



Article submitted to journal

**Subject Areas:**

Psychology

**Keywords:**

Preregistration, Evidence,  
Reproducibility, Credible Interval,  
Bayes Factor

**Author for correspondence:**

Maarten Marsman

e-mail: [m.marsman@uva.nl](mailto:m.marsman@uva.nl)

# A Bayesian Bird's Eye View of "Replications of Important Results in Social Psychology"

Maarten Marsman<sup>1</sup>, Felix D. Schönbrodt<sup>2</sup>,  
Richard D. Morey<sup>3</sup>, Yuling Yao<sup>4</sup>, Andrew  
Gelman<sup>4</sup> and Eric-Jan Wagenmakers<sup>1</sup>

<sup>1</sup> University of Amsterdam

<sup>2</sup> Ludwig-Maximilians-Universität München

<sup>3</sup> Cardiff University

<sup>4</sup> Columbia University

We applied three Bayesian methods to reanalyze the preregistered contributions to the *Social Psychology* special issue "Replications of Important Results in Social Psychology" [1]. First, individual-experiment Bayesian parameter estimation revealed that for directed effect size measures, only 3 out of 44 central 95% credible intervals did not overlap with zero and fell in the expected direction. For undirected effect size measures, only 4 out of 59 credible intervals contained values greater than .10 (10% of variance explained), and only 19 intervals contained values larger than .05. Second, a Bayesian random-effects meta-analysis for all 38 *t*-tests showed that only one out of the 38 hierarchically estimated credible intervals did not overlap with zero and fell in the expected direction. Third, a Bayes factor hypothesis test was used to quantify the evidence for the null hypothesis against a default one-sided alternative. Only 7 out of 60 Bayes factors indicated non-anecdotal support in favor of the alternative hypothesis ( $BF_{10} > 3$ ), whereas 51 Bayes factors indicated at least some support for the null hypothesis. We hope that future analyses of replication success will embrace a more inclusive statistical approach by adopting a wider range of complementary techniques.

Skillfully-conducted replication studies can greatly influence researchers' confidence in the presence, impact, and general nature of a hypothesized effect. But how should replication studies be conducted? A recent special issue in *Social Psychology* showed by example how informative replication studies can be designed [1, 2]. In the special issue, the trinity of replication guidelines was to collaborate with original authors, to use preregistration, and to conduct high-powered studies [see also 3, 4, 5]. Therefore the replication studies in the special issue were conducted under relatively ideal circumstances.

Although the special issue has attracted widespread attention, the data have not yet been analyzed as a whole, across all replication attempts. In addition, the individual replication attempts were analyzed solely with classical statistics (i.e.,  $p$ -values and confidence intervals). Classical methods, however, are unable to quantify evidence; in particular, classical methods cannot distinguish between the absence of evidence (i.e., the data are uninformative) or the evidence of absence (i.e., the data support a point null hypothesis  $\mathcal{H}_0$ ). Hence it is possible that even a high-powered replication with a nonsignificant  $p$ -value can be evidentially uninformative for the question of interest [6, 7].

Consequently, our goals are twofold. Our primary goal is to provide a Bayesian bird's eye perspective on the results from the *Social Psychology* special issue [1]. Specifically, we reanalyze the results from the individual contributions to the special issue using three Bayesian methods. Firstly, an individual-study parameter estimation approach yields posterior distributions of effect size for each study considered in isolation; these posterior distributions quantify our uncertainty about the key quantity of interest – the narrower the posterior distribution for effect size, the more certain we can be about its value. Secondly, a hierarchical parameter estimation approach also yields posterior distributions for effect size, but it does not consider the studies in isolation; instead, the hierarchical approach implements group-level constraints and conceptualizes each individual study's effect size as a draw from a group-level normal distribution whose variance reflects the heterogeneity between studies. Thirdly, a hypothesis testing approach quantifies evidence for the point null hypotheses versus a specific one-sided alternative hypothesis. This cannot be accomplished by Neyman-Pearson style hypothesis testing, whose explicit goal it is to control error rate in repeated use. However, as emphasized by the editors of the *Social Psychology* special issue, confidence in scientific claims is based on “evaluating the evidence” [1] p. 139 (for a summary of other reasons to consider a Bayesian analysis see for example [8, 9]).<sup>1</sup> As explained below, all three Bayesian approaches presented here follow from the same coherent framework in which knowledge about parameters and hypotheses is updated based on predictive success.

Our secondary goal is to highlight the feasibility of analyzing data from standard experiments in social psychology using Bayesian tools. Specifically, we analyze the special issue data using JASP [10, [jasp-stats.org](http://jasp-stats.org)], Stan [11, 12, 13], and R [specifically, the BayesFactor package, 14]. Armed with these software programs, Bayesian methods can be easily applied to a series of common analyses such as the  $t$ -test, contingency tables, regression, and analysis of variance (ANOVA).

## 1. Brief Bayesian Background

At its core, Bayesian inference requires only that the user creates a “generative” statistical model, that is, a model that makes predictions about to-be-observed data. Once a particular data set is observed, Bayes' rule inverts the generative model and updates the uncertainty about the model parameters in a coherent fashion. The updated model makes new predictions about to-be-observed data, and the predict-observe-update cycle of Bayesian inference can continue indefinitely as the data accumulate [15]. The central aspect of Bayesian inference therefore is prediction: it is predictive performance that drives the coherent update of knowledge. We

<sup>1</sup>Although the editors may not have had Bayesian procedures in mind when they used the word “evidence”, the common sense interpretation of evidence as something that causes a change in opinion is consistent with the Bayesian statistical paradigm.

demonstrate this point by application to two key Bayesian tasks: parameter estimation and hypothesis testing.

### (a) Bayesian Parameter Estimation: The Basic Concepts

In order to have a model make predictions, its parameters need to be assigned particular values. Often we do not know exactly what these values are—they are the very entities we want to learn about. Consequently, Bayesians assign prior distributions to parameters  $\theta$  in order to reflect the uncertainty about their true value. These prior distributions  $p(\theta)$  quantify one's knowledge about the unknown parameters  $\theta$  before seeing the data. After seeing the data, the prior distribution  $p(\theta)$  is updated to a posterior distribution  $p(\theta | \text{data})$ : the uncertainty about  $\theta$  "given" the data. The updating proceeds by assessing predictive success, as can be seen by writing Bayes' rule as follows [16, 17]:

$$\underbrace{p(\theta | \text{data})}_{\text{Posterior knowledge}} = \underbrace{p(\theta)}_{\text{Prior knowledge}} \times \underbrace{\frac{p(\text{data} | \theta)}{p(\text{data})}}_{\text{Predictive updating factor}} \quad (1.1)$$

This equation shows that the update from prior to posterior distribution is governed by a predictive updating factor; this factor considers, for each value of  $\theta$ , its predictive success  $p(\text{data} | \theta)$ —that is, the probabilistic forecast for the observed data according to a specific  $\theta$ . This predictive success for a specific  $\theta$  is then assessed relative to the average predictive success  $p(\text{data})$ —the probabilistic forecast for the observed data across all values of  $\theta$ . Hence, the Bayesian updating process is guided by predictive success: parameters that predict well receive a boost in plausibility, whereas parameters that predict poorly suffer a decline [15, 16, 17].

Below we apply the Bayesian parameter estimation framework to the studies published in the *Social Psychology* special issue. For this purpose it is convenient to summarize the posterior distribution by its location (i.e., the posterior median) and spread (i.e., a 95% central credible interval). In the Bayesian framework, the interpretation of these summary values is intuitive and direct: Given the data and the statistical model—which includes the specification of the prior distribution as well as the likelihood—we can be 50% confident that the true value is higher or lower than the median, and we can be 95% confident that the true value lies in the interval [18, 19].

The prior distributions have been assigned by default, depending on general desiderata inherent to the Jeffreys-Zellner-Siow framework [20, 21, 22, 23, 24, 25, 26]. Note that the posterior distribution is a compromise between the prior and the data, and therefore—as long as the data are sufficiently informative—the posterior distribution will be relatively robust to changes in the specification of the prior distribution.

Below we report posterior distributions for both directed and undirected effect size measures. The directed effect size measures include  $\delta$  for  $t$ -tests and  $\rho$  for correlation tests, and the undirected effect size measures include Cramér's  $\phi^2$  for contingency tables [27, p. 282], and  $\rho^2$  for  $t$ -tests, correlation tests, and ANOVAs [e.g., 28]. For the  $t$ -test and ANOVA,  $\rho^2$  has the same interpretation as  $\omega^2$  and  $\eta^2$ —the proportion of variance explained by the experimental design. When more than one experimental factor is used, we report the squared semi-partial correlation to quantify the unique contribution of the primary experimental factor of interest.

We produced estimates for the directed effect sizes  $\delta$  and  $\rho$  using JASP and produced estimates for the undirected effect sizes  $\phi^2$  and  $\rho^2$  using the BayesFactor package [14] in R. We have made the JASP files, data and R-code available at <https://osf.io/bqwzd/>.

### (b) Bayesian Parameter Estimation: Hierarchical Models

The Bayesian parameter estimation approach detailed above can be gracefully extended to a hierarchical model, in which inference for individual studies is informed and constrained by a single overarching distribution: the group-level model [8, 29, 30]. Hierarchical analyses such

as the one applied in our reanalysis have three main benefits [31]. First, the individual-study results contribute to the estimation of group-level parameters that describe both the heterogeneity between studies and the group mean effect. Second, the group-level structure shrinks individual results that are uncertain and relatively extreme towards the group mean [e.g., 32]. Third, the uncertainty about individual studies is generally reduced when information is borrowed from other, statistically similar studies, which is expressed in narrower posterior credible intervals. In our reanalysis below we utilize this hierarchical approach for the available 38 *t*-tests reported in the special issue, as imposing a hierarchical structure on the effect sizes  $\delta$  is simple and intuitive.

For the hierarchical analysis on effect size from the 38 *t*-test studies, we use the model formulation from Rouder et al. [23]. For the one sample *t*-test, the mean in study *s* is parameterized as  $\sigma_s \delta_s$ , where  $\delta_s$  is the standardized effect size. For an independent *t*-test, the mean for group 1 is  $\mu_s + \frac{1}{2} \sigma_s \delta_s$  and the mean for group 2 is  $\mu_s - \frac{1}{2} \sigma_s \delta_s$ , and  $\sigma_s^2$  indicates the common variance. For the one sample *t*-test and the independent *t*-test, the key parameter of interest is the standardized effect size  $\delta_s$ . In the hierarchical model it is assumed that the effect sizes come from a single overarching distribution, that is, the group-level model. Note that many experiments from the special issue concern phenomena that are conceptually unrelated; consequently, the group-level distribution describes the location and heterogeneity of “important results in social psychology” that were deemed suitable for preregistered replication. As will be evident below, the heterogeneity in effect sizes is estimated to be relatively small, and this results in a substantial shrinkage effect.

One reviewer objected to the use of a hierarchical model for effects that are conceptually unrelated. This is a venerable issue (e.g., see [32] for an extended discussion) and we offer the following motivation. First, the hierarchical model contains a parameter that measures the heterogeneity across studies. In our application, the studies turn out to be highly homogeneous. This does not show that the studies are conceptually related, but it does show that their effect sizes are highly similar, and this is all that is required for an application of the model. In other words, the model assumes only that the effect sizes across studies are statistically similar, and does not speak to the degree of conceptual similarity. Second, the extent to which effects are conceptually related is not an all-or-none matter. One may always argue that the individual case is unique and, at some level, conceptually unrelated to the other cases. However, all studies considered here come from social psychology and have been submitted to the same special issue. We believe this commonality warrants the application of the hierarchical model.

