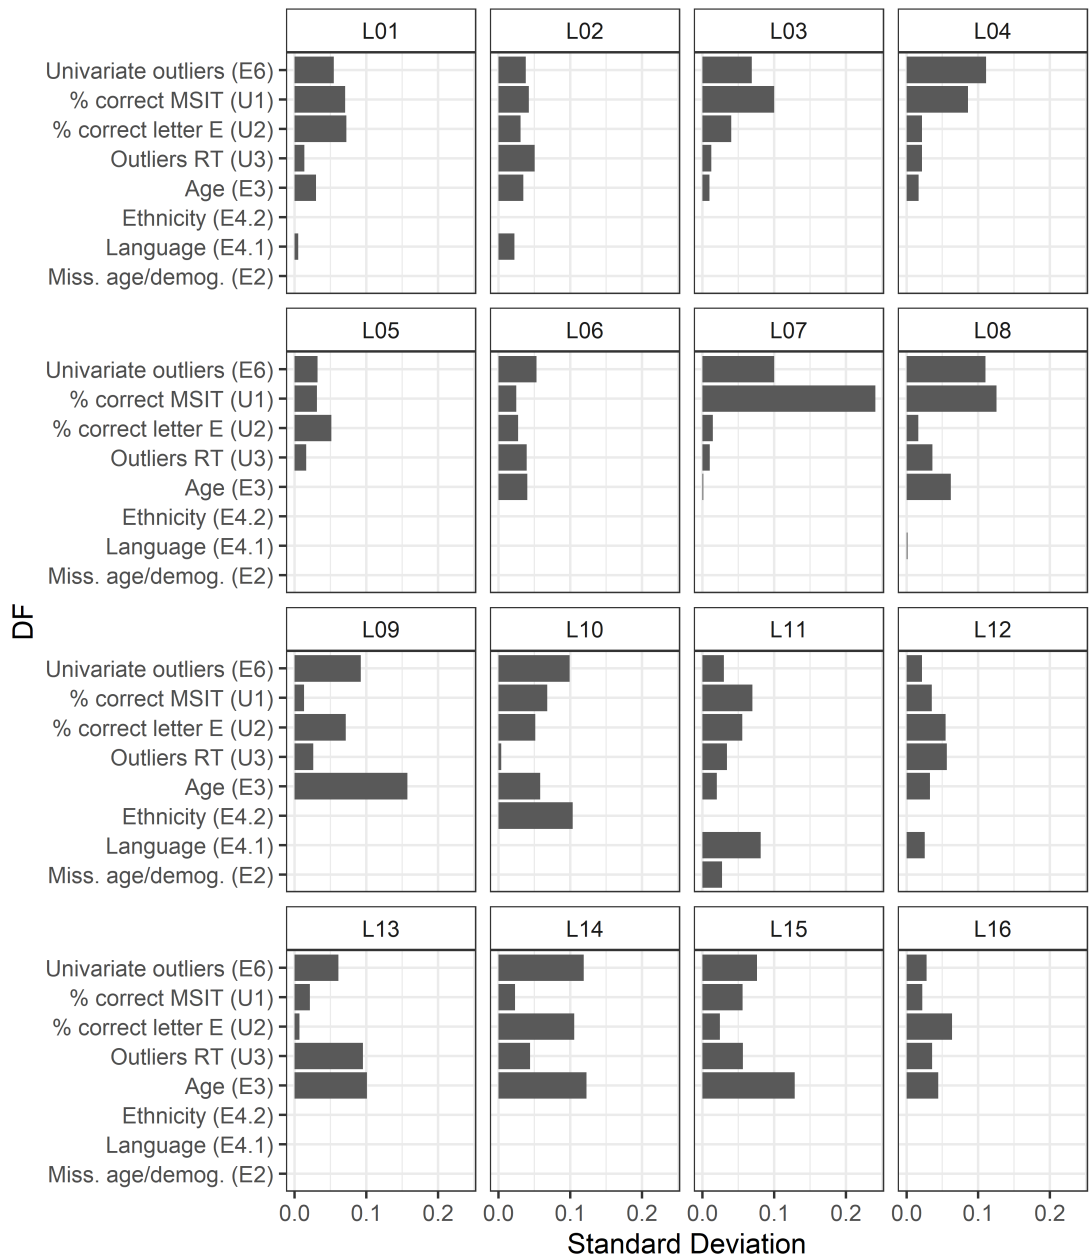
387 *Figure 2.* Arbitrary decisions in research cause underlying multiverse variability (UMV) in effect size
 388 estimates. Funnel plots showing the effect sizes based on the multiverses in 16 labs for RRR04, after
 389 removing cases where $n < 24$ in either experimental group. Values in upper left corner of each facet are
 390 UMV for each lab. For legibility, 16/24 RRR04 labs are shown; the figure including all labs is available
 391 in Supplement E (osf.io/2htc6/). L01 – L16 are lab indicators. Solid lines are funnel lines based on the
 392 t -distribution. Effect sizes falling outside the funnel lines are statistically significant at $\alpha = .05$ using a
 393 two-tailed test. Dotted lines indicate zero effect size. Colors in the funnel plots indicate the frequency of
 394 occurrence of an effect size. Brighter colors indicate that an effect size occurred more often. N = total

395 sample size. Code to reproduce figure: osf.io/thuyk/.

396 Overall in Figure 2, statistically significant observations (indicated by observations falling outside the
397 funnel lines) were rare (median = 0.87%, interquartile range = 0 - 3%). There were labs with a higher
398 proportion of significant outcomes (L14 = 25%, L04 = 24%, L16 = 17%), but in only one case were these
399 in the hypothesized direction (L04). The median underlying multiverse variability (UMV) across labs in
400 Figure 2 was 0.1SD, interquartile range (IQR) = 0.09 - 0.15. Effect sizes could change by as much as $d =$
401 0.97 (L08). Pearson's correlation based on the 16 labs in Figure 2 between UMV and sample size before
402 applying researcher DFs was $r = -0.51$.

403 Across RRRs, the median lab UMV was 0.11SD (IQR = 0.08 - 0.14) for SMD effect sizes and 0.07SD
404 (IQR = 0.04 - 0.12) for log OR, but researcher DFs could change effect sizes in a lab by as much as $d =$
405 1.27 (RRR05 Neglect, L10) and log OR = 1.31 (RRR01, L05). We expected the lab UMV, just as the
406 standard error, to be generally negatively correlated with (original) sample size. However, the median
407 correlation between lab UMV and sample size (before applying researcher DFs) within RRRs was $r =$
408 0.09 (IQR = -0.11 - 0.37). Hence, a large sample size does not ensure a small UMV.

409 There can be substantial variation between labs also in which researcher DF leads to variability in
410 effect sizes (Figure 3). Figure 3 shows the standard deviation (SD) in effect size within the labs from
411 Figure 2 when applying a single researcher DF. Despite identical study design across labs and the same
412 researcher DF being applied, no two bar plots look identical and labs differ in which researcher DF
413 creates the most variation. For example, in Lab 7 (L07) excluding participants based on different accuracy
414 criteria for the main DV (U1) resulted in the largest SD, whereas in Lab 6 (L06) using different criteria
415 for defining and excluding outliers (E6) led to the most variation in estimated effect size. Figures 2 and
416 3 together demonstrate that multiverse analyses should not necessarily be expected to be replicable in
417 new data, because across labs the same researcher DFs can yield different degrees of effect size variance
418 (Figure 2) and this variance can be primarily caused by different researcher DFs (figure 3). Note that the
419 effect sizes in Figure 2 arose from all possible combinations of the researcher DFs and not by applying
420 them separately as in Figure 3 (see 'analysis' methods section), which explains the larger range of effect
421 sizes in Figure 2.




422

423 *Figure 3.* The same arbitrary decisions have a different effect in comparable studies. Standard
 424 deviation (SD) in effect size estimates in 16 labs in RRR04 resulting from applying different researcher
 425 degrees of freedom (DFs) individually, after removing cases where $n < 24$ in either experimental group.
 426 For legibility, 16/24 RRR04 labs are shown; the figure including all labs is available in Supplement E
 427 (osf.io/2htc6/). L01 – L16 are lab indicators. Indicators in parentheses on the Y-axis (E2, E3, E4.1, E4.2,
 428 E6, U1, U2, U3) refer to DF codes for each coded DF (Supplement B osf.io/wj38n/). The Y-axis is
 429 ordered by median SD across labs. Code to reproduce figure: osf.io/thuyk/

430 Some researcher DFs generally contribute more to the UMV than others, and thus constitute larger

431 risk factors when considering selective reporting. Figure 4 shows the SDs in estimated effect sizes in labs
432 resulting from applying each researcher DF individually, disaggregated across RRRs and then aggregated
433 into their respective categories. How the composite score was computed from a scale with multiple
434 items had the largest median effect amongst SMD effects ($d = 0.04$, lower panel top row, Figure 4), and
435 in supplement F we demonstrate how removing (the source of variation in) this researcher DF (S2.c;
436 computing the composite score based on PCA) decreases overall UMV, showing how the removal of
437 sources of variation effectively decreases the risk of selective reporting. Figure 4 also show that excluding
438 participants on age had a relatively strong effect (median upper panel, log OR = 0.05, lower panel, $d =$
439 0.03. In supplement G we show that this effect is driven by the large degree of exclusions across the
440 options of the age researcher DF (E3), by comparing it with a version with more broad inclusion criteria.
441 That is, when researcher DFs result in datasets with less overlap (i.e., that are less correlated), the UMV
442 increases. Hence, we can predict that certain researcher DFs are of more serious concern under selective
443 reporting, although the observed effect in any given sample will depend on random fluctuations in the
444 sample.



