1 Meta Research: Replication of Significant Results—Modeling the Effects of p -Hacking

2 Rolf Ulrich

3 University of Tübingen

4 Jeff Miller

5 University of Otago

6 Please address editorial correspondence to: Rolf Ulrich, Department of Psychology,

7 University of Tübingen, Schleichstr. 4, 72076 Tübingen, Germany.

8 International FAX: +49-7071-29-2410.

9 email: ulrich@uni-tuebingen.de.

10 *Version of April 24, 2020. Please do not cite.*

11 Author Note

12 Address correspondence to Rolf Ulrich, Department of Psychology, University of
13 Tübingen, Schleichstr. 4, 72076 Tübingen, Germany, or Jeff Miller, Department of
14 Psychology, University of Otago, Dunedin, New Zealand. Electronic mail may be sent
15 to ulrich@uni-tuebingen.de or miller@psy.otago.ac.nz. We thank E.J. Wagenmakers for
16 helpful comments on a previous version of this manuscript.

Abstract

This article examines why many studies fail to replicate statistically significant published results. We address this issue within a general statistical framework that also allows us to include various questionable research practices that are thought to reduce replicability. The analyses indicate that the base rate of true effects is the major factor that determines the replication rate of scientific results. Specifically, for purely statistical reasons, replicability is low in research domains where true effects are rare (e.g., search for effective drugs in pharmacology). This point is under-appreciated in current scientific and media discussions of replicability.

Keywords: Replicability, p -hacking, base rate, power, false positives, mathematical modelling of research

Meta Research: Replication of Significant Results—Modeling the Effects of p -Hacking

Most sciences search for lawful data patterns or regularities to serve as the building blocks of theories (e.g., Bunge, 1967; Carnap, 1995; Popper, 2002). Generally, such data patterns must not be singular findings (i.e., chance findings) but instead be replicable by other researchers under similar conditions in order to be scientifically significant (Popper, 2002, p. 23). With this fundamental scientific premise as background, it is understandable that many researchers have become concerned that numerous published results cannot be replicated in independent studies and hence appear to represent chance findings or so-called false-positive results (Baker & Penny, 2016; Ioannidis, 2005b; Pashler & Harris, 2012; Simmons, Nelson, & Simonsohn, 2011; Zwaan, Etz, Lucas, & Donnellan, 2018).¹ For example, only less than 30% of results in social psychology and about 50% in cognitive psychology appear to be reproducible (Open Science Collaboration, 2015). Similarly, the replication rate of 21 systematically selected experimental studies in the social sciences published between 2010 and 2015 in *Nature* and *Science* was estimated to be only about 62% (Camerer et al., 2018). In addition, low replication rates have been reported in medical research (Begley & Ellis, 2012; Ioannidis, 2005a; Prinz, Schlange, & Asadullah, 2011). For example, only 6 out of 53 landmark studies in oncology were reproducible (Begley & Ellis, 2012). Moreover, of 49 highly cited clinical studies, only 20 (44%) could be replicated (Ioannidis, 2005a).

Possible Causes of Low Replication Rates

Understanding the causes of these surprisingly low replication rates has received much attention (e.g., Button & Munafò, 2017; Pashler & Harris, 2012; Schmidt & Oh, 2016), and various possibilities have been discussed. First, scientists may fabricate data to support their hypotheses. However, surveys indicate that this is probably not a major cause because the prevalence of scientific fraud is low—probably smaller than 2%

¹ According to Google Scholar the numbers of documents that mention the phrase “replication crisis” are 84, 283, 568, 979, and 1540 in 2015, 2016, 2017, 2018 and 2019, respectively (as of March 13th, 2020).

(see Fanelli, 2009; Goodstein, 2016; Stroebe, Postmes, & Spears, 2012). Second, in an article with 72 authors (Benjamin et al., 2018), it has recently been argued that the traditional α level of 5% is too large and thus produces too many false positives. These authors suggested changing the critical α level to 0.5%, because this “would immediately improve the reproducibility of scientific research in many fields” (p. 6). Although this change would decrease the false positive rate, however, it would also *increase* the proportion of false negatives unless there were substantial increases in sample size (Fiedler, Kutzner, & Krueger, 2012). Third, another important factor seems to be the typically low statistical power in psychological research (Button & Munafò, 2017; Stanley, Carter, & Doucouliagos, 2018). In a large survey of over 12,000 effect sizes, Stanley et al. (2018) reported that median power was about 36% and that only 8% of all studies had an adequate power of about 80%. Even lower median power of about 21% has been reported for studies in the neurosciences (Button et al., 2013). Low power within a research area reduces replicability for purely statistical reasons, because it reduces the ratio of replicable true positives to non-replicable false positives.

Fourth, the percentage or “base rate” π of true effects within a research area strongly influences the replication rate (Miller, 2009; Miller & Ulrich, 2016; Wilson & Wixted, 2018). When π is small, the relative proportion of false positives within a given research domain will be high (Ioannidis, 2005b; Oberauer & Lewandowsky, 2019), and thus the replication rate will be low. This is easily seen: for $\pi = 0$ the relative proportion of false positives is 100%. In contrast, for $\pi = 1$, no false positives can occur so this proportion is zero. Consequently, replication rates must be higher when the base rate is relatively high than when it is low. For example, Wilson and Wixted (2018) have argued that the fields of cognitive and social psychology differ in the base rate of real effects are investigated, which they call the “prior odds.” On the basis of the results obtained by the Open Science Collaboration (2015), they estimated base rates of $\pi = 0.20$ for cognitive psychology and $\pi = 0.09$ for social psychology, and these estimates are consistent with the finding that the replication rate is lower for social than cognitive psychology. Alternative analyses of replication rates and prediction markets

also suggest similarly low base rates of about 10% (Dreber et al., 2015; Johnson, Payne, Wang, Asher, & Mandal, 2017; Miller & Ulrich, 2016). More generally, it is reasonable to assume that base rates differ between discovery-oriented research and theory-testing research (Lewandowsky & Oberauer, 2020; Oberauer & Lewandowsky, 2019).