A limitation of our Bayesian random-effects meta-analysis is that it assumes that the 38 effect sizes  $\delta$  are independent realizations from a single overarching distribution. This is of course not entirely the case, as several *t*-tests were used to test the same hypothesis (e.g., the analyses by IJzerman et al. [33]), or were used to test an effect on several measures within the same experiment (e.g., the analyses by Johnson et al. [34]). However, a data set of 38 *t*-tests will only admit a model of limited complexity, and we believe that our model achieves the right balance in the inevitable tradeoff between bias and variance (e.g., [35]). Our assessment is bolstered by the fact that for the data set under consideration there are relatively few “duplicates”, and there is relatively little heterogeneity across effect sizes; consequently, there is almost no information available to support the inclusion of additional topic-specific parameters.

We assume here that the standardized effect sizes  $\delta_s$  follow a normal distribution with an unknown group mean  $\theta$  and variance (i.e., study heterogeneity)  $\tau^2$ . To complete the Bayesian hierarchical model, we have used standard non-informative priors on the individual-study means  $\mu_s$ , individual-study variances  $\sigma_s$  and the group-level variance  $\tau^2$  (i.e.,  $p(\mu_s, \sigma_s) \propto \sigma_s^{-2}$  and  $p(\tau) \propto \tau^{-2}$ ), and have used a Cauchy(0,  $1/\sqrt{2}$ ) prior on the group-level mean  $\theta$  that is in line with the prior on effect sizes  $\delta$  in the individual analyses. The analysis was performed using the R-package rstan [11, 12, 13], with the data and R-code available at <https://osf.io/bqwzd/>.

### (c) Bayesian Hypothesis Testing: Bayes Factors

The predictive framework that governs the coherent plausibility updates for parameters carries over seamlessly to plausibility updates for entire models or hypotheses. To see this, we again use Bayes' rule and obtain:

$$\underbrace{\frac{p(\mathcal{H}_1 | \text{data})}{p(\mathcal{H}_0 | \text{data})}}_{\text{Posterior knowledge about hypotheses}} = \underbrace{\frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)}}_{\text{Prior knowledge about hypotheses}} \times \underbrace{\frac{p(\text{data} | \mathcal{H}_1)}{p(\text{data} | \mathcal{H}_0)}}_{\text{Predictive updating factor}} \quad (1.2)$$

As was the case for parameter estimation, the update from prior plausibility to posterior plausibility is governed by predictive success: the hypothesis that predicts the observed data better than the competitor hypothesis receives a boost in plausibility [e.g., 15, 36].

Note that the framework is inherently relative: what matters is which of the two hypotheses does best, not whether a specific hypothesis does well in an absolute sense. Also note that the predictive focus means that the results do not depend on one of the statistical models being "true" in some abstract sense. This latter point is particularly relevant in the context of a point null hypothesis, which many have argued is an unlikely proposition on a priori grounds [37, 38]. For the interpretation of the Bayes factor, however, it does not matter whether the point null hypothesis (or the alternative hypothesis against which it is pitted) is unlikely to be true in an absolute sense; indeed, all models are abstraction of reality and are therefore likely "wrong". However, in a predictive sense the point null hypothesis can be a good approximation for an effect that is so small that it cannot be detected reliably. As remarked by Andrew Gelman, "when effect size is tiny and measurement error is huge, you're essentially trying to use a bathroom scale to weigh a feather—and the feather is resting loosely in the pouch of a kangaroo that is vigorously jumping up and down."<sup>2</sup> In such situations, the point null hypothesis will predictively outperform the alternative hypothesis.

Equation 1.2 quantifies the adage "extraordinary claims require extraordinary evidence"; in Bayesian terms, this translates to the statement "An implausible hypothesis requires substantial predictive success". The quantification of prior implausibility of a hypothesis is subjective and may depend on many unknowns. We therefore follow standard Bayesian practice and quantify only the predictive updating factor, that is, the degree to which the data change the relative plausibility of the hypotheses under consideration.

Thus, for the studies from the special issue in *Social Psychology* we report the predictive updating factor

$$BF_{10} = \frac{p(\text{data} | \mathcal{H}_1)}{p(\text{data} | \mathcal{H}_0)},$$

which is commonly known as the Bayes factor<sup>3</sup> [21, 43]. The result  $BF_{10} = 2$  indicates that the observed data are twice as likely under  $\mathcal{H}_1$  than under  $\mathcal{H}_0$ ; the result  $BF_{10} = 0.5$  indicates the exact opposite. Harold Jeffreys proposed a set of descriptive categories of evidential impact, and proposed that Bayes factors in between 3 and 1/3 are "not worth more than a bare mention" [21, Appendix B]. Although the interpretation of evidence does not require arbitrary category labels, they do facilitate a concise summary. As in the case of parameter estimation, we use the default Jeffreys-Zellner-Siow priors, with the exception that the priors are specified to be one-sided, respecting the fact that in replication research the hypothesis of interest is directional. Other prior choices are possible and, to the extent that they ask a slightly different question, they might

<sup>2</sup>See <http://andrewgelman.com/2015/04/21/feather-bathroom-scale-kangaroo/>. A similar sentiment was expressed by Edwards et al. [39, p. 215-216]: "Convention asks, 'Do these two programs differ at all in effectiveness?' Of course they do. Could any real difference in the programs fail to induce at least some slight difference in their effectiveness? Yet the difference in effectiveness may be negligible compared to the sensitivity of the experiment. In this way, the conventional question can be given meaning, and we shall often ask it without further explanation or apology."

<sup>3</sup>Andrew Gelman wishes to state that he hates Bayes factors. The reasons for his aversion are detailed in Gelman and Rubin [40] and Gelman et al. [41, Chapter 6] and mainly concern the practice of assigning prior mass to a single point from a continuous distribution, and the resulting sensitivity to the prior distribution (see also [42]).

lead to a slightly different answer [e.g., 7, 44, 45, 46, 47]. The default choices in the present work are built into the BayesFactor package for R [14] and the companion software JASP [10].

#### (d) Single Study Example

In order to explain our reanalysis approach in more concrete terms, consider the following example. Shackelford et al. [48] reported that men were more distressed by sexual infidelity than women, an effect that IJzerman et al. [33] sought to replicate in several studies reported in the *Social Psychology* special issue. Their first replication study featured 18 men and 69 women, and the results showed that compared to the women, the men had significantly lower sexual dilemma scores (SDS;  $t_{85} = 4.178$ ,  $p < 0.001$ ,  $d = 1.106$ ), where lower SDS-scores indicate higher distress.<sup>4</sup> Below we report the Bayesian results for single-study parameter estimation and for hypothesis testing.

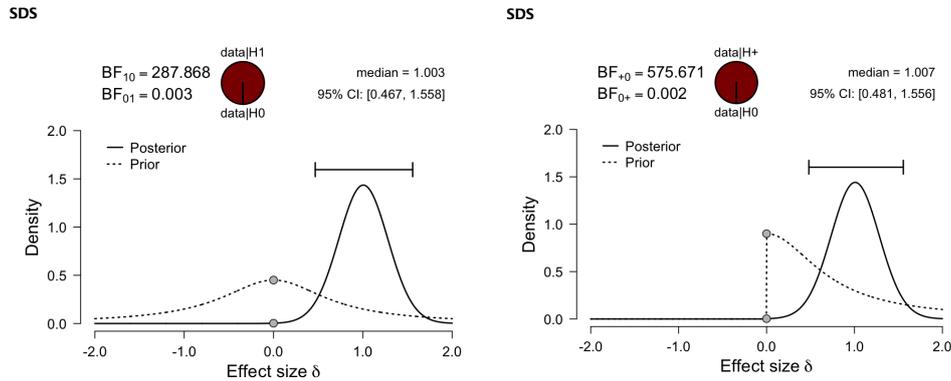
Concerning Bayesian parameter estimation, the left panel of Figure 1 shows the prior distribution (dotted line) and the posterior distribution (solid line) for the effect size  $\delta$  based on the data from the first replication study reported in IJzerman et al. [33]. The data have caused the prior distribution to undergo a substantial update, indicating that the data were highly informative. The posterior median is 1.00 and the central 95% credible interval equals [0.47, 1.56]. In addition to a directed effect size  $\delta$  we also report the undirected effect size  $\rho^2$ , which quantifies the proportion of the variance of sexual dilemma scores that can be explained by gender differences. In this example, the posterior median for  $\rho^2$  is 0.14 (i.e., 14% variance explained) and the central 95% credible interval equals [0.03, 0.29]. These are the posterior summary measures that we report below for all available studies in the *Social Psychology* special issue, and later we also report the results from a hierarchical analysis that considers multiple studies simultaneously.

Concerning Bayesian hypothesis testing, the left panel in Figure 1 shows the Bayes factor  $BF_{10}$ . This Bayes factor contrasts the predictive success of the alternative hypothesis  $\mathcal{H}_1$  (i.e., SDS-scores differ between men and women) with that of the null hypothesis  $\mathcal{H}_0$  (SDS-scores do not differ between men and women). The result,  $BF_{10} \approx 288$ , indicates that the observed data are almost 288 times more likely to occur under  $\mathcal{H}_1$  than under  $\mathcal{H}_0$ , providing very strong support for the alternative hypothesis.

However, the replication study did not aim to test whether SDS-scores differ between men and women; instead, its aim was to test a directional hypothesis, namely that men have lower SDS-scores than women (i.e., men are more distressed than women about sexual infidelity, not less). We denote the directional hypothesis as  $\mathcal{H}_+$ . The directional nature of the hypothesis can be incorporated into the analysis by folding the prior distribution such that it only has mass for effect sizes in the predicted direction [49]. The right panel of Figure 1 shows the prior and posterior distributions produced by this directional hypothesis. The associated Bayes factor  $BF_{+0}$  contrasts the directional alternative hypothesis  $\mathcal{H}_+$  against the null hypothesis  $\mathcal{H}_0$ . The result,  $BF_{+0} \approx 576$ , indicates that the observed data are about 576 times more likely to occur under  $\mathcal{H}_+$  than under  $\mathcal{H}_0$ , again providing very strong support for the alternative directional hypothesis. Note that in this specific case the evidence is almost twice as strong for the directional hypothesis  $\mathcal{H}_+$  (Figure 1, right panel) than it is for the unrestricted hypothesis  $\mathcal{H}_1$  (Figure 1, left panel), despite the fact that the posterior distributions for effect size are virtually identical. The reason is that the directional hypothesis makes predictions that are more daring than those of the unrestricted hypothesis; when the effect goes in the expected direction, the daring predictions are validated and the associated gain in plausibility is therefore higher. This provides an example of how Bayes factors quantify the idea that risky scientific predictions ought to be rewarded more than vague scientific predictions [e.g., 50]. The directional Bayes factor  $BF_{+0}$  is the measure of evidence that we report below for all available studies in the *Social Psychology* special issue.

We have demonstrated that the Bayes factor for the null hypothesis against an alternative hypothesis depends partly on the predictions for effect size under that alternative hypothesis.

<sup>4</sup>Our goal in this paper is comparative, and hence we use models that are similar to the models from the original paper, in this case, a normal error model for ordinal data.