445


446 *Figure 4.* Some arbitrary decisions tend to create more effect size variability than others. Standard
 447 deviation (SD) in effect size estimates in labs resulting from different researcher degrees of freedom
 448 (DFs) applied individually. The top panel shows results for Registered Replication Reports (RRRs) with
 449 an outcome measured as log odds ratio, and the lower panel for RRRs measured as standardized mean
 450 differences. The Y-axis is ordered by median effect size SD. Data are after removing cases where $n <$
 451 24 in either experimental group, disaggregating DFs across Registered Replication Reports (RRRs) and
 452 aggregating into categories. Indicators in parentheses on the Y-axis (S1, S2, E1, E2, E3, E4, E5, E6) refer
 453 to DF codes in Table 2 or (U) to DFs coded as unique for each research project. The Unique (U) category

454 was aggregated across all distinct unique DF. Code to reproduce figure: osf.io/thuyk/.

455 Unique researcher DFs only had the fifth highest median SD ($d = 0.02$) for SMD effects (lower panel
456 Figure 4), likely due to many unique researcher DFs having little effect. However, they also show the
457 largest range in possible outcomes. For example, within RRR07, choosing to exclude participants based
458 on whether they complied with the set time limit or not (U2) resulted in the largest median effect size
459 SD of all researcher DFs for that RRR (SD = 0.21, see also Supplement E). This same researcher DF in
460 RRR07 resulted in 4 out of the 5 highest effect size SDs in Figure 4. The remaining observation (third
461 from the right) belonging to RRR09 Hostility, lab 1, and resulted from choosing whether to exclude
462 participants based on their study major. Unique researcher DFs show less impact in the log odds ratio
463 effects (upper panel Figure 4), which may be due to fewer labs/researcher DFs, and most unique researcher
464 DFs in RRR01/02 only being applicable to a few of the constituent labs (see Supplement B). For example,
465 only three labs included a comprehension check (U1), and only three (different) labs coded ‘familiarity
466 with effect’ (U2).

467 **Meta-analytic multiverses**

468 Variability in effect sizes within labs due to researcher DFs implies that many different meta-analytic
469 outcomes are possible. How and which effect sizes were selected in labs will change meta-analytic
470 results. Figure 5 shows multiple meta-analytic average effect size estimates for all outcome variables,
471 depending on how effect sizes were selected in the constituent labs. The grey density curves indicate the
472 empirical distributions of meta-analytic point estimates across multiverses for each outcome. UMV in
473 point estimates ranged from 0.02SD (RRR09 Hostility) to 0.04SD (RRR03 Intention) for standardized
474 mean differences and for log odds ratios rounded to 0.02 SD for both RRR01 and RRR02. When outcomes
475 were selected through a preregistered decision procedure (purple squares, Figure 5), meta-analytic mean
476 estimates were generally close to the mean of the estimated multiverse distributions, and matched the
477 random draw estimates well (pink crosses, Figure 5). The point estimate of the random draws is by
478 definition identical to the mean of the multiverse distribution, whereas the lower/upper bounds are the
479 average 95% CIs of these draws.



480

481 *Figure 5.* Selective reporting in labs results in overestimates in meta-analysis. Meta-analytic distribu-
 482 tions and estimates after excluding analytic decisions that resulted in $n < 24$ participants per experimen-
 483 tal group at the study-level. Selection mechanism = how effect sizes were selected at the study-level, either
 484 by p -hacking (“Most significant,” “Bounded significant,” “Random significant,” “Below alpha,”), preregis-
 485 tered decisions (“preregistered”), or random selection (“Random draw”). “Random draw” and “Random
 486 significant” are averages across 10^5 draws from the meta-analytic multiverse, whereas other selection
 487 mechanisms are a single outcome. M = Median study multiverse size, N [lower quartile, upper quartile]
 488 = Median study sample sizes across their multiverses, * = effect size sign changed (RRR01, RRR02,

489 RRR05 Exit, RRR05 Neglect, RRR10) so that hypothesized effect size (and *p*-hacking direction) was
490 positive for all meta-analyses. Code to reproduce figure: osf.io/thuyk/.

491 When researcher DFs were combined with biased selection of effect sizes in labs (*p*-hacking), meta-
492 analytic mean estimates were also more extreme in the predicted direction (Figure 5). As expected,
493 selecting the most significant effect size in each lab (yellow stars) and then meta-analyzing resulted in the
494 most extreme mean estimates. Other *p*-hacking approaches (in Figure 5: turquoise triangles, red circles,
495 and green squares) resulted in similar estimates. This similarity in outcome between biased selection
496 mechanisms can be mostly attributed to the low number of significant results across labs and our biased
497 selection procedures resulting in the same results if there were no significant outcomes in a lab (205 / 294
498 studies across all RRRs, counting labs with multiple DVs as separate studies).

499 There is a tendency for projects with larger multiverses (e.g., RRR05, RRR06, and RRR07, as can be
500 seen in Figure 5) to have more extreme estimated effect sizes when *p*-hacked. The difference between the
501 average random draw (pink crosses) and the estimates based on the most significant effect sizes (yellow
502 stars) ranged from 0.1 to 0.48 for SMDs and was about 0.1 for log OR. The correlation between effect
503 size inflation and multiverse size was $r = 0.77$ for SMD effect sizes. The most extreme case corresponded
504 to RRR07 (Figure 5), where the difference in meta-analytic average effect size estimate between the
505 average random draw (pink cross, $d = -0.03$, 95% CI [-0.13, 0.07]) and the estimate based on the most
506 significant effect sizes (yellow star, $d = 0.45$, 95% CI [0.34, 0.56]) was an increase of almost 0.5SD in the
507 predicted direction. Applying publication bias correction methods (PET-PEESE, 3PSM and *p*-uniform*)
508 did not lead to improvements in estimated average effect size estimates (Supplement H), in line with other
509 research that has shown publication bias correction methods as unlikely to be useful in correcting for
510 selective reporting (Carter et al., 2019; van Aert et al., 2016).

511 DISCUSSION

512 In this article, we performed multiverse analysis across multiple direct replication studies, using empirical
513 data from ten Registered Replication Reports (RRRs). Even though the preregistrations used in the
514 actual RRRs limited the effect of researcher DFs in the original analyses, the open data from these
515 extensive studies enabled us to use a counter-factual design to see what biases could have emerged
516 if the studies (and ensuing meta-analyses that included them) had not been preregistered and could
517 have been subjected to selective reporting based on significance of the outcomes in the primary studies
518 (*p*-hacking). We identified researcher DFs based on common decisions in the associated literature for
519 each outcome variable, computed all possible outcomes across direct replications, and examined the
520 variance in these so-called multiverses. We then combined effect sizes from the multiverses of each

521 direct replication in meta-analysis and examined the consequences of different mechanisms for selecting
522 effect sizes for inclusion. Our analyses highlight that multiverse analyses typically yielded thousands of
523 different outcomes within single studies, that multiverse patterns of variation differed across labs using
524 the same protocol, and that selective reporting of outcomes in primary studies could bias meta-analytic
525 results, despite their status as a ‘gold-standard’ of evidence. We also found that original sample size
526 correlates at most weakly with potential for selective reporting (as measured by UMV), suggesting a larger
527 sample size does not protect against selective reporting. Yet interestingly, 205 / 294 studies (counting labs
528 with multiple DVs as separate studies) did not contain any significant (measured as $p \leq .05$) results in
529 their multiverses, suggesting p -hacking null results into significance may be more difficult than expected
530 considering the sheer number of potential outcomes per study. We discuss these results and the limitations
531 of our own study in the remainder of the discussion.

532 **Defining the multiverse**

533 Creating a multiverse is an inherently subjective endeavor given that researchers might disagree about
534 which decisions are (approximately) equally reasonable (Steege et al., 2016). For example, although we
535 created our researcher DFs based on common practice in the associated literature, there are methodological
536 arguments to carefully consider the meaning and impact of outliers, or use outlier robust statistics (e.g.,
537 Rousseeuw & Hubert, 2011) rather than excluding them based on rules of thumb. In the same vein, it may
538 often be preferable to perform multiple or maximum likelihood imputation (e.g., Jakobsen et al., 2017) of
539 missing data points rather than excluding them. For this reason, we have endeavored to structure our data
540 such that a disagreeing reader familiar with R can explore the consequence of only including those of our
541 researcher DFs they consider reasonable, as demonstrated in Supplement F.