Finally, although a certain percentage of false-positive results is an unavoidable by-product of null hypothesis testing, many have argued that replication rates are low because questionable research practices (QRPs) used by scientists chasing after statistically significant results produce an excess of false-positive results beyond the usual nominal significance level of 5% (Ioannidis & Trikalinos, 2007; John, Loewenstein, & Prelec, 2012; Simmons et al., 2011). Bias toward publication of significant results provides a strong incentive to use QRPs (Bakker, Van Dijk, & Wicherts, 2012), especially when competing for academic promotion (Asendorpf et al., 2013) or grant funding (Lilienfeld, 2017). A survey conducted by John et al. (2012) identified several such practices, and the most frequent ones can be grouped into four categories (a) A researcher may capitalize on chance by taking advantage of *multiple studies* to obtain a significant result. For example, the researcher may conduct several similar experiments until one finally yields the hoped-for significant result, and then the researcher only reports the results of the one study that ‘worked’, putting negative results into the file drawer (Rosenthal, 1979). There is convincing evidence that researchers conduct several studies to examine a hypothesis but only report those studies that yielded confirming results (Francis, 2014; Francis, Tanzman, & Matthews, 2014). (b) A researcher may measure multiple dependent measures and report only those that yield significant results. For example, a neuroscientist could record brain activity in hundreds of distinct brain areas and report the results only for those that were sensitive to a specific experimental manipulation (Vul, Harris, Winkielman, & Pashler, 2009). With 10 moderately correlated dependent measures (i.e., $\rho = 0.2$) and one-tailed tests, for example, this strategy of *multiple testing* raises the rate of false positives from 5% to 34%. (c) A researcher may monitor data collection, repeatedly testing for significant results, and stop data collection when a significant result is attained. This strategy of

data peeking can easily raise the rate of false positives up to 20% (Simmons et al., 2011). (d) Finally, *selective outlier removal* can also turn a nonsignificant result into a significant one (Ulrich & Miller, 1994). For example, if an initial analysis produces nonsignificant results, a researcher may try different criteria for excluding outliers in the hope of getting significant results after the data have been ‘cleaned’.

With all four of these QRPs as well as other ones, the researcher exploits the degrees of freedom present in the research process to achieve a statistically significant result — a practice that has been referred to as “*p*-hacking” (Simonsohn, Nelson, & Simmons, 2014a). This clearly inflates the rate of false positives, which would intuitively be expected to decrease replicability. What has received considerably less attention, however, is that *p*-hacking also increases the statistical power for detecting true effects, as noted recently by Witt (2019). Since increasing power also increases replication rates, the influence of QRPs on power tends to counteract its influence on Type 1 error rate with respect to overall replicability. A quantitative model is therefore needed to assess the size of *p*-hacking’s overall effect on replicability.

In the present study, we develop a quantitative model of replication rate that simultaneously takes into account α , power, the base rate of true effects, and *p*-hacking. This model allows us to assess the relative contributions of these factors to the replication rate, with a focus on the influence of QRPs. In contrast, the combined effects of *p*-hacking on Type 1 error rate and power have not previously been modelled at all, and previous studies have generally considered the effects of these factors on replicability one at a time (e.g., α , power, base rate), making it difficult to see their relative contributions. Knowing the relative contributions of these different factors seems important for a better understanding of why the observed replication rate is so low, and this knowledge would be useful in guiding efforts to improve the situation. Because the various different *p*-hacking strategies reviewed above may have different impacts on the replication rate, we conducted separate analyses for each strategy.

Statistical Analysis of the Replication Scenario

The analyses in this manuscript address replication scenarios in which researchers try to replicate studies that reported a statistically significant positive outcome. Figure 1 depicts these scenarios together with all statistically relevant parameters that must be taken into account when computing the rate RR of replicating significant results (Miller, 2009; Miller & Ulrich, 2016; Miller & Schwarz, 2011). First, each original study tests either a true effect (i.e., H_1 is true) or a null effect (i.e., H_0 is true), with base rate probabilities π and $1 - \pi$, respectively, and these probabilities may vary across research fields (Wilson & Wixted, 2018).² If the original study tests a true effect, its statistical power is $1 - \beta_1$ and thus the Type 2 error probability is equal to β_1 . Thus, the compound probability of examining a true effect and rejecting the null hypothesis is $\pi \cdot (1 - \beta_1)$; this outcome is called a “true positive.” In contrast, if the original study tests a null effect, its Type 1 error probability is α_1 . Thus, the probability of testing a null effect and falsely rejecting H_0 is $(1 - \pi) \cdot \alpha_1$; this outcome reflects a “false positive.” Note that, in keeping with accepted procedures for null hypothesis testing, we categorize studies as rejecting the null hypothesis or not based on an all-or-none comparison of computed p -values relative to an α level cutoff. Such a discrete categorization is, for example, how most journals currently evaluate statistical results in publication decisions and how replication success or failure has mainly been operationalized in empirical studies of replication rates (Open Science Collaboration, 2015; Camerer et al., 2018) and consequently conveyed by the news media to the public.

Only true positives and false positives enter into replication projects. The statistical power $1 - \beta_2$ and Type 1 error probability α_2 of the replication studies might differ from those of the original study, especially because replication studies are usually designed to have much higher power than the original studies. Thus, the compound probability of examining a true effect that yields a significant effect in the original and in the replication study is $\pi \cdot (1 - \beta_1) \cdot (1 - \beta_2)$, whereas the compound probability of

² The base rate is also sometimes called the “pre-study probability” (Ioannidis, 2005b) or “prior odds” (Wilson & Wixted, 2018).

examining a null effect and finding significance in both the initial study and the replication study is $(1 - \pi) \cdot \alpha_1 \cdot \alpha_2$. From the above compound probabilities, the rate of replication of initially significant results, RR , can be computed as

$$RR = \frac{\pi \cdot (1 - \beta_1) \cdot (1 - \beta_2) + (1 - \pi) \cdot \alpha_1 \cdot \alpha_2}{\pi \cdot (1 - \beta_1) + (1 - \pi) \cdot \alpha_1}. \quad (1)$$

If p -hacking is performed in the original study, this would increase the Type 1 error rate above the nominal significance level α_1 (usually 5%) to, for example, 10% or even higher. Thus, when a researcher examines a null effect, p -hacking increases the proportion of false positives. The extent of this increase depends on the details of the p -hacking strategy that is used, as we examine in detail below for different strategies. However, and crucially for the analyses that will follow, when a true effect is present, p -hacking also increases the nominal power $1 - \beta_1$, for example, from 0.20 to 0.40. With respect to the overall replication rate RR , this increase in power tends to compensate for the increased Type 1 error probability, making it difficult to determine intuitively how p -hacking would affect the replication rate RR . Fortunately, however, Equation 1 can be used to assess this issue quantitatively.

Besides assessing the effect of these factors on replicability, we will also report computations of the rate of false positives FPR , which is the proportion of false positive results among all significant results within a research area (sometimes also called *false discovery rate* or *false positive report probability*)³

$$FPR = \frac{(1 - \pi) \cdot \alpha_1}{(1 - \pi) \cdot \alpha_1 + \pi \cdot (1 - \beta_1)}. \quad (2)$$

In discussions about replicability—particularly replicability of published research findings—researchers often focus on this proportion (Button et al., 2013; Pashler & Harris, 2012) under the assumption that true positives are replicable but false positives are not. Therefore, it seems useful to include this rate in the analyses.

In the following, we model each of the four common p -hacking strategies that were described above. For each strategy, the inflated Type 1 error probability and the

³ Some researchers instead use the *positive predictive value* (Ioannidis, 2005b), which is $PPV = 1 - FPR$.

statistical power can be computed. These values are then inserted into Equation 1, which allows one to evaluate the effects of base rates and p -hacking on the replication rate, for both true and null effects. In addition, we examined the effects on RR of different levels of α_1 and statistical power, because—as mentioned above—several researchers have recently suggested lowering the α level or increasing power in order to increase the replicability of scientific results (Benjamin et al., 2018; Button et al., 2013). This allows one to judge how these suggested measures would combat low replicability and to compare their effects with those of p -hacking and base rate.