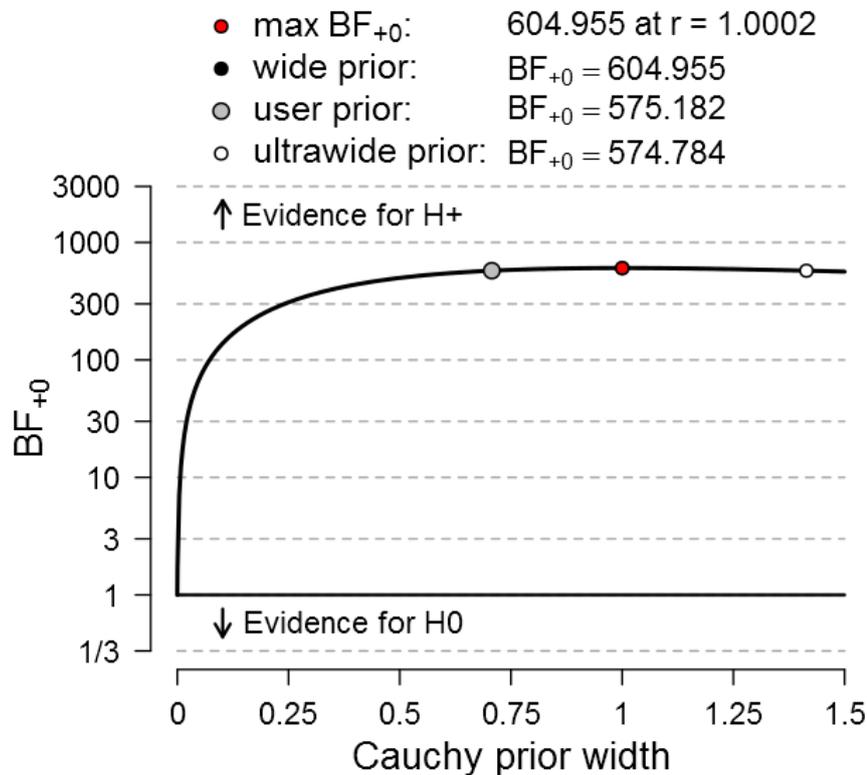
**Figure 1.** JASP output for the  $t$ -test example featuring the first replication study reported in IJzerman et al. [33]. Both panels show the prior and posterior distributions for effect size, the posterior median, the central 95% credible interval, and the Bayes factor. The left panel shows the results for the unrestricted hypothesis  $\mathcal{H}_1$  and the right panel shows the results from the directional hypothesis  $\mathcal{H}_+$ .

These predictions are a direct consequence of the prior distribution that is assigned to effect size; for instance, a two-sided prior yielded  $BF_{10} \approx 288$  whereas the one-sided prior gave  $BF_{+0} \approx 576$ . As mentioned above, for the analysis of the studies from the special issue we use default specifications designed to meet general desiderata [e.g., 20, 51]. However, it is possible to entertain a range of alternative prior specifications and examine the robustness of the conclusions. The standard method to conduct such a robustness check or sensitivity analysis is to vary the width of the prior distribution and consider the resultant change in the Bayes factor.

We discuss the pros and cons of such a robustness check by applying it to the IJzerman experiment. Figure 2 shows the associated output from the JASP “SumStats” module.<sup>5</sup> As dictated by the directionality of the hypothesis under scrutiny, the prior distribution on effect size only assigns mass to positive values. What varies on the  $x$  axis is the width of the prior distribution under  $\mathcal{H}_+$ . Figure 2 indicates that –for all but the smallest values of the prior width– there is compelling evidence for  $\mathcal{H}_+$  over  $\mathcal{H}_0$  in the sense that the observed data are hundreds of times more likely under  $\mathcal{H}_+$  than under  $\mathcal{H}_0$ . The red dot indicates that the Bayes factor is highest for a width of  $r = 1.0002$ , where it equals about 605; coincidentally, this post-hoc prior width is almost exactly the same as the value of the “wide prior” (i.e.,  $r = 1$ ) which was originally proposed as a useful default by Jeffreys [21]; consequently, the red dot obscures the black dot. The grey dot indicates the “user prior”, which is the modern-day default of  $r = 2^{-0.5} \approx 0.707$  [14].

Despite the fact that the evidence is compelling across a wide range of prior widths, Figure 2 also shows that for small values of the width, the evidence is rather weak. As the width approaches zero, the Bayes factor approaches 1, and in extremis, when  $r = 0$ , the Bayes factor equals 1 irrespective of the data. This occurs because when  $r = 0$ , the alternative hypothesis  $\mathcal{H}_+$  has morphed into the null hypothesis  $\mathcal{H}_0$ , and two models that make identical predictions can never be discriminated based on empirical data. This is an important realization because some researchers may feel the need to reduce the default width in order to accommodate a more modest expectation about effect size. Although a careful subjective specification of prior distributions is generally advantageous (particularly when it occurs in advance of seeing the data), Figure 2 indicates that there is a danger when such specification is done solely with respect to the prior width [9]; the default priors are centered on zero and reducing the width will make  $\mathcal{H}_+$  increasingly similar to  $\mathcal{H}_0$ . We surmise that in situations where the researcher is unhappy

<sup>5</sup>The SumStats module conducts Bayesian inference for statistical scenarios that are uniquely defined by a small set of summary statistics (e.g., the binomial,  $t$ -tests, linear regression, and correlation).



**Figure 2.** JASP robustness analysis of the  $t$ -test example featuring the first replication study reported in IJzerman et al. [33]. The evidence in favor of the directional hypothesis  $\mathcal{H}_+$  varies as a function of the width of the prior distribution for effect size (i.e., “Cauchy prior width”, on the  $x$  axis). When the width equals zero,  $\mathcal{H}_0$  and  $\mathcal{H}_+$  are identical and the Bayes factor is 1 regardless of the data. Here the evidence in favor of  $\mathcal{H}_+$  is compelling, with Bayes factors exceeding 100 for all but the most narrow priors. See text for details.

about the default prior width, the researcher is also unhappy about the default prior location – note that under the default specification, the most likely value of effect size under the alternative hypothesis is zero, and absolute values are always more likely the closer they are to zero. Ideally then, a subjective specification ought to take into account both the width and the location of the prior distribution for effect size under the alternative hypothesis.

The current functionality of the JASP program offers only a sensitivity analysis with respect to the prior width. Although informative, this procedure is limited and results for effect sizes near zero should be interpreted with considerable care.<sup>6</sup> For this reason, and in order not to overwhelm the reader, below we report only the results for the default prior setting. However, the JASP analysis files do contain the sensitivity analyses as shown in Figure 2 for the data from the study by IJzerman et al. [33].

<sup>6</sup>We are currently working to expand JASP to accommodate prior distributions that are not centered on zero, thereby facilitating a more complete subjective specification.

## 2. Results for All Studies

Our reanalysis of the results from IJzerman et al. [33, Study 1] indicates compelling statistical support for a replication of the original findings [48]: Bayesian parameter estimation showed the posterior distribution for the effect size  $\delta$  to be away from zero, and Bayesian hypothesis testing confirmed that the default one-sided alternative hypothesis made predictions that were superior to those from the null hypothesis.

We now report the results from a Bayesian reanalysis of the main results across the replication studies in the *Social Psychology* special issue<sup>7</sup> with the exception of the ManyLabs project [54].

The ManyLabs project was excluded for several reasons. Firstly, the ManyLabs project mostly contained replications of benchmark findings outside of social psychology. Secondly, for every finding under scrutiny the ManyLabs project featured many replication attempts, and this demands a different analysis approach from the one that is appropriate for single replication attempts. Finally, the results from the ManyLabs project do not much benefit from a sophisticated statistical reanalysis: the conclusions are already evident from a plot of effect sizes reported across the many participating labs [i.e., 54, Figure 1]. In other words, when a finding has been subjected to multiple replication attempts the data are bound to pass Berkson's interocular traumatic test, when the conclusion hits one straight between the eyes [39].

Most of the Bayesian reanalyses presented below have been produced with JASP [10]. For analyses currently unavailable in JASP, we mostly used the BayesFactor R-package [14]. Specifically, the BayesFactor package was used to produce the  $\rho^2$  and  $\phi^2$  effect size estimates, and to produce some of the Bayes factors involving directional hypotheses for ANOVAs and contingency tables. Finally, the hierarchical analysis was programmed in Stan [11, 12, 13]. For each study, the entire reanalysis—the data set, the R code, and the JASP files with input options—is made available through the Open Science Framework (<https://osf.io/bqwzd/>).

### (a) Results from Bayesian Parameter Estimation: Individual Studies

This section summarizes the results for individual-study parameter estimation of the contributions to the *Social Psychology* special issue. First we discuss directed effect sizes such as  $\delta$  and  $\rho$ , which originate from  $t$ -tests and correlation tests, respectively. Next we discuss undirected effect sizes such as  $\rho^2$  that originate from  $t$ -tests, ANOVAs, correlation test, and contingency tables.

#### (i) Directed Effect Sizes

The results for 44 directed effect sizes are shown in Figures 3 and 4, with posterior medians indicated as dots and the central 95% credible intervals as horizontal lines. The posterior distributions in Figures 3 and 4 were obtained from the unrestricted model (cf. the left panel of Figure 1), recoded such that the prediction of interest stipulates the directed effect sizes to be positive. Furthermore, we have sorted the results in Figures 3 and 4 according to the posterior median values, with the top level entry showing the largest posterior median, and the bottom level entry showing the smallest posterior median.

Figure 3 shows the results for the effect sizes  $\delta$  separately for each of 38 replication studies and analyses utilizing  $t$ -tests. From Figure 3 we see that the estimated effect sizes are generally small, with only 9 out of 38 studies yielding credible intervals that contained values  $|\delta| \geq 0.5$ . Furthermore, we see that despite being designed for high power, many studies still yield credible intervals that are relatively wide—an indication that there remains considerable uncertainty about the true value of effect size under  $\mathcal{H}_1$  [cf. 55]. With small average effects and wide credible

<sup>7</sup>We were unable to perform a Bayesian equivalent of the primary classical analysis in Nauts et al. [52], which involved a signed rank test. Even though there has been some pioneering work for a Bayesian equivalent of this test [see 53], currently this does not involve a Bayes factor. We therefore focused our efforts on the analysis of the open-ended descriptions (i.e., the  $t$ -test reported on page 159).

intervals, only 3 out of 38 central 95% credible intervals do not overlap with zero and fall in the intended direction. All three of these studies come from the IJzerman replication effort.

Figure 4 shows the results for the correlation effects  $\rho$  separately for each of six analyses reported by Sinclair et al. [56] utilizing correlation tests. From Figure 4 we again see relatively small effect size estimates and credible intervals that cover about 10% of the parameter range  $[-1, +1]$ . More importantly, only 2 credible intervals did not overlap with zero, both in the unintended direction.

As is evident from Figures 3 and 4, only 6 out of the 44 credible intervals do not overlap with zero, and only 3 of these are in the predicted (positive) direction.