542 Del Giudice & Gangestad (2021) critically discussed what it means for a researcher DF to be
543 ‘reasonable.’ They rightly argued why some of the researcher DFs as used in previously published
544 multiverse analyses might not be equivalent on prior grounds or might not show equivalence for reasons
545 yet unknown. We admit that our researcher DF S1, which deletes items in measurement scales based
546 on the lowest item-rest correlations, could perhaps be non-equivalent on psychometric grounds. On
547 the one hand, deletion of the item with lowest item-rest correlation could heighten the reliability of the
548 scale, and hence increase genuine effects, if the deleted item were poorly performing. On the other hand,
549 deletion would have little effect on the reliability and hence on actual effects if the item performed as
550 well as other items in the scale. Hence, if there were some item(s) consistently performing poorly (in
551 the psychometric sense) across labs in an RRR, this researcher DF could increase true effects and be
552 of principled non-equivalence, at least for that RRR. Similarly, the researcher DF related to composite
553 scores (S2) could create genuinely different effects under some conditions. In Supplement F, we repeated

554 our analyses without the PCA option of this researcher DF which showed less variation in outcomes, as
555 expected if indeed the way composite scores are computed makes a difference. However, like all our DFs,
556 we included this researcher DF because we feel that it might be used in practice in unregistered studies.

557 One could also argue that multiverses with varying samples sizes are non-equivalent in the statistical
558 sense of sampling variability (Del Giudice & Gangestad, 2021). Smaller sample sizes generally result in
559 less statistical power to detect effects and hence larger p -values, which makes it more difficult to directly
560 compare p -values across multiverses differing in sample size. We do not consider the issue of power to
561 be a major concern for our study, as many of our analyses focus on effect size estimates and variability,
562 and because most RRRs in our sample (apart from RRR1/RRR2, based on our reading of the original
563 articles) appear to study null effects. In addition, although researcher DFs that lead to smaller sample
564 sizes would increase variability of outcomes regardless of true effect size, the strong sample overlap
565 between multiverse samples in a given study creates covariation between multiverses that diminishes total
566 sampling variability. In the end, which multiverses are considered reasonable will depend not only on
567 individual researchers' beliefs, but also on which decisions their research community considers acceptable
568 in terms of theoretical, methodological, and empirical standards (Del Giudice & Gangestad, 2021).

569 Notwithstanding selective reporting, researchers should recognize that researcher DFs, or decisions
570 treated as if they were researcher DFs, by themselves create another layer of uncertainty in study estimates;
571 we found a median UMV of 0.1SD in our SMD data, although this will differ depending on research
572 field and which DFs researchers find reasonable. As such, we advise researchers doing hypothesis-testing
573 research to 1) preregister the (single) analysis they believe is optimal for testing their hypothesis, motivate
574 why this is the case and report uncertainty estimates (e.g., confidence or credibility interval), and 2) include
575 a multiverse analysis as a sensitivity analysis (following the advice of Del Giudice & Gangestad, 2021)
576 and report their UMV. Although some may argue that preregistering a preferred analysis is contradictory
577 if the options in the multiverse analysis are considered equivalent, pragmatically, we believe that most
578 researcher DFs will not be exactly equivalent, and that most researchers will have a preferred analysis that
579 it would be useful to accompany with a sensitivity analysis. The goal of these recommendations is for the
580 research process to be transparent so that results act as credible evidence despite the potential effect of
581 researcher DFs on outcomes.

582 More generally, multi-lab collaboration, or regular, projects may wish to consider incorporating the
583 multiverse perspective already in the design of their studies, identifying which of their decisions are
584 largely arbitrary and collecting data on alternatives. Preregistration of research is likely to be helpful from
585 this perspective, in addition to its transparency enhancing properties, which are helpful when evaluating a
586 study for selective reporting.

587 **Exploring the multiverse**

588 The unique design of our project enabled us to examine the effect of researcher DFs (i.e., perform
589 multiverse analysis) across multiple direct replication studies. We observed that 1) the same researcher
590 DFs applied to direct replication studies resulted in widely varying distributions of effect sizes, and 2)
591 which researcher DF caused the variability within a study differed between direct replications. That
592 is, the effect of researcher DFs both within and across direct replication studies appeared unsystematic.
593 These results demonstrate that results of multiverse analysis in any single study, like other exploratory
594 analyses, are not necessarily replicable in new data. We believe this point is underappreciated among many
595 multiverse analysts. In addition, some researchers may be tempted to directly interpret the existence of
596 researcher DFs and resulting UMV as evidence of ‘hidden moderators’ (Van Bavel et al., 2016); currently
597 unknown moderators that explain why effect sizes differ between studies. However, the existence of the
598 multiverse does not by itself imply moderators as substantial variability and apparent moderating effects
599 may be found through sampling variance alone.

600 That the effect of researcher DFs both within and across direct replication studies was generally non-
601 systematic also corroborates previous findings of ours (Olsson-Collentine et al., 2020) that differences
602 in study results in social and cognitive psychology show little to no between-study heterogeneity, and
603 supports the conclusion that the best explanation for differences between effect sizes in (direct) replication
604 studies is typically the joint effect of sampling error and researcher DFs, possibly in combination with
605 selective reporting.

606 When we have a substantive researcher DF that we suspect of being a moderator, it may be most
607 useful to examine it from an empirical meta-analytic perspective. If we have a researcher DF at the
608 study level (e.g., measurement scale) with sufficient variation between primary studies, it is possible to
609 examine it as a moderator using meta-regression (e.g., Houwelingen et al., 2002). However, individual
610 level researcher DFs (e.g., age) are preferably examined in individual participant data (IPD) meta-analysis
611 to avoid the ecological fallacy (e.g., Stewart & Tierney, 2002). In the case of multiple dependent variables,
612 which might also be a researcher DF, potential systematic differences could be examined in multivariate
613 meta-analysis (e.g., Jackson et al., 2011). As with multiverse analysis, such moderator analyses should
614 primarily be considered exploratory and hypothesis-generating.

615 Researcher DFs in primary studies also add a layer of uncertainty to meta-analysis when those studies
616 are meta-analyzed. Researcher DFs in primary studies can change both point estimates and the associated
617 standard errors and can do so across multiple studies. Consequently, in meta-analysis they can influence
618 not only the meta-analytic point estimate but also the between-study variance. That said, the standard
619 deviations in point estimates over the meta-analytic multiverses in our meta-analyses were quite small,

620 with an UMV of at most 0.04SD amongst SMD effects. This is unsurprising: when researcher's decisions
621 are truly random and in the absence of selective reporting and publication bias, researcher DFs in primary
622 studies can be expected to cancel out across a sufficiently large number of studies. As such, researcher
623 DFs in primary studies (and resulting multiverses) are not a major concern for meta-analysts in the
624 absence of selective reporting. Unfortunately, the availability of such ideal data is not expected in most
625 meta-analyses. Even with ideal data, in a meta-analysis with a small number of primary studies, as is
626 common in medicine (Davey et al., 2011), researcher DFs are less likely to balance each other out and
627 meta-analytic UMV may be a larger concern.

628 Encouragingly, in our data, the meta-analytic point estimates based on preregistered studies often
629 fell close to the center of the meta-analytic distributions. Preregistration may generally lead to less
630 effect size inflation (Schäfer & Schwarz, 2019) by decreasing the risk of selective reporting through
631 increased transparency (although preregistrations are of varying quality or not always suitable, Bakker et
632 al., 2020; Pham & Oh, 2021). However, the connection between multiverse analysis and preregistration
633 may have been enhanced by the nature of our data: large collaborative projects, including researchers with
634 adversary hypotheses, which may have resulted in a 'wisdom of the crowd' selection of decisions amongst
635 researchers DFs. Alternatively, if no decision within a multiverse has a systematic effect, results from
636 any preregistered set of decisions from that multiverse would also be expected to coincide approximately
637 with the mean of the multiverse distribution when analyzed across samples. Regardless, to account for
638 UMV preregistered multi-lab collaborations (e.g., Moshontz et al., 2018) may offer a way forward in the
639 absence of more concrete theory (Fried, 2020), although as we saw in our results even such data is not a
640 guarantee for a point estimate at the center of the meta-analytic multiverse distribution. It is important
641 to be aware that a preregistered set of decisions nonetheless only represents a single universe from the
642 multiverse.

643 The extent of UMV in any given field depends on the multiverse created, and our estimates in this
644 study may only apply to our non-random sample of social and cognitive psychology research. UMV in
645 other fields could be either larger or smaller, but is unlikely to be non-existent, and it may be worthwhile to
646 study the UMV in different subfields to examine their susceptibility to selective reporting given normative
647 research behavior. Generally, the extent of bias introducible by selective reporting will depend on the
648 multiverse size and the UMV, and researcher DFs that affect these two factors to a larger extent will hence
649 contribute more risk of bias to a study.

650 **Selecting from the multiverse**

651 We do not mean to imply that exploration of researcher DFs are problematic per se. We view it as
652 important to study the robustness of conclusions in the context of a sensitivity analysis, of which a

653 multiverse analysis can be seen as an extensive (systematic) variant. Exploring factors that truly moderate
654 an effect or association can be valuable as long as the exploration is transparently reported and employs
655 rigorous statistical controls to guard against overfitting. What is problematic is the selective or incomplete
656 reporting from the multiverse of statistical results. Hence, it is important to evaluate studies for risk of
657 selective reporting when using them to make decisions (e.g., about setting up future research), or including
658 them in systematic reviews (Appelbaum et al., 2018; as recommended by e.g., PRISMA and MARS:
659 Moher et al., 2009). Both preregistrations and multiverse analyses will facilitate evaluation of a study's
660 selective reporting risk by making research decisions more transparent, and there are many selective
661 reporting protocols available (e.g., Page et al., 2018) that may assist in the evaluation.