Multiple Studies

It has been often suspected that researchers tend to selectively report studies that yield positive results, that is, results that are in accordance with the researcher’s hypothesis (e.g., John et al., 2012; Rosenthal, 1979; Simmons et al., 2011; Zwaan et al., 2018). As noted earlier, this tendency will increase the number of reported false positives if researchers publish only the significant outcomes. This section models this p -hacking strategy and examines how it would influence the replication rate.

As a specific example, suppose that a researcher runs a series of experiments, each of which uses a slight variation of the same basic paradigm. This researcher terminates the series when a significant result emerges in support of the researcher’s hypothesis, and in this case the researcher tries to publish that result. However, if no significant result is obtained after conducting k experiments, the researcher abandons the project and concludes that the hypothesis is false. Thus, this researcher has k studies providing opportunities to test the hypothesis, and it would be misleading about the overall α level to publish only the significant outcome but not mention the non-significant attempts (Francis, 2014).

To model this scenario more concretely, assume that the researcher computes a z -value for the outcome of each experiment and considers the outcome to be statistically significant if any z -value exceeds a pre-specified criterion c (e.g., the critical z value of

216 1.96). In general, the probability of rejecting H_0 can be computed for $k \geq 1$ with

$$\Pr(\text{Reject } H_0) = 1 - \prod_{i=1}^k \Pr(Z_i \leq c), \quad (3)$$

217 because the outcomes of the k experiments are statistically independent if a new sample
218 is recruited each time.

219 Figure 2 depicts the probability of rejecting H_0 for two-sample tests as a function
220 of $k \in (1, 2, 3, 4, 5)$, $\alpha \in (5\%, 0.5\%)$, and effect size $d \in (0.0, 0.2, 0.5, 0.8)$. (A parallel
221 analysis for one-sample tests and detailed descriptions of both tests are contained in the
222 Supplementary Materials.) In these examples, the group size is assumed to be $n = 20$, a
223 value that is typical for psychological research (Marszalek, Barber, Kohlhart, & Holmes,
224 2011, Table 3).⁴ The lines for $d = 0$ depict the effective Type 1 error probability. Of
225 course, this probability is equal to α for $k = 1$, but it increases with k because of the
226 greater number of opportunities for getting a significant result by chance when more
227 studies are conducted. This increased Type 1 error probability is problematic because it
228 tends to decrease replication rates (Benjamin et al., 2018). In the worst of these cases,
229 the inflated Type 1 error rate attains a value of about 0.23 with $\alpha = 5\%$ and $k = 5$. As
230 one expects, decreasing the nominal α level from 5% to 0.5% substantially diminishes
231 the Type 1 error probability and thus correspondingly diminishes the probability of
232 obtaining a false positive (Benjamin et al., 2018). Even for $k = 5$, the Type 1 error rate
233 would only be about 0.025 with this smaller nominal α level. It must be stressed,
234 however, that a larger sample would be required for $\alpha = 0.5\%$ than for $\alpha = 5\%$ to
235 achieve the same level of statistical power for both cases (Benjamin et al., 2018).

236 The lines for $d > 0$ reveal the statistical power to reject H_0 when it is false. When
237 researchers follow good scientific practice, the statistical power associated with each
238 value of d can be seen at $k = 1$. As is well known, power generally increases with d , and
239 it is larger with $\alpha = 5\%$ than with $\alpha = 0.5\%$. For the present purposes, however, the
240 most important aspect of the $d > 0$ lines is the strong increase of statistical power with
241 k , which can be seen in both panels, especially when the single-experiment power is well

⁴ The total sample size for a two-sample test is $2 \cdot n$, that is, two times the group size n .

below one. Since replication rates should increase with power (Button et al., 2013; Button & Munafò, 2017), this increase in power will tend to compensate for the increased Type 1 error rate with respect to the overall influence of multiple studies p -hacking on replication rate. It is therefore necessary to use a quantitative model to assess the net effect of the multiple studies practice on the replication rate.

Using the above probabilities of rejecting H_0 , the proportion of false positives associated with this p -hacking scenario can be computed from Equation 2. Figure 3 highlights the false positive rate as a function of d , α , and k . Dashed lines show the rates for researchers engaged in p -hacking. For comparison, the solid lines depict the rates for researchers who follow good scientific practice by just running a single experiment and reaching a conclusion based on its outcome (i.e., $k = 1$). The rates for these researchers were also computed with Equation 2 by inserting the nominal value of α for α_1 and the single-experiment power for $1 - \beta_1$.

Several effects can be observed in Figure 3: (a) As one expects, the false positive rate decreases from one to zero with increasing π , because the proportion of true effects among all significant effects becomes larger when π increases (e.g., Ioannidis, 2005b; Wilson & Wixted, 2018). (b) Not surprisingly, the false positive rate becomes smaller when power increases (Ioannidis, 2005b). (c) Most interestingly and surprisingly, the increase in false positives produced by p -hacking is more pronounced with larger d , where statistical power is higher. This is presumably because p -hacking cannot increase statistical power much when it is already high (i.e., when d is large), so there is little power increase to compensate for the increased Type 1 error rate. Nevertheless, the effect of p -hacking is far from dramatic for $k = 2$, although it can be quite prominent for larger values of k , especially with small base rates.

Figure 4 depicts replication rates computed using the same parameters as in the previous figures. Four features of these computations are especially noticeable. (a) Successful replication depends strongly on the base rate. As one might expect, all rates converge to the statistical power $1 - \beta_2 = 0.90$ of the replication study, because when all significant effects are real, the replication rate simply reflects the statistical power of the

replication study, whether p -hacking was involved in the first study or not. (b) The effect of p -hacking is modest for high base rates, for the smaller α level, and interestingly also for smaller effect sizes and hence for low statistical power. (c) As emphasized by Benjamin et al. (2018), the replication rate is considerably larger for $\alpha = 0.5\%$ than for $\alpha = 5\%$, especially for small base rates.

In summary, the above analysis casts doubt on the idea that this p -hacking strategy is a major contributor to low replicability, even though it seems to be one of the most frequent QRPs (e.g., John et al., 2012). Instead, it seems that using this strategy would have little effect on replicability except in research scenarios where true effects were rare but there was high power to detect them when they were present. The strongest trends suggest that a low base rate of true effects is the major cause of low replicability (Wilson & Wixted, 2018), since changes in base rate can cause replication rates to range across nearly the full 0–1 range.