## (ii) Undirected Effect Sizes

The general pattern of results for the directed effect sizes is corroborated by the 59 undirected effect size estimates that are shown in Figures 5 and 6, with posterior medians indicated as dots and the central 95% credible intervals as vertical lines. Figure 5 summarizes the results for the  $\rho^2$  effect size estimates for the 54 studies that used either a  $t$ -test, a correlation test, or an ANOVA, sorted according to the posterior median values. From Figure 5 we see that only 17 out of the 54 credible intervals contained values larger than .05 (5% variance explained), and only 3 contained values larger than .10 (10% variance explained).<sup>8</sup>

Figure 6 shows the results for the  $\phi^2$  effect size estimates for the five studies that used contingency tables, also sorted according to the posterior median values. We report squared  $\phi$  values here since for a  $2 \times 2$  contingency table Cramér's  $\phi$  reduces to Pearson's  $\rho_\phi$ , making  $\phi^2$  comparable to  $\rho^2$  in this particular case.<sup>9</sup>

Figure 6 shows that only two credible intervals contained values larger than .05 (5% shared variance), and only one (i.e., the interval for the study by Vermeulen et al. [59]) contained values larger than .10 (10% shared variance). Note that the central 95% credible interval from the Vermeulen study is relatively wide, ranging from .01 to .27, indicating considerable uncertainty about the true value of this effect size.

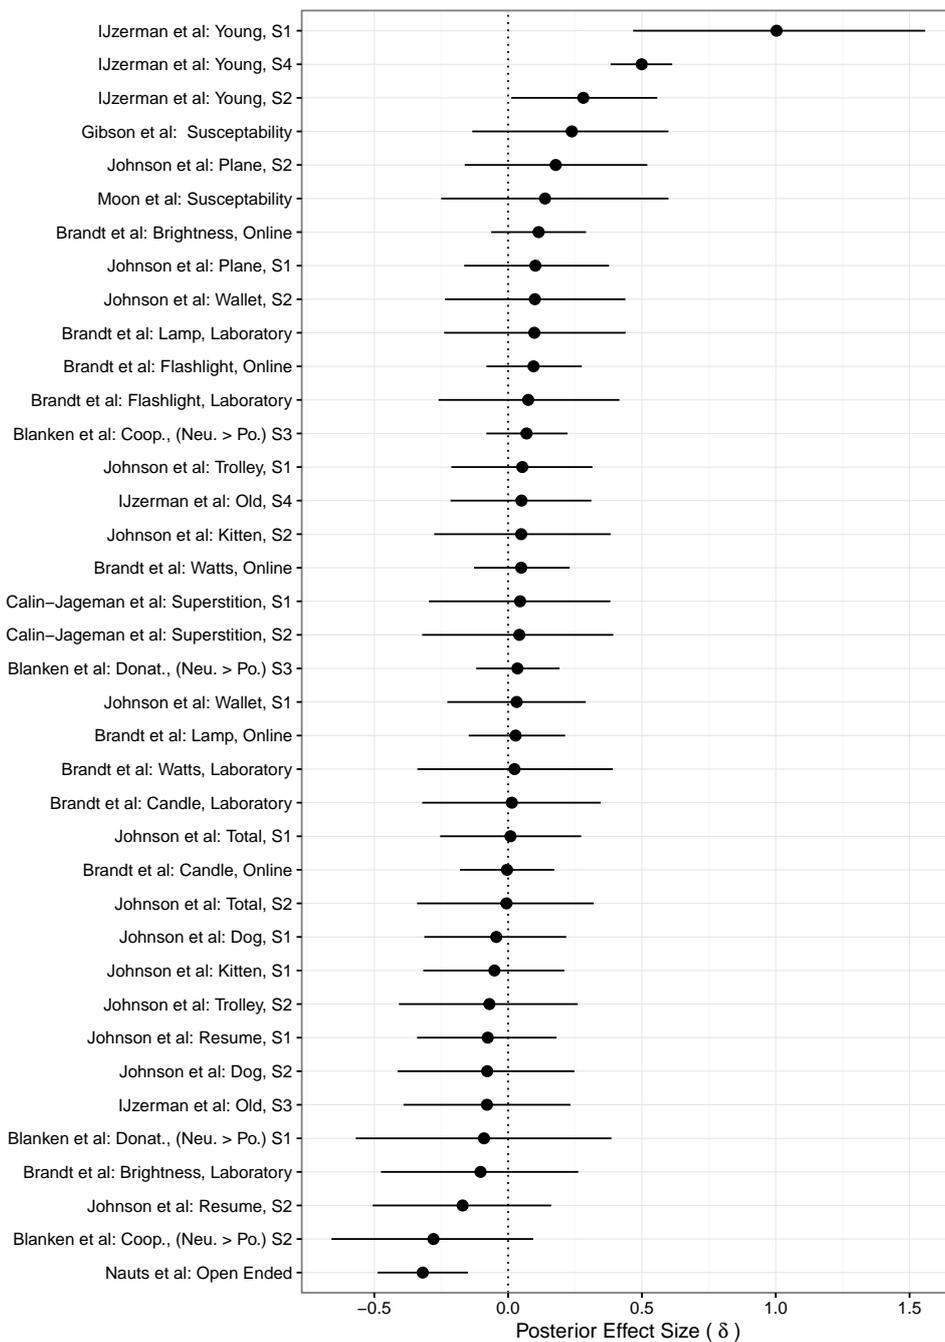
## (b) Results from Bayesian Parameter Estimation: Hierarchical Analysis

This section summarizes the results from the Bayesian random-effects meta-analysis of the effect sizes  $\delta$  obtained from the 38  $t$ -tests in the *Social Psychology* special issue. Figure 7 shows the posterior distributions for the two group-level parameters: the group mean effect  $\theta$  and the between-study heterogeneity  $\tau^2$ . As is evident from Figure 7, there is a relatively small overall group mean Cohen's effect size  $\theta$ , with a posterior median of about 0.05 and a 95% central credible interval that overlaps with zero and ranges from  $-0.01$  to  $0.12$ . Furthermore, the between-study heterogeneity was relatively small; the posterior median for  $\tau^2$  equals 0.019 and the central 95% credible interval equals  $[0.006, 0.035]$ .

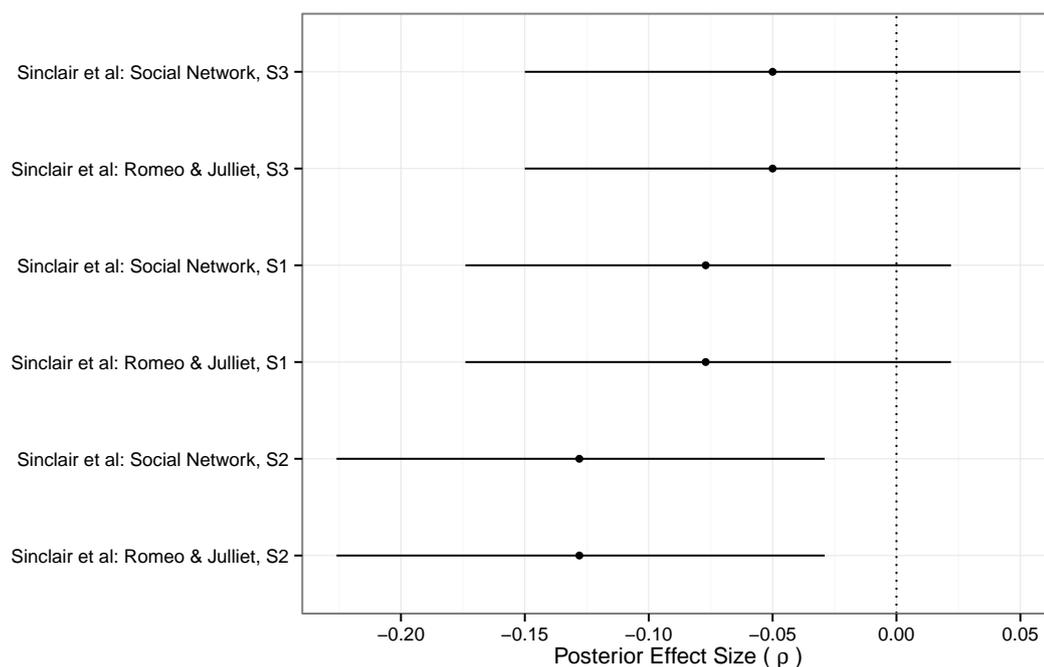
Figure 8 shows the hierarchically estimated credible intervals separately for each of the 38 studies, with posterior medians indicated as dots and the central 95% credible intervals as vertical lines. For a clear comparison, we have retained that same limits on the  $x$ -axis that were used to report the posterior distributions for individual experiments in Figure 3, and sorted the results according to the posterior median values in Figure 3. When we contrast the hierarchical estimates from Figure 8 to the individual estimates from Figure 3, we observe that the hierarchically estimated credible intervals are considerably shorter than the individual estimates. However, despite the substantial decrease in posterior uncertainty, only one of the hierarchically estimated credible intervals reported in Figure 8 did not overlap with zero and fell in the expected direction. This regularity is the result of a substantial shrinkage effect that pulls individual posteriors toward the group mean  $\theta$ . Shrinkage is particularly pronounced for effect size estimates that are relatively extreme and uncertain, as is exemplified by the first replication experiment reported in

<sup>8</sup>The primary effect of interest in Müller and Rothermund [57] could not be properly partialled out.

<sup>9</sup>The Wesselmann et al. [58] study featured a  $3 \times 3$  contingency table; all other studies reported in Figure 6 featured a  $2 \times 2$  contingency table.



**Figure 3.** Individual Bayesian parameter estimation results for directed effect sizes  $\delta$  for each of 38 experiments reported in the *Social Psychology* special issue that used *t*-tests. The posterior medians are indicated as dots and the central 95% credible intervals as vertical lines. The effect sizes were estimated using separate unrestricted models, but recoded such that they are predicted to be positive. Figure available at <https://flic.kr/p/FQJrUr>, under CC license <https://creativecommons.org/licenses/by/2.0/>.



**Figure 4.** Individual Bayesian parameter estimation results for directed effect sizes  $\rho$  for each of 6 experiments reported in the *Social Psychology* special issue that used correlation tests. The posterior medians are indicated as dots and the central 95% credible intervals as vertical lines. The effect sizes were estimated using separate unrestricted models, but recoded such that they are predicted to be positive. Figure available at <https://flic.kr/p/FqBRsm>, under CC license <https://creativecommons.org/licenses/by/2.0/>