662 There is a risk that researchers exploit (intentionally or not) researcher DFs to selectively report those
663 results from the multiverse that most strongly support their hypothesis. In extreme cases researcher DFs
664 and *p*-hacking can provide evidence for any desired conclusion; in one lab the effect size estimate changed
665 by as much as $d = 1.27$. More realistically, we found a median UMV of 0.1SD amongst 294 studies
666 (counting labs with multiple DVs as separate studies). Nonetheless, a median UMV of 0.1SD in a field
667 still implies that selective reporting can turn a statistically non-significant effect into a significant effect.

668 For many studies in our data this was not a concern. We found that in our data and given our researcher
669 DFs, about 70% of study multiverses did not contain a single significant result (measured as $p < .05$) in
670 the hypothesized direction. As most RRRs that made up our data had overall effect size estimates (based
671 on preregistered outcomes) not significantly different from zero, this coincides with previous findings
672 of ours that on average null results also tend to have very little heterogeneity (Olsson-Collentine et al.,
673 2020). This suggests that it may be more difficult to *p*-hack null results into significance than many expect.
674 We caution that this observation may no longer hold when applying other types of researcher DFs than
675 we were able to do (e.g., this may not apply when researcher DF options are less correlated, as in the
676 case of outcome switching), and that 30% of labs did contain multiverses with a mix of significant and
677 non-significant effect size estimates.

678 Relatedly, we found a correlation not significantly different from zero between sample size and
679 potential for selective reporting (as measured by UMV). This implies that sample size should not be
680 taken to be protective against selective reporting, as also corroborated by a simulation study by Stefan &
681 Schönbrodt (2022). We do note, however, that for genuine effects, larger sample sizes would increase
682 power thereby lowering the need to selectively report outcomes based on the multiverse and hence less
683 ensuing bias in estimated effects (Bakker et al., 2012).

684 Our counter-factual design allowed us to see what biases could have emerged if the studies (and
685 ensuing meta-analyses) had not been preregistered and could have been subjected to selective reporting

686 based on significance of the outcomes in the primary studies, as is the case for most meta-analyses. Our
687 analyses demonstrate the substantial bias in the hypothesized direction that may be incorporated into
688 meta-analytic effect size estimates due to selective reporting in primary studies. The possible inflation
689 of average effect size will depend on the proportion of meta-analyzed studies at risk of bias and the
690 strength of this bias. Evidence from Kvarven et al. (2019) based on social and cognitive psychology
691 research suggests meta-analyses may sometimes estimate effect sizes to be as much as a third larger
692 than in comparable multi-lab projects, although this also includes publication bias. Larger or smaller
693 differences may be more typical in other fields.

694 For meta-analysts using retrospective data, minimizing the risk of bias in their included data (i.e., by
695 only including preregistered data and evaluating it for selective reporting) may be the best option until the
696 practices of multiverse analysis or the sharing of raw data become widespread. Our results corroborate
697 recommendations from meta-analysis reporting protocols such as PRISMA and MARS (Appelbaum et al.,
698 2018; Moher et al., 2009) to always evaluate primary studies for risk of bias (here, selective reporting),
699 and we advise meta-analysts to study differences in outcome between studies identified as at high risk of
700 selective reporting bias and those at low risk. In line with previous research on the topic (Carter et al.,
701 2019; van Aert et al., 2016), our supplemental results show that existing publication bias methods should
702 not be relied on to correct for *p*-hacking.

703 The large variance in impact of researcher DFs across direct replication studies demonstrates that the
704 study-specific effect of a researcher DF, and related bias induced by selective reporting, is difficult to
705 predict and dependent on sampling error. Stefan & Schönbrodt (2022), who simulated the effects of many
706 different *p*-hacking methods in single studies, reach a similar conclusion in their simulations: “Apart from
707 the aggressiveness of *p*-hacking itself, our simulations showed that across all strategies, the severity of
708 *p*-hacking also depends on the environment in which *p*-hacking takes place, for example, the correlation
709 structure in the data” (p. 46). Our results using the RRR data indeed show that the correlations between
710 multiverses create variation that is generally smaller than the sampling variation one would expect under
711 independent sampling.

712 That is not to say that we cannot draw some conclusions about the expected (average) impact of
713 *p*-hacking different researcher DFs. The potential for effect size bias is larger when studies allow more
714 analyses to be run and when more variance is created by included researcher DFs. Supplement G
715 demonstrates that more overlapping (sub)samples created using alternative exclusions based on age
716 created less subsequent UMV, which is expected given that more overlap creates higher correlations
717 between alternative outcomes. As such, when considering effect size bias, we should (typically) be more
718 concerned about researcher DFs in which the options are less correlated, although high false positive rates

719 are possible in either case (see discussion by Friese & Frankenbach, 2020).

720 It may be insightful to do more complex modelling of selective reporting from a multiverse perspective,
721 including of non-intentional selective reporting, and we hope our data will also be useful to other
722 researchers interested in more complex modelling of research bias. Our modelling of it in this study was
723 relatively straightforward and we only attempted to model the outcomes of intentional selective reporting
724 (*p*-hacking). Nonetheless, our biased selection methods applied to empirical RRR data are similar to those
725 used in the simulations of a recent compendium of *p*-hacking methods (Stefan & Schönbrodt, 2022) and
726 are on par with other recent simulation studies of *p*-hacking in meta-analysis (Botella et al., 2021; Friese
727 & Frankenbach, 2020).

728 Contrary to these two simulation studies (Botella et al., 2021; Friese & Frankenbach, 2020) we found
729 that *p*-hacking with actual data and using fairly generic researcher DFs could cause substantial inflation
730 of meta-analytic average effect sizes also when the average effect appears to be null. Both Friese &
731 Frankenbach (2020) and Botella et al. (2021) run extensive simulation studies of the effect of *p*-hacking
732 across many conditions and Friese & Frankenbach (2020) consider how it interacts with publication bias,
733 something we did not do. We believe the difference in results is due the choice in both of these simulation
734 studies to *p*-hack results based on the common assumption that *p*-hacking leads to a peak of *p*-values
735 below 0.05 (Hartgerink, 2017). In the case of Friese & Frankenbach (2020), results were *p*-hacked to
736 a distribution with a mode of $p = 0.049$, and Botella et al. (2021), similarly *p*-hacked studies into the
737 region $.025 < p \leq .05$. We know from previous studies that the type of *p*-hacking matters; incremental
738 methods such as optional stopping that result in a peak below $p = .05$ have little effect on effect size
739 inflation, whereas methods such as outcome reporting bias have a large effect and do not result in a
740 peak (Francis, 2012; Kirkham et al., 2010; Stefan & Schönbrodt, 2022). As such, differences between
741 our results and these simulation studies are likely explained by the incremental *p*-hacking methods used
742 in these simulation studies as compared to our methods of selecting the lowest *p*-values or randomly
743 selecting one of the significant outcomes.

744 Unfortunately, there exists little evidence on which method of selection researchers use in practice.
745 Incremental *p*-hacking can still lead to concerning numbers of false positive results, as discussed by
746 both Friese & Frankenbach (2020) and Botella et al. (2021), and it is important to discuss that not all
747 types of *p*-hacking lead to concerning levels of effect size inflation. Nonetheless, our results show that
748 suggesting selective reporting is not a concern for meta-analytic results is inaccurate when considering
749 non-incremental *p*-hacking based on researcher DFs that we consider to be widely applicable across a
750 range of empirical studies.

751 Under some assumptions related to correlational structure of the overlapping data across multiverses,

752 we can be confident that the UMV and hence the potential bias due to selective reporting in a study is less
753 than the reported standard error. The effect sizes in a multiverse are dependent because they are based on
754 the same sample. Due to this dependence, the UMV will normally be smaller than the standard error in
755 a study for a fixed sample size and statistical model, since the variability based on independent data is
756 larger than of dependent data. In other words, if we know that the statistical model and sample size have
757 not changed in a study, and that there is no publication bias, then we can be confident that the UMV in
758 that study is less than its standard error estimate.

759 Consequently, at a fixed standard error the possible bias is always larger with publication bias than
760 with selective reporting due to the dependency between effect sizes in the multiverse. This suggests
761 that while *p*-hacking is likely more common than publication bias in the literature, being more resource
762 efficient, the distortion in the literature may be larger from publication bias when it does occur. Finally,
763 we note that the correlation between effect sizes within a multiverse also means that the independent
764 sample false positive rate (typically .05) should be expected to be lower when sampling effect sizes within
765 a dataset. We can observe this in our funnel plots, where substantially fewer than 1/20 effect sizes are
766 significant for most labs (i.e., fall outside the funnel lines).

767 **Limitations and constraints on generality**

768 Although we have attempted to accompany all our claims and findings in this article together with their
769 caveats, we wish to make explicit the limits of generalizability of claims based on the data and design
770 of our study. The included effects are neither a representative nor a random sample of effects from
771 psychology. We expect our conclusions to be robust for effects in social and cognitive psychology, but
772 specific values that we report (e.g., median UMV of 0.1 among SMD effect sizes) may not generalize
773 beyond our sample. The RRRs in our data overwhelmingly reported average results not significantly
774 different from zero (the exception being RRR1/RRR2). This allowed us to examine selective reporting
775 in its most critical context (i.e., in the likely absence of genuine effects), but means it would be good to
776 focus future research efforts on studying multiverses with non-null effects, as these likely create more
777 heterogeneity in results across labs (see Olsson-Collentine et al., 2020), and hence we would also expect
778 larger overall variability and UMV due to researcher DFs (see also Friese & Frankenbach, 2020). During
779 the process of this project new RRRs have been published some of which report non-null effects and could
780 be used for such further analyses (e.g., Elliott et al., 2021). In addition, we are aware of several projects
781 currently in progress with similar designs that will collect new data and provide additional evidence on
782 the impact of standard researcher DFs in different fields.