Failing to Report All Dependent Measures

Failing to report all of a study’s dependent measures seems to be another common QRP (Fiedler & Schwarz, 2016; John et al., 2012). In this section, we analyze how this practice would affect the replication rate of statistically significant results. In order to model this scenario, we assume that a researcher conducts a study to test a certain hypothesis using a control and experimental condition. After data collection, however, the researcher only reports the outcomes of those tests that surpass the statistical significance threshold and thereby confirm the proposed hypothesis. As examples, multiple dependent measures are usually measured and statistically evaluated in neurosciences and medical research, raising concerns about Type 1 error rates (e.g., Hutton & Williamson, 2002; Vul et al., 2009).

We again employ z -tests to model this scenario. Let $Z_1, \dots, Z_k$ be the outcomes for all k dependent measures of a single study, with each Z -value representing the result of the control/experimental comparison for a single measure. Therefore, the probability

of obtaining at least one significant result is equal to

$$\Pr(\text{Reject } H_0) = 1 - P(Z_1 \leq c, \dots, Z_k \leq c)$$

with c being the critical cutoff value (see Supplementary Materials for computational details). Because such measures are usually correlated across participants, our model incorporates correlations among the Z_i values.

Figure 5 illustrates the effects on Type 1 error probability (i.e., lines with $d = 0$) and statistical power (i.e., lines with $d > 0$) associated with this type of p -hacking. For this illustration, the pair-wise correlations of the different dependent measures were set to 0.2 and the number of observations (per group) was set to 20, which are seemingly typical values in psychological research (Bosco, Aguinis, Singh, Field, & Pierce, 2015; Marszalek et al., 2011). As expected, both the Type 1 error rate and power increase with the number of dependent measures, approximately as was found with multiple studies.

Figures 6 and 7 show the rates of false positives and of replications, respectively, resulting from this type of p -hacking, and these rates are quite similar to those seen with the multiple studies scenario (see Figs. 3 and 4). In particular, both false positive rates and replication rates show strong expected effects of base rate and α level, as well as a clear influence of effect size, d . The effects of p -hacking are again rather modest, however, especially when the effect size is small (i.e., $d = 0.2$) so that increased power is especially helpful.⁵

It should be noted that the extent of both Type 1 error rate inflation and power enhancement depend on the correlations among the different dependent measures. A

⁵ It is instructive to note that the inflation of the Type I error rate cannot exceed a certain upper bound (Hochberg & Tamhane, 1987). Specifically, it has been shown that this rate for k correlated z -statistics can never be larger than $1 - (1 - \alpha)^k$ for a given nominal α level. The current results are consistent with this theoretical result. For example, for $k = 5$ and $\alpha = 0.05$, the above calculations resulted in an inflated error rate of 0.205, which is smaller than the value of $1 - (1 - 0.05)^5 = 0.226$ given by this bound. The upper bound is only attained when the z -statistics are uncorrelated, as they are in the multiple studies scenario.

correlation of zero would yield results identical to those of the scenario with multiple studies, because in this case the outcomes for multiple dependent measures are independent just like the outcomes of multiple independent studies. In contrast, larger correlations (e.g., larger than the 0.2 used in Figs. 6 and 7) weaken the effects of this *p*-hacking strategy, because the measures become increasingly redundant as the intercorrelations increase, and this lowers the possibility of capitalizing on chance. In other words, increasing the intercorrelations would lower the Type 1 error rate inflation and decrease power enhancement. Moreover, increased intercorrelations would decrease the false positive rate and increase the replication rate, that is, moving the dashed lines in Figures 6 and 7) toward the solid reference lines (see the Supplementary Material for a parallel analysis with intercorrelations of 0.8).

Data Peeking

Another frequently-used QRP is data peeking (Fiedler & Schwarz, 2016; John et al., 2012). This practice occurs when a researcher collects additional data after noting that the results of initially collected data have not yielded statistical significance. A researcher may even peek at the results several times and increase the sample with additional observations each time a nonsignificant result is obtained. Data collection is finally terminated only if the study yields no significant result after k peeks. It is known that this practice increases the Type 1 error rate (Armitage, McPherson, & Rowe, 1969; Francis, 2012; McCarroll, Crays, & Dunlap, 1992; Simmons et al., 2011; Strube, 2006). For example, Monte-Carlo simulations by Simmons et al. (2011) revealed that this strategy can increase the error rate up to 14.3% with a first peek at $n = 10$ and four subsequent peeks (each time increasing the sample by 10 observations). However, this practice increases not only the Type 1 error rate but also the effective statistical power to reject a false H_0 (Strube, 2006), so a quantitative analysis is needed to determine its effect on replication rate.

An analysis similar to that of the preceding sections was conducted to examine how data peeking affects Type 1 error rates, power levels, false positive rates, and

replication rates. The Supplementary Material contains the computational details of this analysis, which follows an extension of Armitage et al. (1969)’s procedure described there. In brief, the probability of rejecting H_0 with a maximum of k peeks at successive sample sizes $n_1 < n_2 < \dots < n_k$ is again given by the multivariate normal distribution for z -tests

$$\Pr(\text{Reject } H_0) = 1 - P(Z_1 \leq c, \dots, Z_k \leq c).$$

The correlations among the different Z_i values are determined by the amount of shared data used in computing them (e.g., all observations used in computing Z_1 are also included in the computation of Z_2).

Figure 8 depicts the probability of rejecting H_0 for various effect sizes and two-sample tests. The abscissa represents the maximal number of peeks k at which a researcher would give up recruiting additional participants. For this example, it is assumed that data peeking occurs after 10, 15, 20, 25, or 30 observations. Thus, a researcher with a maximum of $k = 2$ peeks will check statistical significance the first time at $n = 10$ and if the first peek does not reveal a significant result, the data will be examined a second and final time at $n = 15$. For $k = 3$, data will be examined a first time at $n = 10$ and—depending on the outcome of the first peek—a second time at $n = 15$; if the second peek also does not reveal a significant result, a final peek occurs at $n = 20$. Figure 8 shows quantitatively how the probability of rejecting H_0 increases with the maximum number of peeks. In particular, the increase can be quite strong in situations with only moderate power (e.g., $\alpha = 0.5\%$ and $d = 0.8$) due to the extra chances of detecting the true effect. In contrast to the multiple dependent measures with intercorrelations of 0.2 as discussed in the previous section, the Type 1 error rate inflation is smaller in the present case, because $Z_1, \dots, Z_k$ are more strongly correlated under this scenario (cf. the correlation matrix in the Supplementary Material).

Given the probabilities of rejecting H_0 , the replication rate and false positive rate are again computed using Equations 1 and 2, respectively. Since the results with respect to the false positive rate and the replication rate are quite similar to those of the preceding scenarios, we only present the replication rate in Figure 9. We compare

this p -hacking strategy with researchers who conform to good scientific practice and thus examine the data only once at a preplanned n . In order to enable a conservative comparison with p -hackers, we used a preplanned n corresponding to the maximum number of observations a p -hacker would try when using the indicated number of peeks (i.e., this preplanned size would be $n = 15$ for the comparison with $k = 2$, $n = 20$ for the comparison with $k = 3$, etc.). As can be seen, the pattern of results is quite comparable to the previous scenarios. Overall, the data peeking strategy again seems to have little effect on replication rate except in research scenarios where true effects are infrequent and there is high power to detect them when they do occur, just like multiple studies p -hacking.