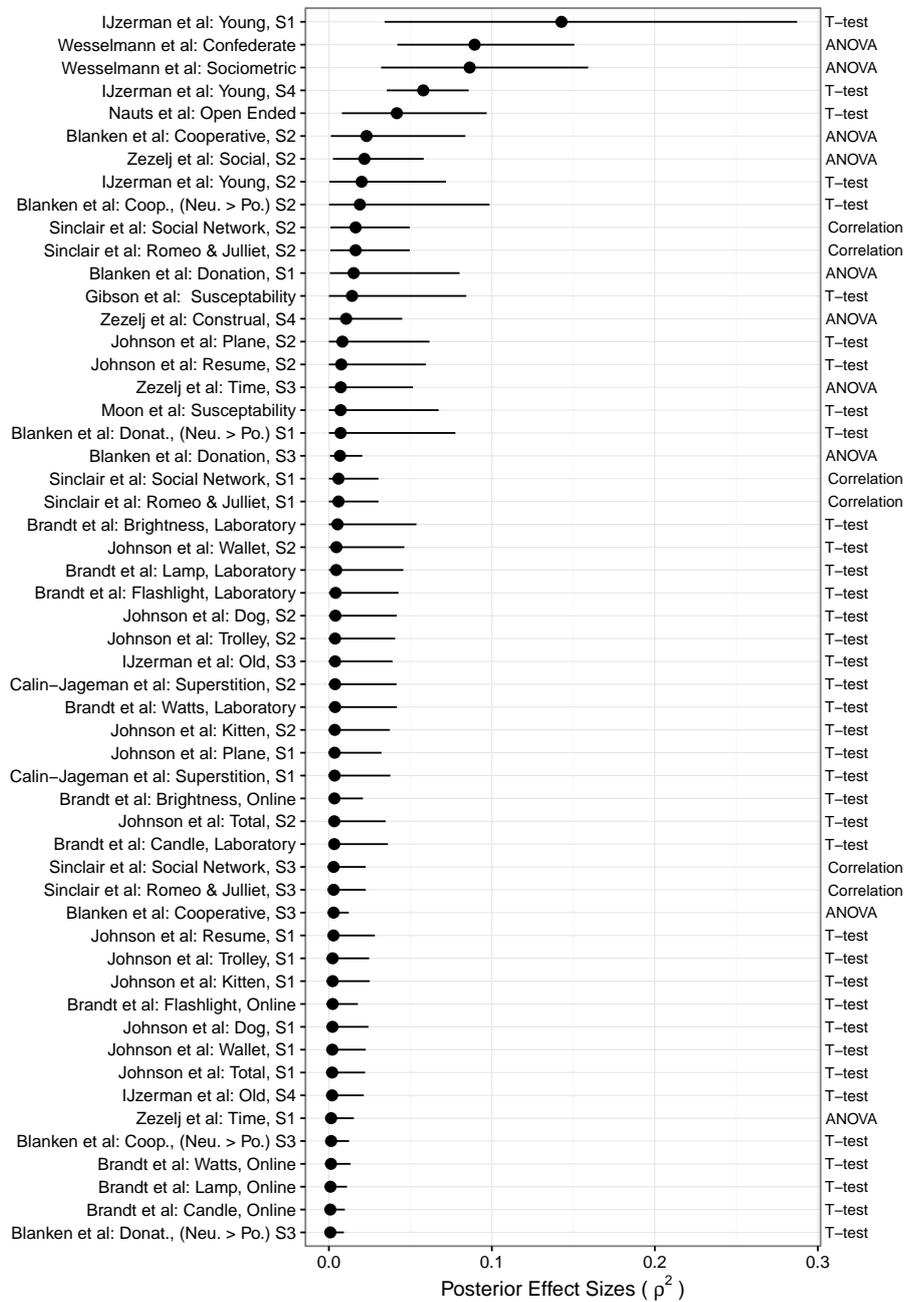
[33]. The individual estimate (shown as the top entry in Figure 3) yields a posterior median of 1.0 and a central 95% credible interval of [0.467, 1.556]; in contrast, the hierarchical estimate yields a much smaller posterior median of 0.26 and a central 95% credible interval of [-0.002, 0.559].

### (c) Results from Bayesian Hypothesis Testing

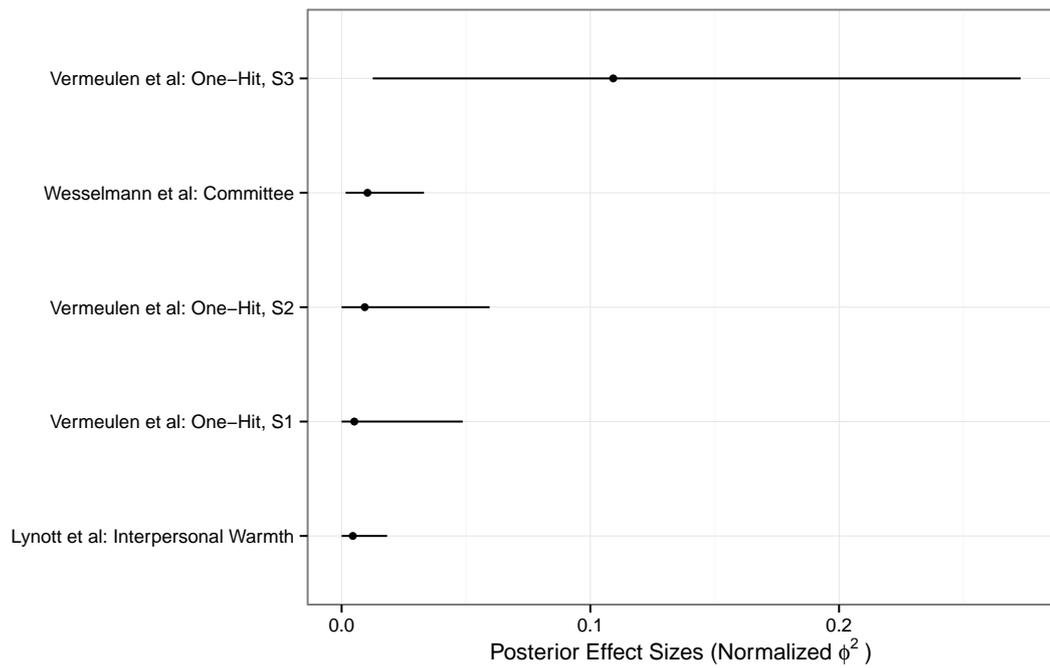
This section reports the 60 default Bayes factors that test the directional hypothesis  $\mathcal{H}_+$  against the null hypothesis  $\mathcal{H}_0$ . Figure 9 shows the Bayes factors sorted according to the degree to which they support  $\mathcal{H}_+$ , with the top level entry showing the most support for  $\mathcal{H}_+$ , and the bottom level entry showing the most support for  $\mathcal{H}_0$ . The overall results are qualitatively consistent with those from the credible intervals: 9 of the 60 Bayes factors showed evidence in favor of the alternative hypothesis, but only 7 showed evidence for the alternative hypothesis that is more than anecdotal (i.e.,  $\text{BF}_{+0} > 3$ ). The remaining 51 Bayes factors showed evidence in favor of the null hypothesis. Out of these 51, a total of 6 indicated anecdotal support for  $\mathcal{H}_0$  (i.e.,  $\text{BF}_{+0} \in (1/3, 1)$ ), 34 indicated moderate support for  $\mathcal{H}_0$  (i.e.,  $\text{BF}_{+0} \in (1/10, 1/3)$ ), and 11 indicate support for  $\mathcal{H}_0$  which is strong to extreme (i.e.,  $\text{BF}_{+0} < (1/10)$ ).

## 3. General Discussion

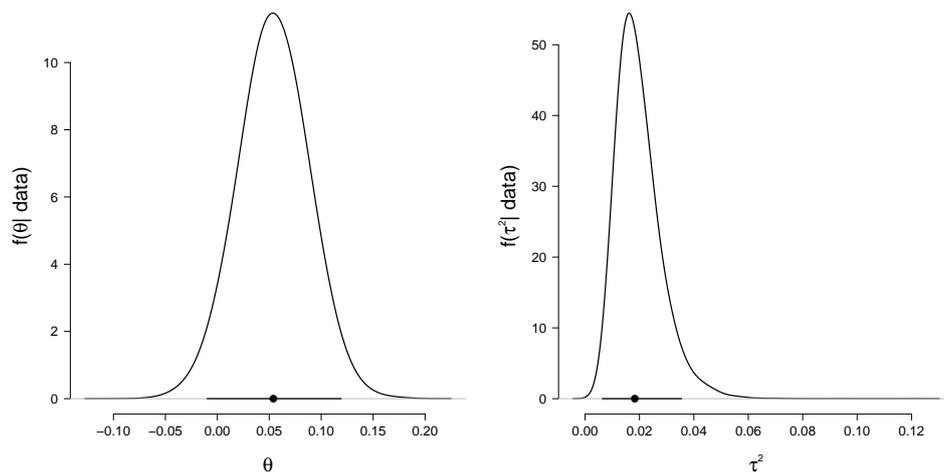
We reanalyzed the replication studies from the *Social Psychology* special issue on “Replications of Important Results in Social Psychology”. Our primary aim was to reanalyze and summarize the results from the individual contributions using Bayesian methods. We selected the *Social*



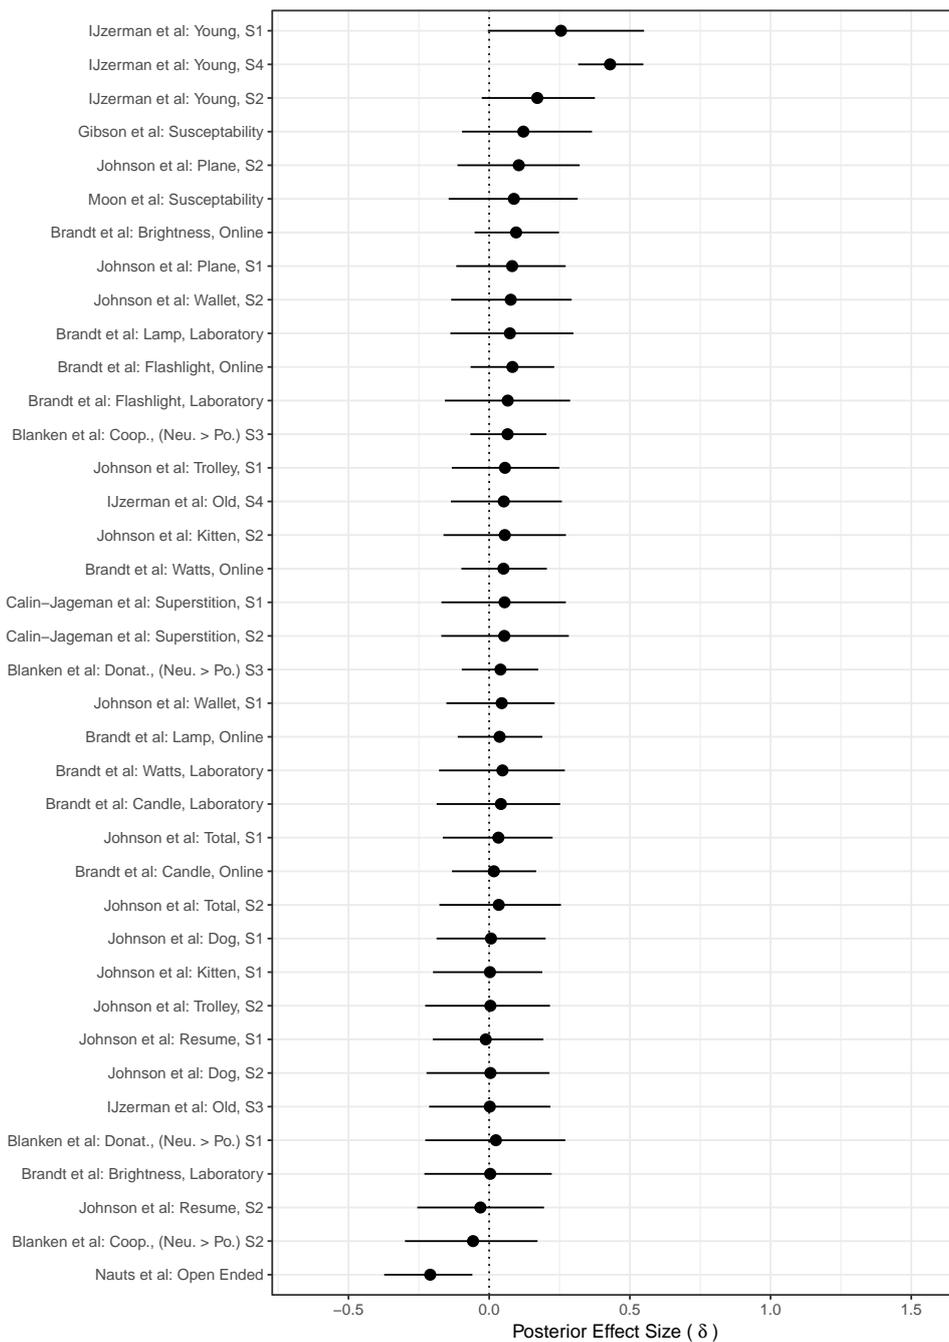
**Figure 5.** Individual Bayesian parameter estimation results for undirected effect sizes  $\rho^2$  for each of 54 experiments reported in the *Social Psychology* special issue that used either *t*-tests, correlation tests, or ANOVAs. The posterior medians are indicated as dots and the central 95% credible intervals as vertical lines. The effect sizes were estimated using an unrestricted model. Figure available at <https://flic.kr/p/FJSqWH>, under CC license <https://creativecommons.org/licenses/by/2.0/>