783 Researchers interested in the theoretical implications of the specific effects studied here should
784 carefully consider which researcher DFs they find reasonable (see Del Giudice & Gangestad, 2021)

785 before drawing conclusions. Our researcher DFs were chosen to match standard decisions in social and
786 cognitive psychology treated as researcher DFs, meaning they were not guided by the substantive theory
787 of the studied effects, except as reflected by decisions made by the RRRs in data collection and analysis.
788 Moreover, if for instance age is theoretically expected to moderate an effect it should preferably be tested
789 formally instead of being used in a multiverse analysis. van Aert (2022) demonstrated how this can be
790 done by looking at the effect of age across labs in RRR9 (McCarthy et al., 2018) using IPD meta-analysis,
791 finding a small positive interaction ($p = .038$). Individual studies often do not have the power to detect
792 moderating effects, which also affects multiverse analyses.

793 We found that in only about 30% of studies did significant effect sizes in the hypothesized direction
794 emerge in their multiverses, despite the typically thousands of analyses in every study. This finding
795 suggests it is more difficult to turn apparent null results into significant results than might be expected
796 but is dependent on our selection of researcher DFs. Although we implemented an extensive number of
797 researcher DFs that we consider representative of researcher DFs that could be used in practice across a
798 range of social and cognitive psychological studies, our use of secondary (real) data means there were
799 many researcher DFs (e.g., Wicherts et al., 2016) that we could not apply but that might be applied in
800 real situations (e.g., outcome switching, which is known to have a large impact, or changing the analytic
801 model). As such, our selection of DFs and the resulting multiverse variances are unlikely to represent a
802 worst-case scenario. It is feasible that in real life more extreme statistical results are found. Generally, we
803 can expect researcher DFs with lower correlation between options, because of less sample overlap and/or
804 weaker correlations between (in)dependent variables, to result in larger multiverse variance.

805 Finally, it may be informative to analyze other multi-lab replications studies than those we included in
806 our study such as Many Labs 1 – 5 (e.g., Klein et al., 2018). The studies in our sample (mostly) studied a
807 single effect across multiple labs, whereas the Many Labs projects study many effects at the same time
808 across multiple labs. We examined the RRRs to be able to apply study-unique researcher DFs, but the
809 Many Labs design would allow examining the impact of applying a single set of researcher DFs on a large
810 sample of effects from social and cognitive psychology.

811 **Conclusion**

812 We have shown that researcher degrees of freedom offer a wide array of potential outcomes in relatively
813 standard psychological studies and demonstrated how selective reporting based on these researcher degrees
814 of freedom creates bias in meta-analytic effect size estimates that may undermine the credibility of many
815 meta-analyses. Preregistration is a methodological solution to researcher degrees of freedom enabling
816 selective reporting, whereas a statistical solution is to perform multiverse analysis of results. These two
817 transparency-enhancing practices can also be applied together, although our analyses of multiverses across

818 direct replications highlight that multiverse analyses in single studies should not necessarily be expected to
819 replicate in new data. Due to dependencies between effect sizes within multiverses, exploring multivariate
820 approaches to multiverse analysis may be a useful next step in helping to address uncertainties and biases
821 in primary studies due to researcher degrees of freedom.

822 REFERENCES

- 823 Agnoli, F., Wicherts, J. M., Veldkamp, C. L. S., Albiero, P., & Cubelli, R. (2017). Questionable
824 research practices among Italian research psychologists. *PLOS ONE*, *12*(3), e0172792. <https://doi.org/10.1371/journal.pone.0172792>
825
- 826 Alogna, V. K., Attaya, M. K., Aucoin, P., Bahník, Š., Birch, S., Birt, A. R., Bornstein, B. H., Bouwmeester,
827 S., Brandimonte, M. A., Brown, C., Buswell, K., Carlson, C., Carlson, M., Chu, S., Cislak, A.,
828 Colarusso, M., Colloff, M. F., Dellapaolera, K. S., Delvenne, J.-F., ... Zwaan, R. A. (2014).
829 Registered Replication Report: Schooler and Engstler-Schooler (1990). *Perspectives on Psychological*
830 *Science*, *9*, 556–578. <https://doi.org/10.1177/1745691614545653>
- 831 Appelbaum, M., Cooper, H., Kline, R. B., Mayo-Wilson, E., Nezu, A. M., & Rao, S. M. (2018).
832 Journal article reporting standards for quantitative research in psychology: The APA Publications
833 and Communications Board task force report. *The American Psychologist*, *73*(1), 3–25. <https://doi.org/10.1037/amp0000191>
834
- 835 Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The Rules of the Game Called Psychological
836 Science. *Perspectives on Psychological Science*, *7*(6), 543–554. [https://doi.org/10.1177/](https://doi.org/10.1177/1745691612459060)
837 [1745691612459060](https://doi.org/10.1177/1745691612459060)
- 838 Bakker, M., Veldkamp, C. L. S., Assen, M. A. L. M. van, Cromptvoets, E. A. V., Ong, H. H., Nosek, B.
839 A., Soderberg, C. K., Mellor, D., & Wicherts, J. M. (2020). Ensuring the quality and specificity of
840 preregistrations. *PLOS Biology*, *18*(12), e3000937. [https://doi.org/10.1371/journal.](https://doi.org/10.1371/journal.pbio.3000937)
841 [pbio.3000937](https://doi.org/10.1371/journal.pbio.3000937)
- 842 Bakker, M., & Wicherts, J. M. (2014). Outlier removal, sum scores, and the inflation of the type I error
843 rate in independent samples t tests: The power of alternatives and recommendations. *Psychological*
844 *Methods*, *19*(3), 409–427. <https://doi.org/10.1037/met0000014>
- 845 Borenstein, M. (2009). Effect sizes for continuous data. In H. Cooper, L. Hedges, & J. Valentine
846 (Eds.), *The handbook of research synthesis and meta-analysis* (2nd ed., pp. 221–235). Russel Sage
847 Foundation.
- 848 Botella, J., Suero, M., Durán, J. I., & Blazquez, D. (2021). The small impact of p-hacking marginally
849 significant results on the meta-analytic estimation of effect size. *Anales de Psicología*, *37*(1), 178–187.

850 <https://doi.org/10.6018/analesps.433051>

851 Botvinik-Nezer, R., Holzmeister, F., Camerer, C. F., Dreber, A., Huber, J., Johannesson, M., Kirchler,
852 M., Iwanir, R., Mumford, J. A., Adcock, R. A., Avesani, P., Baczkowski, B. M., Bajracharya,
853 A., Bakst, L., Ball, S., Barilari, M., Bault, N., Beaton, D., Beitner, J., ... Schonberg, T. (2020).
854 Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810), 84–88.
855 <https://doi.org/10.1038/s41586-020-2314-9>

856 Bouwmeester, S., Verkoeijen, P. P. J. L., Aczel, B., Barbosa, F., Bègue, L., Brañas-Garza, P., Chmura, T.
857 G. H., Cornelissen, G., Døssing, F. S., Espín, A. M., Evans, A. M., Ferreira-Santos, F., Fiedler, S.,
858 Flegr, J., Ghaffari, M., Glöckner, A., Goeschl, T., Guo, L., Hauser, O. P., ... Wollbrant, C. E. (2017).
859 Registered Replication Report: Rand, Greene, and Nowak (2012). *Perspectives on Psychological*
860 *Science*, 12(3), 527–542. <https://doi.org/10.1177/1745691617693624>

861 Breznau, N., Rinke, E. M., Wuttke, A., Adem, M., Adriaans, J., Alvarez-Benjumea, A., Andersen, H.
862 K., Auer, D., Azevedo, F., Bahnsen, O., Balzer, D., Bauer, G., Bauer, P. C., Baumann, M., Baute,
863 S., Benoit, V., Bernauer, J., Berning, C., Berthold, A., ... Nguyen, H. H. V. (2021). Observing
864 Many Researchers using the Same Data and Hypothesis Reveals a Hidden Universe of Uncertainty.
865 *MetaArXiv*. <https://doi.org/10.31222/osf.io/cd5j9>

866 Burton, A., & Altman, D. G. (2004). Missing covariate data within cancer prognostic studies: A review
867 of current reporting and proposed guidelines. *British Journal of Cancer*, 91(1), 4–8. <https://doi.org/10.1038/sj.bjc.6601907>

868

869 Carter, E. C., Schönbrodt, F. D., Gervais, W. M., & Hilgard, J. (2019). Correcting for Bias in Psychology:
870 A Comparison of Meta-Analytic Methods: *Advances in Methods and Practices in Psychological*
871 *Science*, 2(2), 115–144. <https://doi.org/10.1177/2515245919847196>