Selective Outlier Removal

Another QRP identified by John et al. (2012) is to analyze the same overall data set several times, each time excluding “outlier” data points identified by different criteria. The researcher may be tempted to conclude that a real effect has been found if any analysis yields a significant result, but this practice inflates the Type 1 error rate, because each of the analyses provides a further opportunity to obtain a significant result by chance. On the positive side, though, this practice again increases power, because each of the analyses also provides a further opportunity for detecting a real effect.

Because the effects of this type of p -hacking are not computable, we conducted Monte-Carlo simulations to see how multiple attempts at outlier removal would affect the Type 1 error rate, power, rate of false positives, and replication rate. Specifically, we examined the common practice of excluding scores more than a given number of standard deviations from the sample mean. We simulated researchers who carried out a sequence of at most five separate analyses on a single data set. The first analysis included all scores, and the next three analyses included only the scores within 3, 2.5, and 2 standard deviations of the mean, respectively. The fifth analysis used a nonparametric test, which could potentially be used as an analysis in an attempt to minimize the influence of outliers even further.

We simulated experiments for two-sample tests⁶ with a sample size of $n = 20$ per group using standard normally distributed scores and true effect sizes of $d = 0, 0.2, 0.5$, and 0.8 . Researchers were modelled as using either $\alpha = 0.5\%$ or 5% , one-tailed. The nonparametric test was the Mann-Whitney U test, and this test was used only if none of the previous analyses had produced significant results. We simulated 10,000 experiments with outliers by adding a random noise value to 5% of the data values, where these noise values came from a normal distribution with $\mu = 0$ and $\sigma = 10$.

Figure 10 shows the probabilities of rejecting H_0 . As with the other p -hacking methods, this probability increases with the number of analyses conducted, increasing the probability of a Type 1 error when $d = 0$ and increasing power when $d > 0$. In fact, the extra analyses produce relatively large increases in rejection probability when a true effect is present (i.e., $d > 0$), consistent with the idea that the power of the analysis is improved by removing outliers. Figures 11 and 12 show the false positive and replication rates. Interestingly, both measures indicate slightly better results (i.e., lower false positive rates and higher replication rates) when researchers perform multiple analyses to remove the effects of possible outliers than when they do not.

Of course the story is different when no outliers are present in the data set. Making multiple attempts to remove outliers in this case would actually increase the false positive rate and lower the replication rate (see Supplementary Material for a parallel simulation with no outliers). In fact, extreme data points in data sets without outliers appear to be especially diagnostic for testing the equality of locations between populations, as the Tukey pocket test demonstrates (Tukey, 1959), so throwing away extreme observations that are not outliers reduces the information in the data set.

General Discussion

Reproducibility of research results is indispensable for scientific progress (Popper, 2002) and promotes public trust in science (Saltelli & Funtowicz, 2017). Therefore,

⁶ A parallel analysis was performed for one-sample tests. Because the results were extremely similar to the results presented in this paper, we will not report the results of these additional simulations.

replication failures are problematic not only scientifically but also from a public relations viewpoint. The ongoing crisis involving perceived reproducibility concerns virtually all sciences, including, for example, psychology (Open Science Collaboration, 2015), quantitative social science (Freese & Peterson, 2017), economics (Brown & Wood, 2018), neuroscience (Button et al., 2013), medicine and pharmacology (Begley & Ellis, 2012), as well as traditional sciences like biology, chemistry, and physics (Baker & Penny, 2016).

This unfortunate situation naturally prompts questions about how replication rates can be improved, and numerous articles and hundreds of blog posts have been written to address this question. Several measures have been advocated as ways to raise reproducibility, such as (a) preregistration of studies (Nosek, Ebersole, DeHaven, & Mellor, 2018), (b) increasing the transparency of research by making data and research materials publicly available (e.g., stimulus materials and computer codes; Nosek et al., 2015), (c) reducing the α level for classifying results as statistically significant (Benjamin et al., 2018), (d) increasing statistical power (Button & Munafò, 2017), (e) improving statistical training (Asendorpf et al., 2013), (f) replacing traditional statistical inference with Bayesian approaches (Etz & Vandekerckhove, 2016), and (h) even overhauling the standard methodological approaches (Feldman Barrett, 2020). The variety of these proposed measures demonstrates that the failures to replicate research can result from a multitude of causes (e.g., fraud, QRPs, *p*-hacking, erroneous computer codes, misunderstanding of statistical concepts) that may come into play at various steps along the “entire analysis pipe line” (Leek & Peng, 2015).

The present article focused on the statistical consequences of QRPs with respect to replication rate. Equation 1 in the Introduction shows that—even when the researchers in a field use best practices—the rate of successful replications depends on the base rate π of true effects being tested within the field, as well as the α level and statistical power of both the original study and the replication study. Of course, the statistical power in turn is determined by effect size, sample size, and the chosen α level. This nominal α level becomes inflated when researchers adopt QRPs, however,

which increases the rate of false positives within their field (Simmons et al., 2011). It is quite understandable that many researchers are concerned about this type of inflation (e.g., Armitage et al., 1969; Button et al., 2013; Pashler & Harris, 2012), because it could threaten the replicability of results. However, little attention has been paid to the fact that p -hacking can also increase the statistical power of detecting true effects (Witt, 2019). Since increases in Type 1 error rate and power have counteracting effects on replicability, a quantitative analysis is needed to assess p -hacking's likely contribution to currently low replication rates (Open Science Collaboration, 2015).

To the best of our knowledge, the impacts of the various statistical factors affecting replication rate (i.e., α , power, π , p -hacking) have so far been examined almost exclusively in isolation, which does not allow a complete assessment of their mutual influence. A problem with this approach is that considering each of these effects in isolation may lead to incompatible suggestions. As mentioned earlier, for example, some researchers have recently suggested lowering the traditional α -level cutoff from 5% to 0.5% in order to decrease the false positive rate and thus to increase the probability of successful replications (Benjamin et al., 2018). Others, however, have suggested increasing statistical power (e.g., Button et al., 2013; Button & Munafò, 2017). Since reducing α tends to decrease power, however, these two suggestions can only be implemented simultaneously if sample sizes are greatly increased, which would generally be relatively costly (Miller & Ulrich, 2019). In order to develop a better quantitative picture of the different influences on replicability, we modelled several apparently-frequent p -hacking strategies (i.e., multiple studies, testing multiple measures, data peeking, and selective outlier removal) to examine their impacts on replication rate.