**Figure 6.** Individual Bayesian parameter estimation results for undirected effect sizes  $\phi^2$  for each of 5 experiments reported in the *Social Psychology* special issue utilizing contingency tables. The posterior medians are indicated as dots and the central 95% credible intervals as vertical lines. The effect sizes were estimated using an unrestricted model. Figure available at <https://flic.kr/p/EVhMSS>, under CC license <https://creativecommons.org/licenses/by/2.0/>



**Figure 7.** Posterior distributions for the group-level parameters of the hierarchical normal model describing the distribution of the 38  $t$ -test effect sizes in our reanalysis. Posterior medians are indicated with dots and central 95% credible intervals are indicated with horizontal lines.



**Figure 8.** Individual results from a Bayesian random-effects analysis for the directed effect sizes  $\delta$  from each of 38 experiments that used  $t$ -tests. The posterior medians are indicated as dots and the central 95% credible intervals as vertical lines. The effect sizes were estimated using an hierarchical model. Figure available at <https://flic.kr/p/FqBRto>, under CC license <https://creativecommons.org/licenses/by/2.0/>

*Psychology* special issue for several reasons. First, the data are publicly available on the Open Science Framework, greatly facilitating reanalysis. Second, the special issue has attracted much

attention, but no Bayesian analysis or systematic overview of the results is currently available. Finally, the contributions to the special issue met several key desiderata: the findings under scrutiny were judged to be important (i.e., nontrivial, theoretically valuable); the experiments were designed taking into account feedback from the original authors; the data analysis plan was preregistered; and the studies were designed for high power. The set of replications from the *Social Psychology* special issue appear to represent an ideal scenario, one in which all prerequisites for successful replication have been met.

Our reanalysis featured three complementary Bayesian methods. *Individual-experiment parameter estimation* revealed that (1) for only 3 out of 44 directed effect sizes did the estimates go in the intended direction with central 95% credible intervals excluding zero, and (2) for only 4 out of 59 undirected effect sizes did the central 95% credible intervals contain values greater than .1 (10% variance explained). *Hierarchical random-effects meta-analysis* for 38 *t*-tests revealed that (1) the group-level mean effect size is near zero, and (2) the hierarchically estimated credible intervals for the individual experiments showed a considerable shrinkage effect, underscoring the uncertainty surrounding the results obtained from individual experiments. Only one study yielded a central 95% credible interval that did not overlap with zero and fell in the intended direction. Finally, *Bayes factor hypothesis tests* revealed that for only 7 out of 60 hypotheses did the Bayes factor indicate non-anecdotal evidence in favor of the hypothesis of interest.

### (a) Dangers of Generalizing the Results Beyond the Special Issue

Across the empirical sciences there are recent signs of a “crisis of confidence” [60] and a “crisis of reproducibility” [61]. In psychology, several carefully conducted, large-scale replication initiatives have generally produced disappointing outcomes (e.g., [62]). In light of these developments it is tempting to view our results as another nail in the coffin for experimental psychology in general and social psychology specifically. However, such skepticism may be misplaced.

The strongest argument against blindly generalizing the present results to the entire field is that the studies from the special issue may not be representative. The special issue authors may have proposed the studies because they had prior knowledge—obtained from pilot experiments, colleagues, or expert assessment of the plausibility of a given claim—suggesting the effect at hand may not replicate. This need not imply that the authors were intent on demonstrating a failure to replicate. Instead, the special issue authors may have been reluctant to propose a replication for well-established effects such as confirmation bias and social exclusion. Implicitly or explicitly, the authors may have felt that their time and effort would be spent more wisely on effects whose replication success was more uncertain.

Another argument against overgeneralizing the present results is that “No single replication provides the definitive word for or against the reality of an effect, just as no original study provides definitive evidence for it.” [1], p. 139. Indeed, when the replicability of specific findings is under scrutiny, a “many-labs” approach is preferable (e.g., [54, 63, 64, 65, 66, 67]). Finally, one always has to keep in mind that “different results between original and replication research could mean that there are unknown moderators or boundary conditions that differentiate the two studies. As such, the replication can raise more questions than it answers.” [1], p. 138.

The skeptic might retort that the results from the special issue are consistent with those from the Reproducibility Project: Psychology [62] which featured a more random selection of studies in social psychology. Moreover, as the editors acknowledged, “This special issue contains several replications of textbook studies, sometimes with surprising results (Nauts, Langner, Huijsmans, Vonk, & Wigboldus, 2014; Sinclair, Hood, & Wright, 2014; Vermeulen, Batenburg, Beukeboom, & Smits, 2014; Wesselmann et al., 2014).” [1], p. 139. Indeed, Figure 9 confirms that these studies cluster at the bottom, indicating that they provide evidence against the textbook effect. Lastly, even when the effects do not generalize to the field as a whole, the skeptic may argue that the general impression is still cause for concern.

In sum, we believe that it is premature, imprudent, and unwarranted to generalize the pattern of results obtained from our Bayesian reanalysis for the studies from the *Social Psychology* special

issue to all of social psychology, or even to a subdiscipline of social psychology. At the same time, it is also unwise to ignore the general message, which is that previously published results—even in psychology textbooks—may not replicate to the degree that one may hope. Cornerstone research demands careful replication, and the fact that previous research once “found” the effect (e.g.,  $N = 20$ ,  $p = .04$ ) is no reason to put blind faith in the result and consider it a proven fact of life (e.g., [68], their Figure 4).

## (b) Alternative Statistical Analyses

Our analysis efforts have demonstrated the ease with which default Bayesian analyses can be carried out using readily available software such as JASP, Stan, and the BayesFactor package in R. In these default or “objective” analyses, the prior distribution on effect size under  $\mathcal{H}_1$  is relatively broad and centered on zero [e.g., 20, 21, 69].

An alternative “subjective” Bayesian procedure is to use substantive knowledge and assign effect size under  $\mathcal{H}_1$  a more informative prior distribution. Compared to the default analyses, these subjective prior distributions are likely to be less spread out and are likely not to be centered on zero [e.g., 70].<sup>10</sup> The challenges and advantages of a subjective Bayesian analysis are beyond the scope of this article.

In general, there exist additional Bayesian methods to assess the degree of replication success [e.g., 47]. For instance, one may compare the effect size of the replication study to that of the original study [71, 72], or one may use the information about effect size from the original study to set up a prior distribution for the hypothesis test in the replication attempt (e.g., [47, 67]). We decided to report the results from the current set of methods because these methods are standard, fully developed for the tests of interest, and easy to extend to research efforts that do not focus on replication.

Finally, it should be mentioned that the special issue included statistical methodology that is not yet available in its complete Bayesian form (e.g., [52]). In order to take full advantage of the Bayesian paradigm, it is imperative that Bayesian procedures are developed for the run-of-the-mill scenarios that confront researchers every day.

In closing, we believe that our Bayesian bird’s eye view has provided an unambiguous overview of the results from the contributions to the *Social Psychology* special issue on “Replications of Important Results in Social Psychology”. This overview may motivate the field to take measures that ensure that published findings replicate at a higher rate than they do now [e.g., 73, 74]: a stronger focus on replication studies, more use of high-powered designs, standard adoption of preregistration, and data sharing by default [75]. We also hope that future analyses of replication studies will be more inclusive by employing a range of different, complementary techniques. When different statistical procedures support the same inference this can only serve to reinforce one’s confidence in the robustness and validity of the results.

## Acknowledgment

This work was supported by the ERC grant “Bayes or Bust!” from the European Research Council. This project was feasible only because the data from the *Social Psychology* special issue are openly available on the Open Science Framework (<https://osf.io/>). The code for all of our analyses is available at <https://osf.io/bqwzd/>. We thank Don van der Bergh for his assistance with the robustness checks.

## References

- 1 B. A. Nosek and D. Lakens. Registered reports: A method to increase the credibility of published results. *Social Psychology*, 45:137–141, 2014.

<sup>10</sup>Note that it is imprudent to change the default analysis by drastically lowering the width of the prior distribution while keeping it centered at zero; such a change makes  $\mathcal{H}_1$  similar to  $\mathcal{H}_0$ , and all but guarantees that the data are not diagnostic.

- 2 B. A. Nosek and D. Lakens. Call for proposals: Special issue of *social psychology* on “replications of important results in social psychology”. *Social Psychology*, 44:59–60, 2013.
- 3 K. S. Button, J. P. A. Ioannidis, C. Mokrysz, B. A. Nosek, J. Flint, E. S. J. Robinson, and M. R. Munafò. Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14:1–12, 2013.
- 4 C. D. Chambers. Registered Reports: A new publishing initiative at Cortex. *Cortex*, 49:609–610, 2013.
- 5 E.-J. Wagenmakers, R. Wetzels, D. Borsboom, H. L. J. van der Maas, and R. A. Kievit. An agenda for purely confirmatory research. *Perspectives on Psychological Science*, 7:627–633, 2012.
- 6 E.-J. Wagenmakers, A. J. Verhagen, A. Ly, M. Bakker, M. D. Lee, D. Matzke, J. N. Rouder, and R. D. Morey. A power fallacy. *Behavior Research Methods*, 47:913–917, 2015.
- 7 E.-J. Wagenmakers, A. J. Verhagen, and A. Ly. How to quantify the evidence for the absence of a correlation. *Behavior Research Methods*, 48:413–426, 2016.
- 8 A. Gelman, J. Hwang, and A. Vehtari. Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24:997–1016, 2014.
- 9 E.-J. Wagenmakers, M. Marsman, T. Jamil, A. Ly, A. J. Verhagen, J. Love, R. Selker, Q. F. Gronau, M. Šmíra, S. Epskamp, D. Matzke, J. N. Rouder, and R. D. Morey. Bayesian statistical inference for psychological science. Part I: Theoretical advantages and practical ramifications. *Accepted pending minor revision, Psychonomic Bulletin & Review*, 2016.
- 10 JASP Team. JASP (Version 0.8)[Computer software], 2016. URL <https://jasp-stats.org/>.
- 11 B. Carpenter, A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancour, M.A. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, in press.
- 12 Stan Development Team. *Stan Modeling Language User’s Guide and Reference Manual, Version 2.8.0*, 2015. URL <http://mc-stan.org/>.
- 13 M.D. Hoffman and A. Gelman. The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15:1351–1381, 2014.
- 14 Richard D. Morey and Jeffrey N. Rouder. BayesFactor 0.9.11-1. Comprehensive R Archive Network, 2015. URL <http://cran.r-project.org/web/packages/BayesFactor/index.html>.
- 15 E.-J. Wagenmakers, R. D. Morey, and M. D. Lee. Bayesian benefits for the pragmatic researcher. *Current Directions in Psychological Science*, 25:169–176, 2016.
- 16 R. D. Morey, J. W. Romeijn, and J. N. Rouder. The philosophy of Bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, 72:6–18, 2016.
- 17 J. N. Rouder, R. D. Morey, A. J. Verhagen, A. R. Swagman, and E.-J. Wagenmakers. Bayesian analysis of factorial designs. *Psychological Methods*, in press.
- 18 R. Hoekstra, R. D. Morey, J. N. Rouder, and E.-J. Wagenmakers. Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review*, 21:1157–1164, 2014.
- 19 R. D. Morey, R. Hoekstra, J. N. Rouder, M. D. Lee, and E.-J. Wagenmakers. The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review*, 23:103–123, 2016.
- 20 M. J. Bayarri, J. O. Berger, A. Forte, and G. García-Donato. Criteria for Bayesian model choice with application to variable selection. *The Annals of Statistics*, 40:1550–1577, 2012.
- 21 H. Jeffreys. *Theory of Probability*. Oxford University Press, Oxford, UK, 3 edition, 1961.
- 22 F. Liang, R. Paulo, G. Molina, M. A. Clyde, and J. O. Berger. Mixtures of  $g$  priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103:410–423, 2008.
- 23 J. N. Rouder, P. L. Speckman, D. Sun, R. D. Morey, and G. Iverson. Bayesian  $t$  tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16:225–237, 2009.
- 24 J. N. Rouder, R. D. Morey, P. L. Speckman, and J. M. Province. Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56:356–374, 2012.
- 25 A. Zellner. On assessing prior distributions and Bayesian regression analysis with  $g$ -prior distributions. In P. Goel and A. Zellner, editors, *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, pages 233–243. North-Holland, Amsterdam, 1986.