872 Cheung, I., Campbell, L., LeBel, E. P., Ackerman, R. A., Aykutoğlu, B., Bahník, Š., Bowen, J. D., Bredow,
873 C. A., Bromberg, C., Caprariello, P. A., Carcedo, R. J., Carson, K. J., Cobb, R. J., Collins, N. L.,
874 Corretti, C. A., DiDonato, T. E., Ellithorpe, C., Fernández-Rouco, N., Fuglestad, P. T., ... Yong, J.
875 C. (2016). Registered Replication Report: Study 1 From Finkel, Rusbult, Kumashiro, & Hannon
876 (2002). *Perspectives on Psychological Science*, 11(5), 750–764. <https://doi.org/10.1177/1745691616664694>

877

878 Coburn, K. M., & Vevea, J. L. (2019). *Weightr: Estimating weight-function models for publication bias*
879 [Manual]. <https://CRAN.R-project.org/package=weightr>

880 Curran, P. G. (2016). Methods for the detection of carelessly invalid responses in survey data. *Journal*
881 *of Experimental Social Psychology*, 66, 4–19. [https://doi.org/10.1016/j.jesp.2015.](https://doi.org/10.1016/j.jesp.2015.07.006)
882 07.006

- 883 Davey, J., Turner, R. M., Clarke, M. J., & Higgins, J. P. (2011). Characteristics of meta-analyses and their
884 component studies in the Cochrane Database of Systematic Reviews: A cross-sectional, descriptive
885 analysis. *BMC Medical Research Methodology*, *11*(1), 160. [https://doi.org/10.1186/](https://doi.org/10.1186/1471-2288-11-160)
886 [1471-2288-11-160](https://doi.org/10.1186/1471-2288-11-160)
- 887 Del Giudice, M., & Gangestad, S. W. (2021). A Traveler's Guide to the Multiverse: Promises, Pitfalls,
888 and a Framework for the Evaluation of Analytic Decisions. *Advances in Methods and Practices in*
889 *Psychological Science*, *4*(1), 22–26. <https://doi.org/10.1177/2515245920954925>
- 890 Eekhout, I., de Boer, R. M., Twisk, J. W. R., de Vet, H. C. W., & Heymans, M. W. (2012). Missing
891 Data: A Systematic Review of How They Are Reported and Handled. *Epidemiology*, *23*(5), 729–732.
892 <https://doi.org/10.1097/EDE.0b013e3182576cdb>
- 893 Eerland, A., Sherrill, A. M., Magliano, J. P., Zwaan, R. A., Arnal, J. D., Aucoin, P., Berger, S. A.,
894 Birt, A. R., Capezza, N., Carlucci, M., Crocker, C., Ferretti, T. R., Kibbe, M. R., Knepp, M. M.,
895 Kurby, C. A., Melcher, J. M., Michael, S. W., Poirier, C., & Prenoveau, J. M. (2016). Registered
896 Replication Report: Hart & Albarracín (2011). *Perspectives on Psychological Science*, *11*(1), 158–171.
897 <https://doi.org/10.1177/1745691615605826>
- 898 Elliott, E. M., Morey, C. C., AuBuchon, A. M., Cowan, N., Jarrold, C., Adams, E. J., Attwood, M.,
899 Bayram, B., Beeler-Duden, S., Blakstvedt, T. Y., Büttner, G., Castelain, T., Cave, S., Crepaldi, D.,
900 Fredriksen, E., Glass, B. A., Graves, A. J., Guitard, D., Hoehl, S., . . . Voracek, M. (2021). Multilab
901 Direct Replication of Flavell, Beach, and Chinsky (1966): Spontaneous Verbal Rehearsal in a Memory
902 Task as a Function of Age. *Advances in Methods and Practices in Psychological Science*, *4*(2).
903 <https://doi.org/10.1177/25152459211018187>
- 904 Flake, J. K., Pek, J., & Hehman, E. (2017). Construct Validation in Social and Personality Research:
905 Current Practice and Recommendations. *Social Psychological and Personality Science*, *8*(4), 370–378.
906 <https://doi.org/10.1177/1948550617693063>
- 907 Francis, G. (2012). The Psychology of Replication and Replication in Psychology. *Perspectives on*
908 *Psychological Science*, *7*(6), 585–594. <https://doi.org/10.1177/1745691612459520>
- 909 Franco, A., Malhotra, N., & Simonovits, G. (2016). Underreporting in Psychology Experiments: Evidence
910 From a Study Registry. *Social Psychological and Personality Science*, *7*(1), 8–12. [https://doi.](https://doi.org/10.1177/1948550615598377)
911 [org/10.1177/1948550615598377](https://doi.org/10.1177/1948550615598377)
- 912 Fried, E. I. (2020). Lack of Theory Building and Testing Impedes Progress in The Factor and Network
913 Literature. *Psychological Inquiry*, *31*(4), 271–288. [https://doi.org/10.1080/1047840X.](https://doi.org/10.1080/1047840X.2020.1853461)
914 [2020.1853461](https://doi.org/10.1080/1047840X.2020.1853461)

- 915 Friese, M., & Frankenbach, J. (2020). P-Hacking and publication bias interact to distort meta-analytic
916 effect size estimates. *Psychological Methods*, 25(4), 456–471. [https://doi.org/10.1037/
917 met0000246](https://doi.org/10.1037/met0000246)
- 918 Fusar-Poli, P., Radua, J., Frascarelli, M., Mechelli, A., Borgwardt, S., Fazio, F. D., Biondi, M., Ioannidis,
919 J. P. A., & David, S. P. (2014). Evidence of reporting biases in voxel-based morphometry (VBM)
920 studies of psychiatric and neurological disorders. *Human Brain Mapping*, 35(7), 3052–3065. <https://doi.org/10.1002/hbm.22384>
- 922 Goldacre, B., Drysdale, H., Dale, A., Milosevic, I., Slade, E., Hartley, P., Marston, C., Powell-Smith,
923 A., Heneghan, C., & Mahtani, K. R. (2019). COMPare: A prospective cohort study correcting and
924 monitoring 58 misreported trials in real time. *Trials*, 20(1), 118. [https://doi.org/10.1186/
925 s13063-019-3173-2](https://doi.org/10.1186/s13063-019-3173-2)
- 926 Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., Brand, R.,
927 Brandt, M. J., Brewer, G., Bruyneel, S., Calvillo, D. P., Campbell, W. K., Cannon, P. R., Carlucci, M.,
928 Carruth, N. P., Cheung, T., Crowell, A., De Ridder, D. T. D., Dewitte, S., ... Zwieneberg, M. (2016).
929 A Multilab Preregistered Replication of the Ego-Depletion Effect. *Perspectives on Psychological
930 Science*, 11(4), 546–573. <https://doi.org/10.1177/17456916166652873>
- 931 Hartgerink, C. H. J. (2017). Reanalyzing Head et al. (2015): Investigating the robustness of widespread
932 p-hacking. *PeerJ*, 5, e3068. <https://doi.org/10.7717/peerj.3068>
- 933 Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and
934 Brain Sciences*, 33(2-3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- 935 Henry, L., & Wickham, H. (2020). *Purrr: Functional programming tools* [Manual]. [https://CRAN.
936 R-project.org/package=purrr](https://CRAN.R-project.org/package=purrr)
- 937 Houwelingen, H. C. van, Arends, L. R., & Stijnen, T. (2002). Advanced methods in meta-analysis:
938 Multivariate approach and meta-regression. *Statistics in Medicine*, 21(4), 589–624. [https://doi.
939 org/10.1002/sim.1040](https://doi.org/10.1002/sim.1040)
- 940 Huntington-Klein, N., Arenas, A., Beam, E., Bertoni, M., Bloem, J. R., Burli, P., Chen, N., Grieco, P.,
941 Ekpe, G., Pugatch, T., Saavedra, M., & Stopnitzky, Y. (2021). The influence of hidden researcher
942 decisions in applied microeconomics. *Economic Inquiry*, 59(3), 944–960. [https://doi.org/
943 10.1111/ecin.12992](https://doi.org/10.1111/ecin.12992)
- 944 Ioannidis, J. P. A. (2008). Why Most Discovered True Associations Are Inflated. *Epidemiology*, 19(5),
945 640–648. <https://doi.org/10.1097/EDE.0b013e31818131e7>
- 946 Ioannidis, J. P. A. (2005). Why Most Published Research Findings Are False. *PLoS Medicine*, 2(8), e124.
947 <https://doi.org/10.1371/journal.pmed.0020124>