With respect to p -hacking, our quantitative analyses suggest that its effects on replicability are unlikely to be massive. Consistent with many previous cautionary articles (e.g., Simmons et al., 2011), p -hacking does clearly inflate the effective Type 1 error rate, which tends to reduce replicability, but the corresponding increase in power compensates for this inflation to a fair degree. Compared to the strong effect of the

base rate on replicability, the reduction in replication rate caused by p -hacking appears much smaller for all four of the p -hacking strategies that we assessed in this article. Unsurprisingly, this impact is larger when p -hacking is relatively extensive (i.e., $k = 5$) rather than mild (i.e., $k = 2$). Moreover, p -hacking affects the replication rate most when the base rate is small. This makes sense, because p -hacking is harmful primarily when H_0 is true, which is more common with small base rates. The net influence of p -hacking on replicability appears to be smallest with small effect sizes, which is presumably the situation where p -hacking is most likely to be used. With small effects, the power increases associated with p -hacking are especially helpful for replicability. Interestingly, however, not all strategies have a negative impact on the replication rate. For example, selective outlier removal may in fact even increase it. Finally and somewhat surprisingly, p -hacking tends to have a smaller effect on the replication rate when the nominal α level is 0.5% rather than 5%.

Of course, these conclusions are restricted to the limited extent of p -hacking that we examined (i.e., $k = 2, \dots, 5$), and more extensive p -hacking—or combining multiple p -hacking strategies—would presumably have larger effects on replicability. Nonetheless, we think that five is a reasonable upper bound on the number of attempts to obtain statistical significance that most researchers would make using the p -hacking methods discussed here and that researchers would rarely use multiple p -hacking methods in combination. In addition, when estimating the actual effect of p -hacking on observed replication rates (e.g., Open Science Collaboration, 2015), it is important to note that the effects shown in our figures are *upper bounds* that would only be approached if nearly all researchers employed these p -hacking methods. If only 10% of researchers use these methods, then the overall effects on empirical replication rates would be only 10% as large as those suggested by our model. Even the highest estimates of the prevalence of QRPs are only approximately 50% (John et al., 2012), and these may be serious overestimates (Fiedler & Schwarz, 2016).

Another limitation concerning our conclusions is that our list of p -hacking strategies was not exhaustive. For example, we did not examine the possibility that

researchers might try several covariates until a significant result is obtained (e.g., Simonsohn, Nelson, & Simmons, 2014b). As another example, suppose a researcher conducts a multi-factor analysis of variance (ANOVA) that invites the examination of multiple main effects and interactions, any one of which might be cherry picked as a “finding” in the absence of a specific *a priori* hypothesis. For instance, a three-factorial ANOVA allows the examination of seven potential effects (i.e., three main effects and four interactions). Assuming that all seven sources are independent, the probability of at least one significant result when H_0 holds in all cases is $1 - (1 - \alpha)^7$ —about 30% with $\alpha = 0.05$ —which would simply emulate the multiple studies scenario that we analysed in this article.⁷ Thus, analyses similar to the present ones would be needed to analyze the consequences of these other strategies, but it would be surprising if the results were drastically different.

We supplemented the analyses reported in this manuscript by two further analyses (see Supplementary Material), each of which approached the replication issue from a different angle. The first supplementary analysis assessed the effect of *p*-hacking on power while controlling for the overall Type 1 error rate. The outcome of this analysis demonstrated that some *p*-hacking strategies can actually produce higher statistical power than good practice at each level of Type 1 error. This superiority can be explained by the fact that *p*-hacking sometimes involves the collection of additional data (e.g., as with data peeking or measuring additional variables), and in these cases the additional data can cause statistical power to increase faster than the Type 1 error rate. The second supplementary analysis compared the overall research payoff associated with good practice versus data peeking using the payoff model of Miller and Ulrich (2016). This analysis replicated several features of our previous studies (see also Miller & Ulrich, 2019) concerning the effects of α level, base rate, and sample size on total research

⁷ This calculation is only valid for a completely within-subjects design in which different error terms apply to each source of variance and thus the F values for the various effects are uncorrelated. For a between-subjects design, however, the same error term is used for all F tests and the resulting F -values are thus positively correlated (Schwarz, 1993). In that case, $1 - (1 - \alpha)^7$ would provide an upper bound for the inflated α level.

541 payoff. With respect to *p*-hacking, however, the expected total payoff can actually be
542 larger with data peeking than with good practice, evidently because data peeking tends
543 to make more efficient use of limited sample sizes when true effects are common.

544 If *p*-hacking is not a major contributor to low replicability, then what is? In
545 keeping with previous analyses (Dreber et al., 2015; Johnson et al., 2017; Miller, 2009;
546 Miller & Ulrich, 2016; Wilson & Wixted, 2018), our results suggest that low base rates
547 of true effects—not too-large α levels, too-low power, or *p*-hacking—are most likely to
548 be the major causes of poor replicability, so researchers concerned about replicability
549 should pay special attention to the issue of base rates. Clearly, low base rates can lead
550 to disappointingly low replication rates even in the absence of *p*-hacking (e.g., Figs. 4, 7,
551 9, and 12, “good practice”). It follows from our analyses that research fields with
552 inherently low base rates (e.g., biomedical research) simply cannot improve their
553 replication rates much by focusing exclusively on methodological issues.

554 As was mentioned in the introduction, there is increasing evidence—based on
555 several different ways of estimating the base rates of true effects (e.g., Dreber et al.,
556 2015; Johnson et al., 2017; Miller & Ulrich, 2016, 2019; Wilson & Wixted, 2018)—that
557 base rates are low in many fields and that they vary across research fields in a manner
558 that is consistent with replication rates. For example, Wilson and Wixted (2018) noted
559 the different replication rates of 25% and 50% respectively, for social and cognitive
560 psychology (Open Science Collaboration, 2015), and they noted that the difference
561 could be explained by differing base rates of 9% and 20%, respectively, in the two areas,
562 even in the absence of any *p*-hacking in either field. The lower base rate for social
563 psychology could reflect a greater emphasis on the discovery of surprising results.

564 Looking beyond replication rates, meta-scientists should consider exactly what
565 measure of research productivity they want to optimize. For example, if the goal is to
566 minimize false positives, they should use small α levels and eliminate *p*-hacking. If the
567 goal is to minimize false negatives, however, they should do exactly the opposite. The
568 major problem in statistical decision making is that one cannot maximize all of the
569 desirable goals at the same time. Thus, focusing on only one goal—even that of

maximizing replicability—will not yield an optimal research strategy. Identifying the optimal strategy requires considering all of the goals simultaneously and integrating them into a composite measure of research productivity. For example, one way to do this is to analyze the probabilities and payoffs for a set of possible research outcomes and to identify research parameters maximizing the expected research payoff (Miller & Ulrich, 2016).