- 26 A. Zellner and A. Siow. Posterior odds ratios for selected regression hypotheses. In J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, editors, *Bayesian Statistics*, pages 585–603. University Press, Valencia, 1980.
- 27 H. Cramér. *Mathematical Methods of Statistics*. Asia Publishing House, Bombay, 1962.
- 28 A. Gelman and I. Pardoe. Bayesian measures of explained variance and pooling in multilevel (hierarchical) models. *Technometrics*, 48(2):241–251, 2006.
- 29 A. Gelman and J. Hill. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, Cambridge, 2007.
- 30 D. V. Lindley and A. F. M. Smith. Bayes estimates for the linear model. *Journal of the Royal Statistical Society Series B (Methodological)*, 34:1–41, 1972.
- 31 R. M. Shiffrin, M. D. Lee, W. Kim, and E.-J. Wagenmakers. A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32:1248–1284, 2008.
- 32 B. Efron and C. Morris. Stein’s paradox in statistics. *Scientific American*, 236:119–127, 1977.
- 33 H. IJzerman, I. Blanken, M. J. Brandt, J. M. Oerlemans, M. M. W. van der Hoogenhof, S. J. Franken, and M. W. G. Oerlemans. Sex differences in distress from infidelity in early adulthood and in later life: A replication and meta-analysis of Shackelford et al. (2004). *Social Psychology*, 45:202–208, 2014.
- 34 D. J. Johnson, F. Chueng, and M. B. Donnellan. Does cleanliness influence moral judgments? A direct replication of Schnall, Benton, and Harvey (2008). *Social Psychology*, 45(3):209–215, 2014.
- 35 D. van Ravenzwaaij, C. Donkin, and J. Vandekerckhove. The EZ diffusion model provides a powerful test of simple empirical effects. *Psychonomic Bulletin & Review*, in press.
- 36 E.-J. Wagenmakers, P. Grünwald, and M. Steyvers. Accumulative prediction error and the selection of time series models. *Journal of Mathematical Psychology*, 50:149–166, 2006.
- 37 J. Cohen. The earth is round ( $p < .05$ ). *American Psychologist*, 49:997–1003, 1994.
- 38 P. E. Meehl. Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46:806–834, 1978.
- 39 W. Edwards, H. Lindman, and L. J. Savage. Bayesian statistical inference for psychological research. *Psychological Review*, 70:193–242, 1963.
- 40 A. Gelman and D. B. Rubin. Avoiding model selection in bayesian social research. *Sociological Methodology*, 25:165–173, 1995.
- 41 A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis (2nd ed.)*. Chapman & Hall/CRC, Boca Raton (FL), 2004.
- 42 C. P. Robert. The expected demise of the Bayes factor. *Journal of Mathematical Psychology*, 72:33–37, 2016.
- 43 R. E. Kass and A. E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90:773–795, 1995.
- 44 W. Boekel, E.-J. Wagenmakers, L. Belay, A. J. Verhagen, S. D. Brown, and B. Forstmann. A purely confirmatory replication study of structural brain-behavior correlations. *Cortex*, in press.
- 45 Z. Dienes. Bayesian versus orthodox statistics: Which side are you on? *Perspectives on Psychological Science*, 6:274–290, 2011.
- 46 Z. Dienes. Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5:781, 2014.
- 47 A. J. Verhagen and E.-J. Wagenmakers. Bayesian tests to quantify the result of a replication attempt. *Journal of Experimental Psychology: General*, 143:1457–1475, 2014.
- 48 T. M. Shackelford, M. Voracek, D. P. Schmitt, D. M. Buss, V. A. Weekes-Shackelford, and R. L. Michalski. Romantic jealousy in early adulthood and in later life. *Human Nature*, 15, 2004.
- 49 R. D. Morey and E.-J. Wagenmakers. Simple relation between Bayesian order-restricted and point-null hypothesis tests. *Statistics and Probability Letters*, 92:121–124, 2014.
- 50 S. Roberts and H. Pashler. How persuasive is a good fit? A comment on theory testing in psychology. *Psychological Review*, 107:358–367, 2000.
- 51 A. Ly, A. J. Verhagen, and E.-J. Wagenmakers. Harold Jeffreys’s default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72:19–32, 2016.

- 52 S. Nauts, O. Langner, I. Huijsmans, R. Vonk, and D.H.J. Wigboldus. Forming impressions of personality: A replication and review of Asch's (1946) evidence for a primacy-of-warmth effect in impression formation. *Social Psychology*, 45(3):153–163, 2014.
- 53 A. Benavoli, G. Corani, F. Mangili, M. Zaffalon, and F. Ruggeri. A bayesian wilcoxon signed-rank test based on the dirichlet process. In T. Jebara and E.P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1026–1034. JMLR Workshop and Conference Proceedings, 2014. URL <http://jmlr.org/proceedings/papers/v32/benavoli14.pdf>.
- 54 R.A. Klein, K.A. Ratliff, M. Vianello, R. B. Adams, Jr., Š. Bahník, M.J. Bernstein, K. Bocian, M.J. Brandt, B. Brooks, C.C. Brumbaugh, Z. Cemalčiar, J. Chandler, W. Cheong, W.E. Davis, T. Devos, M. Eisner, N. Frankowska, D. Furrow, E. M. Galliani, F. Hasselman, J.A. Hicks, J.F. Hovermale, S.J. Hunt, J.R. Huntsinger, H. IJzerman, M.-S. John, J.A. Joy-Gaba, H. Barry Kappes, L.E. Krueger, J. Kurtz, C.A. Levitan, R.K. Mallett, W.L. Morris, A.J. Nelson, J.A. Nier, G. Packard, R. Pilati, A.M. Rutchick, K. Schmidt, J.L. Skorinko, R. Smith, T.G. Steiner, J. Storbeck, L.M. Van Swol, D. Thompson, A.E. van 't Veer, L.A. Vaughn, M. Vranka, A.L. Wichman, J.A. Woodzicka, and B.A. Nosek. Investigating variation in replicability: A "many labs" replication project. *Social Psychology*, 45:142–152, 2014.
- 55 A. Etz and J. Vandekerckhove. A Bayesian perspective on the Reproducibility Project: Psychology. *PLoS ONE*, 11:e0149794, 2016.
- 56 H.C. Sinclair, K.B. Hood, and B.L. Wright. Revisiting the Romeo and Juliet effect (Driscoll, Davis, & Lipetz, 1972): Reexamining the links between social network opinions and romantic relationship outcomes. *Social Psychology*, 45:170–178, 2014.
- 57 F. Müller and K. Rothermund. What does it take to activate stereotypes? simple primes don't seem enough: A replication of stereotype activation (Banaji & Hardin, 1996; Blair & Banaji, 1996). *Social Psychology*, 45(3):187–193, 2014.
- 58 E.D. Wesselmann, K.D. Williams, J.B. Pryor, F.A. Eichler, D.M. Gill, and J.D. Hogue. Revisiting Schachter's research on rejection, deviance, and communication (1951). *Social Psychology*, 45(3):164–169, 2014.
- 59 I. Vermeulen, A. Batenburg, C.J. Beukeboom, and T. Smits. Breakthrough or one-hit wonder? Three attempts to replicate single-exposure musical conditioning effects on choice behavior (Gorn, 1982). *Social Psychology*, 45(3):179–186, 2014.
- 60 H. Pashler and E.-J. Wagenmakers. Editors' introduction to the special section on replicability in psychological science: A crisis of confidence? *Perspectives on Psychological Science*, 7:528–530, 2012.
- 61 M. Baker. Is there a reproducibility crisis? *Nature*, 533:452–454, 2016.
- 62 The Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*, 349:aac4716, 2015.
- 63 C.R. Ebersole, O.E. Atherton, A.L. Belanger, H.M. Skulborstad, J.M. Allen, J.B. Banks, E. Baranski, M.J. Bernstein, D.B.V. Bonfiglio, L. Boucher, E.R. Brown, N.I. Budiman, A.H. Cairo, C.A. Capaldi, C.R. Chartier, J.M. Chung, D.C. Cicero, J.A. Coleman, J.G. Conway, W.E. Davis, T. Devos, M.M. Fletcher, K. German, J.E. Grahe, A.D. Hermann, J.A. Hicks, N. Honeycutt, B. Humphrey, M. Janus, D.J. Johnson, J.A. Joy-Gaba, H. Juzeler, A. Keres, D. Kinney, J. Kirshenbaum, R.A. Klein, R.E. Lucas, C.J.N. Lustgraaf, D. Martin, M. Menon, M. Metzger, J.M. Moloney, P.J. Morse, R. Prislín, T. Razza, D.E. Re, N.O. Rule, D.F. Sacco, K. Sauerberger, E. Shrider, M. Shultz, C. Siemsen, K. Sobocko, R.W. Sternglanz, A. Summerville, K.O. Tskhay, Z. van Allen, L.A. Vaughn, R.J. Walker, A. Weinberg, J.P. Wilson, J.H. Wirth, J. Wortman, and B.A. Nosek. Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67:68–82, 2016.
- 64 V. K. Alogna, M. K. Attaya, P. Aucoin, Š. Bahník, S. Birch, A. R. Birt, B. H. Bornstein, S. Bouwmeester, M. A. Brandimonte, C. Brown, K. Buswell, C. Carlson, M. Carlson, S. Chu, A. Cislak, M. Colarusso, M. F. Colloff, K. S. Dellapaolera, J.-F. Delvenne, A. Di Domenico, A. Drummond, G. Echterhoff, J. E. Edlund, C. M. Eggleston, B. Fairfield, G. Franco, F. Gabbert, B. W. Gamblin, M. Garry, R. Gentry, E. A. Gilbert, D. L. Greenberg, J. Halberstadt, L. Hall,