- 948 Jackson, D., Riley, R., & White, I. R. (2011). Multivariate meta-analysis: Potential and promise. *Statistics*
949 *in Medicine*, 30(20), 2481–2498. <https://doi.org/10.1002/sim.4172>
- 950 Jakobsen, J. C., Gluud, C., Wetterslev, J., & Winkel, P. (2017). When and how should multiple im-
951 putation be used for handling missing data in randomised clinical trials – a practical guide with
952 flowcharts. *BMC Medical Research Methodology*, 17(1), 162. <https://doi.org/10.1186/s12874-017-0442-1>
- 954 John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the Prevalence of Questionable Research
955 Practices With Incentives for Truth Telling. *Psychological Science*, 23(5), 524–532. <https://doi.org/10.1177/0956797611430953>
- 957 Kerr, N. L. (1998). HARKing: Hypothesizing After the Results are Known. *Personality and Social*
958 *Psychology Review*, 2(3), 196–217. https://doi.org/10.1207/s15327957pspr0203_4
- 959 Kirkham, J. J., Dwan, K. M., Altman, D. G., Gamble, C., Dodd, S., Smyth, R., & Williamson, P. R.
960 (2010). The impact of outcome reporting bias in randomised controlled trials on a cohort of systematic
961 reviews. *BMJ*, 340, c365. <https://doi.org/10.1136/bmj.c365>
- 962 Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M.,
963 Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R.,
964 Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., ... Nosek, B. A. (2018).
965 Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in*
966 *Methods and Practices in Psychological Science*, 1(4), 443–490. <https://doi.org/10.1177/2515245918810225>
- 968 Kvarven, A., Strømmland, E., & Johannesson, M. (2019). Comparing meta-analyses and preregistered
969 multiple-laboratory replication projects. *Nature Human Behaviour*, 4, 423–434. <https://doi.org/10.1038/s41562-019-0787-z>
- 971 Kyzas, P. A., Loizou, K. T., & Ioannidis, J. P. A. (2005). Selective Reporting Biases in Cancer Prognostic
972 Factor Studies. *JNCI: Journal of the National Cancer Institute*, 97(14), 1043–1055. <https://doi.org/10.1093/jnci/dji184>
- 974 Madden, K., Tai, K., Ali, Z., Schneider, P., Singh, M., Ghert, M., & Bhandari, M. (2019). Published
975 intimate partner violence studies often differ from their trial registration records. *Women & Health*,
976 59(1), 13–27. <https://doi.org/10.1080/03630242.2017.1421287>
- 977 McCarthy, R. J., Skowronski, J. J., Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R.,
978 Acar, O. A., Aczel, B., Bakos, B. E., Barbosa, F., Baskin, E., Bègue, L., Ben-Shakhar, G., Birt, A.
979 R., Blatz, L., Charman, S. D., Claesen, A., Clay, S. L., ... Yıldız, E. (2018). Registered Replication
980 Report on Srull and Wyer (1979). *Advances in Methods and Practices in Psychological Science*, 1(3),

981 321–336. <https://doi.org/10.1177/2515245918777487>

982 Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., & Group, T. P. (2009). Preferred Reporting Items for
983 Systematic Reviews and Meta-Analyses: The PRISMA Statement. *PLOS Medicine*, *6*(7), e1000097.
984 <https://doi.org/10.1371/journal.pmed.1000097>

985 Moshontz, H., Campbell, L., Ebersole, C. R., IJzerman, H., Urry, H. L., Forscher, P. S., Grahe, J.
986 E., McCarthy, R. J., Musser, E. D., Antfolk, J., Castille, C. M., Evans, T. R., Fiedler, S., Flake,
987 J. K., Forero, D. A., Janssen, S. M. J., Keene, J. R., Protzko, J., Aczel, B., ... Chartier, C. R.
988 (2018). The Psychological Science Accelerator: Advancing Psychology Through a Distributed
989 Collaborative Network. *Advances in Methods and Practices in Psychological Science*, *1*(4), 501–515.
990 <https://doi.org/10.1177/2515245918797607>

991 Nickerson, R. S. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of*
992 *General Psychology*, *2*(2), 175–220. <https://doi.org/10.1037/1089-2680.2.2.175>

993 O'Donnell, M., Nelson, L. D., Ackermann, E., Aczel, B., Akhtar, A., Aldrovandi, S., Alshaif, N.,
994 Andringa, R., Aveyard, M., Babincak, P., Balatekin, N., Baldwin, S. A., Banik, G., Baskin, E., Bell,
995 R., Białobrzaska, O., Birt, A. R., Boot, W. R., Braithwaite, S. R., ... Zrubka, M. (2018). Registered
996 Replication Report: Dijksterhuis and van Knippenberg (1998): *Perspectives on Psychological Science*,
997 *13*(2), 268–294. <https://doi.org/10.1177/1745691618755704>

998 Olsson-Collentine, A., Wicherts, J., Bakker, M., & Aert, R. C. M. van. (2019). *Meta multiverse OSF*
999 *repository*. <https://osf.io/j8yg2/>

1000 Olsson-Collentine, Anton, Wicherts, Jelte, Bakker, Marjan, & van Aert, Robbie. (2022). Meta multiverse
1001 replication material [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.7341292>

1002 Olsson-Collentine, A., Wicherts, J. M., & van Assen, M. A. L. M. (2020). Heterogeneity in direct
1003 replications in psychology and its association with effect size. *Psychological Bulletin*, *146*(10),
1004 922–940. <https://doi.org/10.1037/bul0000294>

1005 Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting
1006 satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*(4), 867–872.
1007 <https://doi.org/10.1016/j.jesp.2009.03.009>

1008 Page, M. J., McKenzie, J. E., & Higgins, J. P. T. (2018). Tools for assessing risk of reporting biases
1009 in studies and syntheses of studies: A systematic review. *BMJ Open*, *8*(3), e019703. <https://doi.org/10.1136/bmjopen-2017-019703>

1010

1011 Page, M. J., Sterne, J. A. C., Higgins, J. P. T., & Egger, M. (2020). Investigating and dealing with
1012 publication bias and other reporting biases in meta-analyses of health research: A review. *Research*
1013 *Synthesis Methods*, *12*, 248–259. <https://doi.org/10.1002/jrsm.1468>

- 1014 Palpacuer, C., Hammas, K., Duprez, R., Laviolle, B., Ioannidis, J. P. A., & Naudet, F. (2019). Vibration
1015 of effects from diverse inclusion/exclusion criteria and analytical choices: 9216 different ways to
1016 perform an indirect comparison meta-analysis. *BMC Medicine*, *17*(1), 174. [https://doi.org/](https://doi.org/10.1186/s12916-019-1409-3)
1017 [10.1186/s12916-019-1409-3](https://doi.org/10.1186/s12916-019-1409-3)
- 1018 Patel, C. J., Burford, B., & Ioannidis, J. P. A. (2015). Assessment of vibration of effects due to model
1019 specification can demonstrate the instability of observational associations. *Journal of Clinical Epi-*
1020 *demiology*, *68*(9), 1046–1058. <https://doi.org/10.1016/j.jclinepi.2015.05.029>
- 1021 Pham, M. T., & Oh, T. T. (2021). Preregistration Is Neither Sufficient nor Necessary for Good Science.
1022 *Journal of Consumer Psychology*, *31*(1), 163–176. <https://doi.org/10.1002/jcpy.1209>
- 1023 Pigott, T. D., Valentine, J. C., Polanin, J. R., Williams, R. T., & Canada, D. D. (2013). Outcome-Reporting
1024 Bias in Education Research. *Educational Researcher*, *42*(8), 424–432. [https://doi.org/10.](https://doi.org/10.3102/0013189X13507104)
1025 [3102/0013189X13507104](https://doi.org/10.3102/0013189X13507104)
- 1026 R Core Team. (2020). *R: A language and environment for statistical computing* [Manual]. R Foundation
1027 for Statistical Computing. <https://www.R-project.org/>
- 1028 Rankin, J., Ross, A., Baker, J., O'Brien, M., Scheckel, C., & Vassar, M. (2017). Selective outcome
1029 reporting in obesity clinical trials: A cross-sectional review: Reporting outcomes in obesity clinical
1030 trials. *Clinical Obesity*, *7*(4), 245–254. <https://doi.org/10.1111/cob.12199>
- 1031 Revelle, W. (2020). *Psych: Procedures for psychological, psychometric, and personality research*
1032 [Manual]. <https://CRAN.R-project.org/package=psych>
- 1033 Robinson, D., Hayes, A., & Couch, S. (2020). *Broom: Convert statistical objects into tidy tibbles*
1034 [Manual]. <https://CRAN.R-project.org/package=broom>
- 1035 Roese, N. J., & Vohs, K. D. (2012). Hindsight Bias. *Perspectives on Psychological Science*, *7*(5), 411–426.
1036 <https://doi.org/10.1177/1745691612454303>
- 1037 Rombach, I., Rivero-Arias, O., Gray, A. M., Jenkinson, C., & Burke, Ó. (2016). The current practice of
1038 handling and reporting missing outcome data in eight widely used PROMs in RCT publications: A
1039 review of the current literature. *Quality of Life Research*, *25*(7), 1613–1623. [https://doi.org/](https://doi.org/10.1007/s11136-015-1206-1)
1040 [10.1007/s11136-015-1206-1](https://doi.org/10.1007/s11136-015-1206-1)
- 1041 Rongen, J. J., & Hannink, G. (2016). Comparison of Registered and Published Primary Outcomes in
1042 Randomized Controlled Trials of Orthopaedic Surgical Interventions: *The Journal of Bone and Joint*
1043 *Surgery*, *98*(5), 403–409. <https://doi.org/10.2106/JBJS.15.00400>
- 1044 Rosati, P., Porzsolt, F., Ricciotti, G., Testa, G., Inglese, R., Giustini, F., Fiscarelli, E., Zazza, M.,
1045 Carlino, C., Balassone, V., Fiorito, R., & D'Amico, R. (2016). Major discrepancies between what
1046 clinical trial registries record and paediatric randomised controlled trials publish. *Trials*, *17*(1), 430.