Conclusion

We modelled different causes (alpha level, power, base rate of true effects, QRPs) of low replication rates within a general statistical framework. Our analyses indicate that a low rate of true effects—not *p*-hacking—is mainly responsible for low replication rates—a point that is often under-appreciated in current debates about how to improve replicability. Of course, we do not wish to transmit the message that *p*-hacking is tolerable just because it might increase power when a researcher examines a true effect. As has often been discussed previously (Simmons et al., 2011), *p*-hacking should always be avoided because it inflates Type 1 error rates above stated levels and thus undermines scientific progress. Rather, our message is that scientists and others concerned about low replication rates should look beyond *p*-hacking for its primary causes. The current analyses suggest that even massive campaigns against *p*-hacking (e.g., researcher education, pre-registration initiatives) may produce only modest improvements in replicability. To make large changes in this important scientific measure, it will likely be necessary to address other aspects of the scientific culture. Unfortunately, that may not happen if attention and blame are focused too narrowly on *p*-hacking as a major cause of the current problems in this area.

References

- Armitage, P., McPherson, C. K., & Rowe, B. (1969). Repeated Significance Tests on Accumulating Data. *Journal of the Royal Statistical Society: Series A (General)*, 132, 235–244.
- Asendorpf, J. B., Conner, M., De Fruyt, F., De Houwer, J., Denissen, J. J. A., Fiedler, K., ... Wicherts, J. M. (2013). Recommendations for increasing replicability in psychology. *European Journal of Personality*, 27, 108–119.
- Baker, M., & Penny, D. (2016). Is there a reproducibility crisis? *Nature*, 533, 452–454.
- Bakker, M., Van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, 7, 543–554.
- Begley, C. G., & Ellis, L. M. (2012). Raise standards for preclinical cancer research. *Nature*, 483, 531–533.
- Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E. J., Berk, R., ... Johnson, V. E. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2, 6–10.
- Bosco, F. A., Aguinis, H., Singh, K., Field, J. G., & Pierce, C. A. (2015). Correlational effect size benchmarks. *Journal of Applied Psychology*, 100, 431–449.
- Brown, A. N., & Wood, B. D. K. (2018). Which tests not witch hunts: a diagnostic approach for conducting replication research. *Economics: The Open-Access, Open-Assessment E-Journal*, 12, 1–26.
- Bunge, M. (1967). *Scientific Research II: The search for truth* (Vol. 3/II). New York: Springer.
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews. Neuroscience*, 14, 365–376.
- Button, K. S., & Munafò, M. R. (2017). Powering Reproducible Research. In S. O. Lilienfeld & I. D. Waldman (Eds.), *Psychological Science Under Scrutiny: Recent Challenges and Proposed Solutions* (pp. 22–33). Chichester, England:

Wiley Online Library.

Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., ...

Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2, 637–644.

Carnap, R. (1995). *An introduction to the philosophy of science*. Toronto: General Publishing Company.

Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., ... Wachter, K. W. (2015). Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences of the United States of America*, 112, 15343–15347.

Etz, A., & Vandekerckhove, J. (2016). A Bayesian perspective on the reproducibility project: Psychology. *PLoS ONE*, 11, 1–12.

Fanelli, D. (2009). How many scientists fabricate and falsify research? A systematic review and meta-analysis of survey data. *PLoS ONE*, 4(5), e5738.

Feldman Barrett, L. (2020). Forward into the past. *Observer*, 33(3), 5-7.

Fiedler, K., Kutzner, F., & Krueger, J. I. (2012). The Long Way From Error Control to Validity Proper: Problems With a Short-Sighted False-Positive Debate.

Perspectives on Psychological Science, 7(6), 661–669.

Fiedler, K., & Schwarz, N. (2016). Questionable research practices revisited. *Social Psychology and Personality Science*, 7, 45–52.

Francis, G. (2012). Publication bias and the failure of replication in experimental psychology. *Psychonomic Bulletin & Review*, 19, 975–991.

Francis, G. (2014). The frequency of excess success for articles in Psychological Science. *Psychonomic Bulletin & Review*, 21, 1180–1187.

Francis, G., Tanzman, J., & Matthews, W. J. (2014). Excess success for psychology articles in the journal Science. *PLoS ONE*, 9, e114255.

Freese, J., & Peterson, D. (2017). Replication in Social Science. *Annual Review of Sociology*, 43, 147–165.

Goodstein, C. (2016). Scientific Misconduct. *Annual Review of Psychology*, 67,

683–711.

Hochberg, Y., & Tamhane, A. C. (1987). *Multiple Comparison Procedures*. New York: John Wiley & Sons. doi: 10.4135/9781412985178.n2

Hutton, J. L., & Williamson, P. R. (2002). Bias in meta-analysis due to outcome variable selection within studies. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 49, 359–370.

Ioannidis, J. P. A. (2005a). Contradicted and Initially Stronger Effects in Highly Cited Clinical Research. *JAMA: The Journal of the American Medical Association*, 294, 218–228.

Ioannidis, J. P. A. (2005b). Why most published research findings are false. *PLoS Medicine*, 2, e124.

Ioannidis, J. P. A., & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical trials (London, England)*, 4, 245–253.

John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23, 524–532.

Johnson, V. E., Payne, R. D., Wang, T., Asher, A., & Mandal, S. (2017). On the Reproducibility of Psychological Science. *Journal of the American Statistical Association*, 112, 1–10.

Leek, J. T., & Peng, R. D. (2015). P values are just the tip of the iceberg. *Nature*, 520, 612.

Lewandowsky, S., & Oberauer, K. (2020). Low replicability can support robust and efficient science. *Nature Communications*, 11, 1–12.

Lilienfeld, S. O. (2017). Psychology’s Replication Crisis and the Grant Culture: Righting the Ship. *Perspectives on Psychological Science*, 12, 660–664.

Marszalek, J. M., Barber, C., Kohlhart, J., & Holmes, C. B. (2011). Sample size in psychological research over the past 30 years. *Perceptual and Motor Skills*, 112, 331–348.

McCarroll, D., Crays, N., & Dunlap, W. P. (1992). Sequential ANOVAs and Type I

Error Rates. *Educational and Psychological Measurement*, 52, 387–393.

Miller, J. (2009). What is the probability of replicating a statistically significant effect?

Psychonomic Bulletin & Review, 16, 617–640.

Miller, J., & Schwarz, W. (2011). Aggregate and individual replication probability within an explicit model of the research process. *Psychological Methods*, 16, 337–360.

Miller, J., & Ulrich, R. (2016). Optimizing Research Payoff. *Perspectives on Psychological Science*, 11, 664–691.

Miller, J., & Ulrich, R. (2019). The quest for an optimal alpha. *Plos One*, 14(1), e0208631.

Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences of the United States of America*, 115, 2600–2606.

Nosek, B. A., et al. (2015). Promoting an open research culture: Author guidelines for journals could help to promote transparency, openness, and reproducibility. *Science*, 348, 1422–1425.

Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review*, 1596–1618.

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.

Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7, 531–536.

Popper, K. (2002). *The logic of scientific discovery*. London: Routledge Classics.

Prinz, F., Schlange, T., & Asadullah, K. (2011). Believe it or not: How much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery*, 10, 712–713.

Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86, 638–641.

Saltelli, A., & Funtowicz, S. (2017). What is science’s crisis really about? *Futures*, 91,

5–11.