- P. J. B. Hancock, D. Hirsch, G. Holt, J. C. Jackson, J. Jong, A. Kehn, C. Koch, R. Kopietz, U. Körner, M. A. Kunar, C. K. Lai, S. R. H. Langton, F. P. Leite, N. Mammarella, J. E. Marsh, K. A. McConaughy, S. McCoy, A. H. McIntyre, C. A. Meissner, R. B. Michael, A. A. Mitchell, M. Mugayar-Baldocchi, R. Musselman, C. Ng, A. L. Nichols, N. L. Nunez, M. A. Palmer, J. E. Pappagianopoulos, M. S. Petro, C. R. Poirier, E. Portch, M. Rainsford, A. Rancourt, C. Romig, E. Rubínová, M. Sanson, L. Satchell, J. D. Sauer, K. Schweitzer, J. Shaheed, F. Skelton, G. A. Sullivan, K. J. Susa, J. K. Swanner, W. B. Thompson, R. Todaro, J. Ulatowska, T. Valentine, P. P. J. L. Verkoeijen, M. Vranka, K. A. Wade, C. A. Was, D. Weatherford, K. Wiseman, T. Zaksaitė, D. V. Zuj, and R. A. Zwaan. Registered Replication Report: Schooler and Engstler-Schooler (1990). *Perspectives on Psychological Science*, (5):556–578, 2014.
- 65 A. Eerland, A. M. Sherrill, J. P. Magliano, R. A. Zwaan, J. D. Arnal, P. Aucoin, S. A. Berger, A. R. Birt, N. Capezza, M. Carlucci, C. Crocker, T. R. Ferretti, M. R. Kibbe, M. M. Knepp, C. A. Kurby, J. M. Melcher, S. W. Michael, C. Poirier, and J. M. Prenoveau. Registered Replication Report: Hart & Albarracín (2011). *Perspectives on Psychological Science*, 11:158–171, 2016.
- 66 I. Cheung, L. Campbell, and E. P. LeBel. Registered Replication Report: Study 1 from Finkel, Rusbult, Kumashiro, & Hannon (2002). *Perspectives on Psychological Science*, 11:750–764, 2016.
- 67 E.-J. Wagenmakers, T. Beek, L. Dijkhoff, Q. F. Gronau, A. Acosta, R. B., Jr. Adams, D. N. Albohn, E. S. Allard, S. D. Benning, E.-M. Blouin-Hudon, L. C. Bulnes, T. L. Caldwell, R. J. Calin-Jageman, C. A. Capaldi, N. S. Carfagno, K. T. Chasten, A. Cleeremans, L. Connell, J. M. DeCicco, K. Dijkstra, A. H. Fischer, F. Foroni, U. Hess, K. J. Holmes, J. L. H. Jones, O. Klein, C. Koch, S. Korb, P. Lewinski, J. D. Liao, S. Lund, J. Lupiáñez, D. Lynott, C. N. Nance, S. Oosterwijk, A. A. Özdoğru, A. P. Pacheco-Unguetti, B. Pearson, C. Powis, S. Riding, T.-A. Roberts, R. I. Rumiati, M. Senden, N. B. Shea-Shumsky, K. Sobocko, J. A. Soto, T. G. Steiner, J. M. Talarico, Z. M. van Allen, M. Vandekerckhove, B. Wainwright, J. F. Wayand, R. Zeelenberg, E. E. Zetzer, and R. A. Zwaan. Registered Replication Report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, in press.
- 68 S. S. Young and A. Karr. Deming, data and observational studies: A process out of control and needing fixing. *Significance*, 8:116–120, 2011.
- 69 J. Berger. The case for objective Bayesian analysis. *Bayesian Analysis*, 1:1–17, 2004.
- 70 V. E. Johnson and D. Rossell. On the use of non-local prior densities in Bayesian hypothesis tests. *Journal of the Royal Statistical Society, Series B*, 72:143–170, 2010.
- 71 M. J. Bayarri and A. M. Mayoral. Bayesian analysis and design for comparison of effect-sizes. *Journal of Statistical Planning and Inference*, 103:225–243, 2002.
- 72 M. J. Bayarri and A. M. Mayoral. Bayesian design of “successful” replications. *The American Statistician*, 56:207–214, 2002.
- 73 D. S. Lindsay. Replication in psychological science. *Psychological Science*, 26:1827–1832, 2015.
- 74 B. A. Nosek, G. Alter, G. C. Banks, D. Borsboom, S. D. Bowman, S. J. Breckler, S. Buck, C. D. Chambers, G. Chin, G. Christensen, M. Contestabile, A. Dafoe, E. Eich, J. Freese, R. Glennerster, D. Goroff, D. P. Green, B. Hesse, M. Humphreys, J. Ishiyama, D. Karlan, A. Kraut, A. Lupia, P. Mabry, T. A. Madon, N. Malhotra, E. Mayo-Wilson, M. McNutt, E. Miguel, E. Levy Paluck, U. Simonsohn, C. Soderberg, B. A. Spellman, J. Turitto, G. VandenBos, S. Vazire, E.-J. Wagenmakers, R. Wilson, and T. Yarkoni. Promoting an open research culture. *Science*, 348:1422–1425, 2015.
- 75 R. D. Morey, C. D. Chambers, P. J. Etchells, C. R. Harris, R. Hoekstra, D. Lakens, S. Lewandowsky, C. Coker Morey, D. P. Newman, F. D. Schönbrodt, W. Vanpaemel, E.-J. Wagenmakers, and R. A. Zwaan. The peer reviewers’ openness initiative: Incentivizing open research practices through peer review. *Royal Society Open Science*, 3:150547, 2016.

## Research ethics

We have not needed or received ethical approval from a local ethics committee to carry out our study.

## Animal ethics

We have not needed or received licenses or approvals from our Animal Ethics committee.

## Permission to carry our fieldwork

We have not needed or received permission to carry out fieldwork.

## Competing interests

The authors declare no competing interests.

## Data accessibility

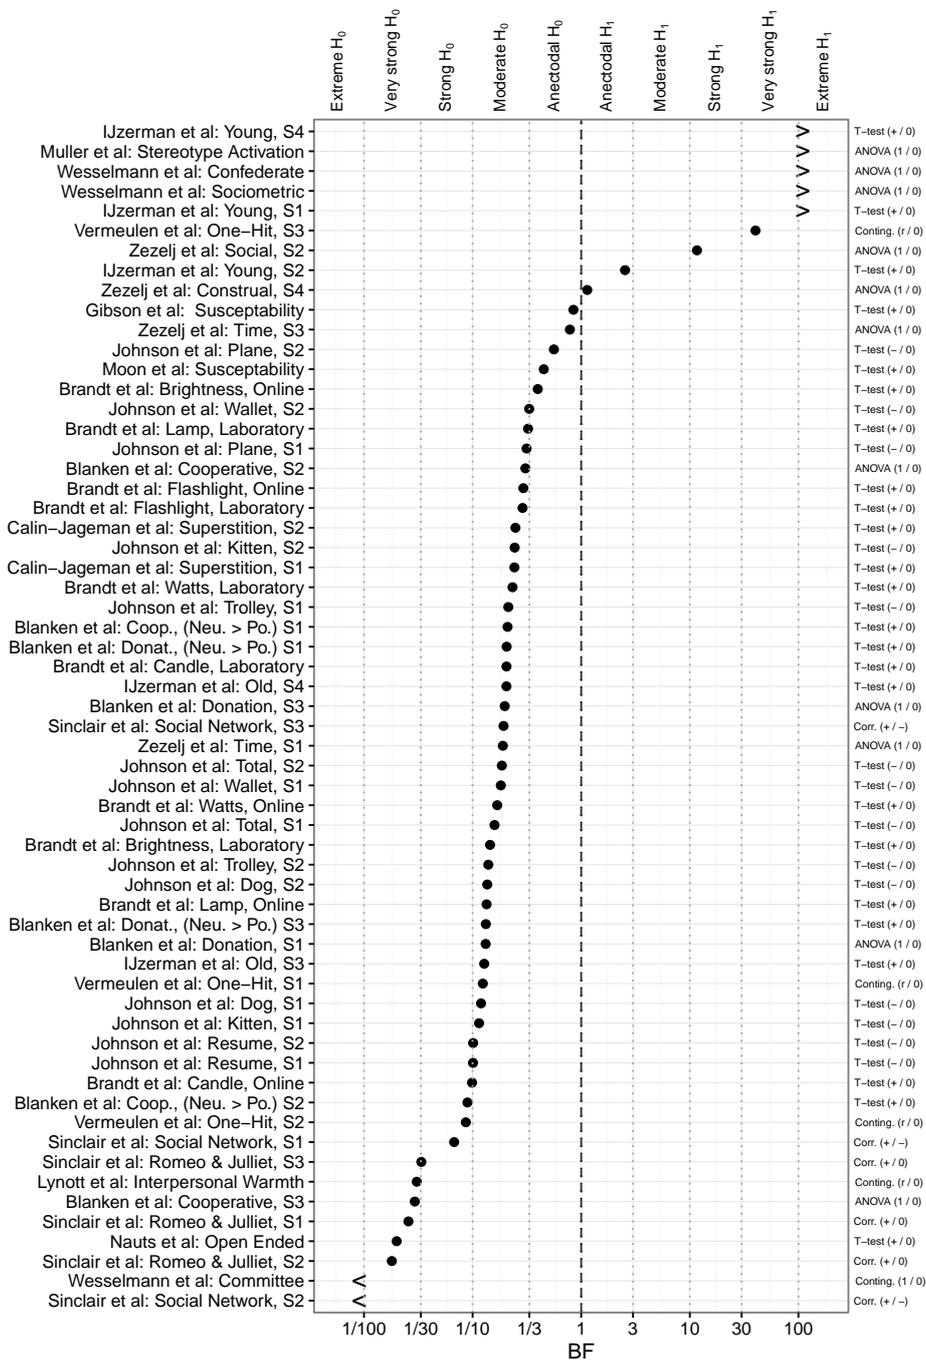
Our analyses and data files are available at <https://osf.io/bqwzd/>.

## Authors' contributions

All authors conceived and designed the statistical methods and analyses. MM, FS, RM and YY performed the statistical analyses. MM and EW drafted the manuscript. All authors gave final approval for publication.

## Funding

The work of MM and EW was supported by the ERC grant "Bayes or Bust" from the European Research Council (no. 283876).



**Figure 9.** Default Bayes factors for 60 analyses reported in the *Social Psychology* special issue. The direction of the hypothesis is indicated in the right margin: “(1 / 0)” refers to the two-sided  $BF_{10}$ , “(+ / 0)” refers to  $BF_{+0}$ , “(- / 0)” refers to the one-sided  $BF_{-0}$ , “(+ / -)” refers to the one-sided  $BF_{+-}$ , and “(r / 0)” refers to  $BF_{r0}$  where  $\mathcal{H}_r$  refers to a predicted ordering of parameters. The top margin indicates the evidence categories proposed by Jeffreys [21]. Figure available at <https://flic.kr/p/DZBTBj>, under CC license <https://creativecommons.org/licenses/by/2.0/>