- 1047 <https://doi.org/10.1186/s13063-016-1551-6>
- 1048 Rousseuw, P. J., & Hubert, M. (2011). Robust statistics for outlier detection. *WIREs Data Mining and*
1049 *Knowledge Discovery*, 1(1), 73–79. <https://doi.org/10.1002/widm.2>
- 1050 Sassenberg, K., & Ditrich, L. (2019). Research in Social Psychology Changed Between 2011 and
1051 2016: Larger Sample Sizes, More Self-Report Measures, and More Online Studies. *Advances in*
1052 *Methods and Practices in Psychological Science*, 2(2), 107–114. [https://doi.org/10.1177/](https://doi.org/10.1177/2515245919838781)
1053 [2515245919838781](https://doi.org/10.1177/2515245919838781)
- 1054 Schäfer, T., & Schwarz, M. A. (2019). The Meaningfulness of Effect Sizes in Psychological Research:
1055 Differences Between Sub-Disciplines and the Impact of Potential Biases. *Frontiers in Psychology*, 10.
1056 <https://doi.org/10.3389/fpsyg.2019.00813>
- 1057 Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., Bahník, Š., Bai, F.,
1058 Bannard, C., Bonnier, E., Carlsson, R., Cheung, F., Christensen, G., Clay, R., Craig, M. A., Dalla
1059 Rosa, A., Dam, L., Evans, M. H., Flores Cervantes, I., ... Nosek, B. A. (2018). Many Analysts,
1060 One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in*
1061 *Methods and Practices in Psychological Science*, 1(3), 337–356. [https://doi.org/10.1177/](https://doi.org/10.1177/2515245917747646)
1062 [2515245917747646](https://doi.org/10.1177/2515245917747646)
- 1063 Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed
1064 Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological*
1065 *Science*, 22(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>
- 1066 Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2020). Specification curve analysis. *Nature Human*
1067 *Behaviour*, 4(11), 1208–1214. <https://doi.org/10.1038/s41562-020-0912-z>
- 1068 Stanley, T. D., & Doucouliagos, H. (2014). Meta-regression approximations to reduce publication selection
1069 bias. *Research Synthesis Methods*, 5, 60–78. <https://doi.org/10.1002/jrsm.1095>
- 1070 Steegen, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing Transparency Through a
1071 Multiverse Analysis. *Perspectives on Psychological Science*, 11(5), 702–712. [https://doi.org/](https://doi.org/10.1177/1745691616658637)
1072 [10.1177/1745691616658637](https://doi.org/10.1177/1745691616658637)
- 1073 Stefan, A., & Schönbrodt, F. (2022). *Big Little Lies: A Compendium and Simulation of p-Hacking*
1074 *Strategies*. PsyArXiv. <https://doi.org/10.31234/osf.io/xy2dk>
- 1075 Stewart, L. A., & Tierney, J. F. (2002). To IPD or not to IPD?: Advantages and Disadvantages of
1076 Systematic Reviews Using Individual Patient Data. *Evaluation & the Health Professions*, 25(1),
1077 76–97. <https://doi.org/10.1177/0163278702025001006>
- 1078 Tabachnick, B. G., Fidell, L. S., & Ullman, J. B. (2007). *Using multivariate statistics* (Vol. 5). MA:
1079 Pearson.

- 1080 van Aert, R. C. M. (2022). Analyzing Data of a Multilab Replication Project With Individual Participant
1081 Data Meta-Analysis: A Tutorial. *Zeitschrift für Psychologie*, 230(1), 60–72. [https://doi.org/](https://doi.org/10.1027/2151-2604/a000483)
1082 10.1027/2151-2604/a000483
- 1083 van Aert, R. C. M. (2020). *Puniform: Meta-analysis methods correcting for publication bias* [Manual].
1084 <https://CRAN.R-project.org/package=puniform>
- 1085 van Aert, R. C. M., & van Assen, M. A. L. M. (2020). Correcting for Publication Bias in a Meta-Analysis
1086 with the P-uniform* Method. *MetaArXiv*. <https://doi.org/10.31222/osf.io/zqj9r>
- 1087 van Aert, R. C. M., Wicherts, J. M., & van Assen, M. A. L. M. (2016). Conducting Meta-Analyses
1088 Based on p Values: Reservations and Recommendations for Applying p -Uniform and p -Curve.
1089 *Perspectives on Psychological Science*, 11(5), 713–729. [https://doi.org/10.1177/](https://doi.org/10.1177/1745691616650874)
1090 1745691616650874
- 1091 Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Reply to Inbar: Contextual
1092 sensitivity helps explain the reproducibility gap between social and cognitive psychology. *Pro-*
1093 *ceedings of the National Academy of Sciences*, 113(34). [https://doi.org/10.1073/pnas.](https://doi.org/10.1073/pnas.1609700113)
1094 1609700113
- 1095 van Ginkel, J. R., Sijtsma, K., van der Ark, L. A., & Vermunt, J. K. (2010). Incidence of Missing Item
1096 Scores in Personality Measurement, and Simple Item-Score Imputation. *Methodology*, 6(1), 17–30.
1097 <https://doi.org/10.1027/1614-2241/a000003>
- 1098 Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R., McCarthy, R. J., Skowronski, J. J.,
1099 Acar, O. A., Aczel, B., Bakos, B. E., Barbosa, F., Baskin, E., Bègue, L., Ben-Shakhar, G., Birt, A.
1100 R., Blatz, L., Charman, S. D., Claesen, A., Clay, S. L., ... Yıldız, E. (2018). Registered Replication
1101 Report on Mazar, Amir, and Ariely (2008). *Advances in Methods and Practices in Psychological*
1102 *Science*, 1(3), 299–317. <https://doi.org/10.1177/2515245918781032>
- 1103 Vevea, J. L., & Hedges, L. V. (1995). A general linear model for estimating effect size in the presence of
1104 publication bias. *Psychometrika*, 60(3), 419–435. <https://doi.org/10.1007/BF02294384>
- 1105 Viechtbauer, W. (2010). Conducting Meta-Analyses in R with the **Metafor** Package. *Journal of Statistical*
1106 *Software*, 36(3). <https://doi.org/10.18637/jss.v036.i03>
- 1107 Voracek, M., Kossmeier, M., & Tran, U. S. (2019). Which Data to Meta-Analyze, and How? A
1108 Specification-Curve and Multiverse-Analysis Approach to Meta-Analysis. *Zeitschrift für Psychologie*,
1109 227(1), 64–82. <https://doi.org/10.1027/2151-2604/a000357>
- 1110 Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R. B., Albohn, D.
1111 N., Allard, E. S., Benning, S. D., Blouin-Hudon, E.-M., Bulnes, L. C., Caldwell, T. L., Calin-
1112 Jageman, R. J., Capaldi, C. A., Carfagno, N. S., Chasten, K. T., Cleeremans, A., Connell, L.,

1113 DeCicco, J. M., ... Zwaan, R. A. (2016). Registered Replication Report: Strack, Martin, & Stepper
1114 (1988). *Perspectives on Psychological Science*, 11(6), 917–928. [https://doi.org/10.1177/
1115 1745691616674458](https://doi.org/10.1177/1745691616674458)

1116 Wagenmakers, E.-J., Sarafoglou, A., & Aczel, B. (2022). One statistical analysis must not rule them all.
1117 *Nature*, 605(7910), 423–425. <https://doi.org/10.1038/d41586-022-01332-8>

1118 Wayant, C., Scheckel, C., Hicks, C., Nissen, T., Leduc, L., Som, M., & Vassar, M. (2017). Evidence of
1119 selective reporting bias in hematology journals: A systematic review. *PLOS ONE*, 12(6), e0178379.
1120 <https://doi.org/10.1371/journal.pone.0178379>

1121 Wicherts, J. M. (2012). *Cohort Differences in Big Five Personality Factors Over a Period of 25*
1122 *Years*. Data Archiving and Networked Services (DANS). [https://easy.dans.knaw.nl/
1123 ui/datasets/id/easy-dataset:51655](https://easy.dans.knaw.nl/ui/datasets/id/easy-dataset:51655)

1124 Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., Aert, R. C. M. van, & Assen,
1125 M. A. L. M. van. (2016). Degrees of Freedom in Planning, Running, Analyzing, and Reporting
1126 Psychological Studies: A Checklist to Avoid p-Hacking. *Frontiers in Psychology*, 7. [https:
1127 //doi.org/10.3389/fpsyg.2016.01832](https://doi.org/10.3389/fpsyg.2016.01832)

1128 Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. [https:
1129 //ggplot2.tidyverse.org](https://ggplot2.tidyverse.org)

1130 Wickham, H., & Bryan, J. (2019). *Readxl: Read excel files* [Manual]. [https://CRAN.R-project.
1131 org/package=readxl](https://CRAN.R-project.org/package=readxl)

1132 Wickham, H., François, R., Henry, L., & Müller, K. (2020). *Dplyr: A grammar of data manipulation*
1133 [Manual]. <https://CRAN.R-project.org/package=dplyr>

1134 Wickham, H., & Miller, E. (2020). *Haven: Import and export 'SPSS', 'stata' and 'SAS' files* [Manual].
1135 <https://CRAN.R-project.org/package=haven>

1136 Zhu, H. (2019). *kableExtra: Construct complex table with 'kable' and pipe syntax* [Manual]. [https:
1137 //CRAN.R-project.org/package=kableExtra](https://CRAN.R-project.org/package=kableExtra)