Schmidt, F. L., & Oh, I.-S. (2016). The crisis of confidence in research findings in psychology: Is lack of replication the real problem? Or is it something else? *Archives of Scientific Psychology*, 4, 32–37.

Schwarz, W. (1993). On the degree of dependence of F tests based on a common denominator. *Methodika*, 7, 120–129.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 1359–1366.

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014a). p-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143, 534–547.

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014b). p-curve and effect size: Correcting for publication bias using only significant results. *Psychological Science*, 9, 666–681.

Stanley, T. D., Carter, E. C., & Doucouliagos, H. (2018). What meta-analyses reveal about the replicability of psychological research. *Psychological Bulletin*, 144, 1325–1346.

Stroebe, W., Postmes, T., & Spears, R. (2012, nov). Scientific Misconduct and the Myth of Self-Correction in Science. *Perspectives on Psychological Science*, 7, 670–688.

Strube, M. J. (2006). SNOOP: A program for demonstrating the consequences of premature and repeated null hypothesis testing. *Behavior Research Methods*, 38, 24–27.

Tukey, J. W. (1959). A Quick, Compact, Two-Sample Test to Duckworth's Specifications. *Technometrics*, 1, 31–48.

Ulrich, R., & Miller, J. (1994). Effects of Truncation on Reaction Time Analysis. *Journal for Experimental Psychology: General*, 123, 34–80.

Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science*, 274–290.

- 738 Wilson, B. M., & Wixted, J. T. (2018). The Prior Odds of Testing a True Effect in
739 Cognitive and Social Psychology. *Advances in Methods and Practices in*
740 *Psychological Science*, 1, 186–197.
- 741 Witt, J. K. (2019). Insights into Criteria for Statistical Significance from Signal
742 Detection Analysis. *Meta-Psychology*, 3.
- 743 Zwaan, R. A., Etz, A., Lucas, R. E., & Donnellan, M. B. (2018). Making replication
744 mainstream. *Behavioral and Brain Sciences*, e120, 1–61.

Figure 1. (Probability Tree of the Replication Scenario). The base rates of examining an alternative hypothesis H_1 or a null hypothesis H_0 are π and $1 - \pi$, respectively. The statistical power and the Type 1 error rate of the original study are $1 - \beta_1$ and α_1 . There are four possible outcomes of an original study, with the researcher deciding to reject the null hypothesis (i.e., $'H_1''$) in two outcomes and failing to reject it (i.e., $'H_0''$) in the other two. If H_1 is true, the outcomes associated with these decisions are called true positives and false negatives. By contrast, if H_0 holds, they are called false positives and true negatives. Replication studies replicate original studies that reported a significant positive result. The statistical power and the Type 1 error probability of the replication study are $1 - \beta_2$ and α_2 , respectively. The replication study may either reject H_0 (which denotes a successful replication of the original positive result, R) or fail to reject it (which denotes a failure to replicate the original result, $\bar{R}$).

Figure 2. (Multiple Studies). Each panel depicts the probability of rejecting H_0 as a function of the number of studies k , nominal α level, and effect size d for a two-sample test.

Figure 3. (Multiple Studies). False positive rate (FPR) as a function of base rate π , number of studies k , effect size d , and nominal α level (0.5% vs. 5%). The nominal α level and power of the replication study are $\alpha_2 = 5\%$ and $1 - \beta_2 = 90\%$. All results are based on $n = 20$ per group. Dashed lines give the results for p -hacking whereas solid lines depict the results of researchers who act in accord with good scientific practice. Note that the solid lines are the same in all rows of a single column because these constant reference lines do not depend on k .

Figure 4. (Multiple Studies). Replication rate (RR) as a function of base rate π , number of studies k , effect size d , and nominal α level (0.5% or 5%). The nominal α level and power of the replication study are $\alpha_2 = 5\%$ and $1 - \beta_2 = 90\%$. All results are based on $n = 20$ per group. Dashed lines give the results for p -hacking whereas solid lines depict the results of researchers who act in accord with good scientific practice. Note that the solid lines are the same in all rows of a single column because these constant reference lines do not depend on k .

Figure 5. (Failing to Report all Dependent Measures). Each panel depicts the probability of rejecting H_0 as a function of the number of dependent measures k , nominal α level, and effect size d for a two-sample test.

Figure 6. (Failing to Report all Dependent Measures). False positive rate (FPR) as a function of base rate π , number of studies k , effect size d , and nominal α level (0.5% vs. 5%). The nominal α level and power of the replication study are $\alpha_2 = 5\%$ and $1 - \beta_2 = 90\%$. All results are based on $n = 20$ per group. Dashed lines give the results for *p*-hacking whereas solid lines depict the results of researchers who act in accord with good scientific practice. Note that the solid lines are the same in all rows of a single column because these constant reference lines do not depend on k .

Figure 7. (Failing to Report all Dependent Measures). Replication rate (RR) as a function of base rate π , number of studies k , effect size d , and nominal α level (0.5% vs. 5%). The nominal α level and power of the replication study are $\alpha_2 = 5\%$ and $1 - \beta_2 = 90\%$. All results are based on $n = 20$ per group. Dashed lines give the results for p -hacking whereas solid lines depict the results of researchers who act in accord with good scientific practice. Note that the solid lines are the same in all rows of a single column because these constant reference lines do not depend on k .

Figure 8. (Data Peeking). Each panel depicts the probability of rejecting H_0 as a function of the number of maximal peeks k , nominal α level, and effect size d for a two-sample test.

Figure 9. (Data Peeking). Replication Rate (RR) as a function of base rate π , number of maximal data peeks k , and nominal α level (0.5% vs. 5%). The nominal α level and power of the replication study are $\alpha_2 = 0.05$ and $1 - \beta_2 = 0.90$. Dashed lines give the results for p -hacking whereas solid lines depict the results of researchers who act in accord with good scientific practice. Note that the solid lines are the same in all rows of a single column because these constant reference lines do not depend on k .

Figure 10. (Selective Outlier Removal). Estimated probability of rejecting H_0 as a function of the number k of outlier rejection methods attempted for various effect sizes d , and nominal α level (0.5% or 5%). Probability estimates were based on 10,000 simulated experiments. Simulated data included 5% outliers.

Figure 11. (Selective Outlier Removal). False positive rate (FPR) as a function of the number k of outlier rejection methods attempted, effect size d , and nominal α level (0.5% or 5%). The nominal α level and power of the replication study is $\alpha_2 = 0.05$ and $1 - \beta_2 = 0.90$. Dashed lines gives the results for p -hacking whereas solid lines depict the results for researchers who act according to good scientific practice. Simulated data included 5% outliers.

Figure 12. (Selective Outlier Removal). Replication rate (RR) as a function of the number k of outlier rejection methods attempted, effect size d , and nominal α level (0.5% or 5%). The nominal α level and power of the replication study is $\alpha_2 = 0.05$ and $1 - \beta_2 = 0.90$. Dashed lines gives the results for p -hacking whereas solid lines depict the results for researchers who act according to good scientific practice. Simulated data included 5% outliers.