A Bayesian Model-Averaged Meta-Analysis of the Power Pose Effect with Informed and Default Priors: The Case of Felt Power

Quentin F. Gronau[1], Sara van Erp[2], Daniel W. Heck[3], Joseph Cesario[4], Kai Jonas[5], &

Eric-Jan Wagenmakers[1]

1) University of Amsterdam

2) Tilburg University

3) University of Mannheim

4) Michigan State University

5) Maastricht University

Corresponding Author:

Quentin F. Gronau, University of Amsterdam, Nieuwe Achtergracht 129 B, 1018 WT

Amsterdam, The Netherlands. E-mail may be sent to Quentin.F.Gronau@gmail.com.

Abstract

Carney, Cuddy, and Yap (2010) found that --compared to participants who adopted constrictive body postures-- participants who adopted expansive body postures reported feeling more powerful, showed an increase in testosterone and a decrease in cortisol, and displayed an increased tolerance for risk. However, these power pose effects have recently come under considerable scrutiny. Here we present a Bayesian meta-analysis of six preregistered studies from this special issue, focusing on the effect of power posing on felt power. Our analysis improves on standard classical meta-analyses in several ways. First and foremost, we considered only preregistered studies, eliminating concerns about publication bias. Second, the Bayesian approach enables us to quantify evidence for both the alternative and the null hypothesis. Third, we use Bayesian model-averaging to account for the uncertainty with respect to the choice for a fixed-effect model or a random-effect model. Fourth, based on a literature review we obtained an empirically informed prior distribution for the between-study heterogeneity of effect sizes. This empirically informed prior can serve as a default choice not only for the investigation of the power pose effect, but for effects in the field of psychology more generally. For effect size, we considered a default and an informed prior. Our meta-analysis yields very strong evidence for an effect of power posing on felt power. However, when the analysis is restricted to participants unfamiliar with the effect, the meta-analysis yields evidence that is only moderate.

Introduction

Could adopting a powerful body posture make us more powerful? Carney, Cuddy, and Yap (2010) found that participants who adopted expansive, high-power body postures (Figure 1, top row) as opposed to constrictive, low-power body postures (Figure 1, bottom row) reported feeling more powerful and in charge, showed an increase in testosterone and a decrease in cortisol, and displayed an increased tolerance for risk. The power-pose effect has attracted a lot of attention, partly due to the anticipated consequences for day-to-day life suggesting that it might be possible to "fake it 'til you make it".
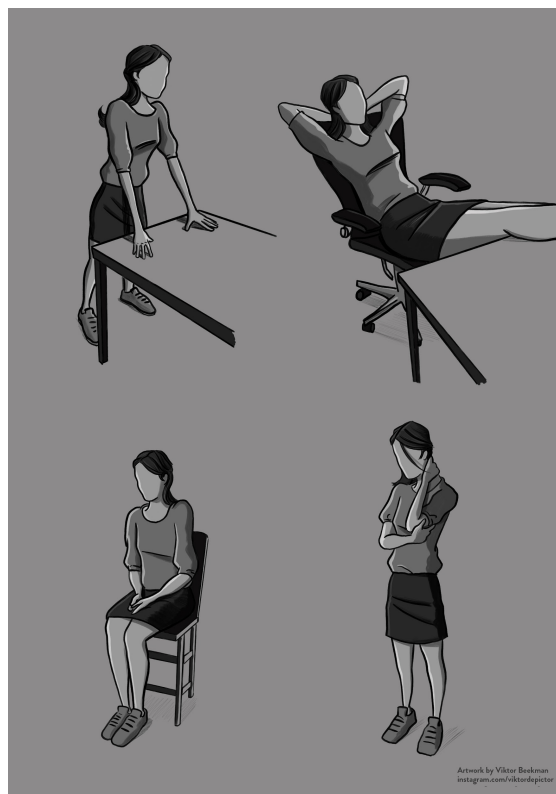
Figure 1: High-power poses (top row) and low-power poses (bottom row). CC-BY: Artwork by Viktor Beekman, commissioned by Eric-Jan Wagenmakers.

However, this power pose effect has recently come under scrutiny. When Ranehill et al. (2015) attempted to replicate the effect, they found—similar to the original study—that adopting high-power poses increased participants' self-reported feelings of power; nevertheless, they did not find an effect on testosterone or cortisol nor on behavioral measures such as risk taking. Carney, Cuddy, and Yap (2015) pointed out a number of methodological differences that they believe might have been the cause for the diverging results. Recently, Garrison, Tang, and Schmeichel (2016) conducted a preregistered replication and extension of the power pose

study, and they failed to identify an effect of power posing on risk taking behavior. Furthermore, in contrast to Ranehill et al. (2015), these authors did not find evidence for a power pose effect on subjective feelings of power.

In the present special issue, seven preregistered studies investigated the effect of power posing under various circumstances (i.e., Bailey, LaFrance, & Dovidio, this issue; Bombari, Schmid Mast, & Pulfrey, this issue; Jackson, Nault, Smart Richman, LaBelle, & Rohleder, this issue; Keller, Johnson, & Harder, this issue; Klaschinski, Schröder-Abé, & Schnabel, this issue; Latu, Duffy, Pardal, & Alger, this issue; Ronay, Tybur, van Huijstee, & Morssinkhof, this issue). Here we present a meta-analysis of the effect of power posing on self-reported felt power, which was included as a dependent variable in six of the seven studies in this special issue.

Our analysis improves upon classical analyses in several ways. First, we only consider a set of preregistered studies which comes with the advantage that publication bias can be ruled out a priori (cf. the concept of a *prospective meta-analysis* in medicine). Second, the Bayesian approach enables us to quantify evidence for both the alternative hypothesis *and* for the null hypothesis; note that this evidence can be seamlessly updated as future studies on the effect become available. Third, Bayesian model-averaging enables us to fully acknowledge uncertainty with respect to the choice of a fixed-effect or random-effect model; in the fixed-effect model, the effect is assumed to be identical across studies; in the random-effect model, the effect is assumed to vary across studies. Instead of adopting one model for inference and ignoring the other model entirely, we can weight the results of both models according to their posterior plausibilities. This yields a model-averaged measure of evidence and a model-averaged estimate for the meta-analytic effect size. Fourth, the Bayesian approach enables us to incorporate existing knowledge into our analysis (e.g., Rhodes, Turner, & Higgins, 2015). Based on an extensive literature review of meta-analyses in the field of psychology, we obtained an informed prior distribution for the between-study heterogeneity. This informed prior distribution can serve as an informed default not only for the investigation of the power pose effect in the present meta-analysis, but for the field of psychology more generally. For effect size we also consider an informed prior distribution based on knowledge about effect sizes in the field of psychology. As a robustness check with respect to the prior choice we show that qualitatively similar results are obtained when we instead use a default prior for the effect size parameter.

The outline of this article is as follows: first, we explain the details of our analysis. Second, we present the results of an extensive literature review that allowed us to specify an informed prior distribution for the between-study heterogeneity. Third, we present the results of the model-averaged Bayesian meta-analysis for two different prior choices for effect size. Finally, we investigate whether the results change when only participants unaware of the power pose effect are included in the analysis.

Method

In our meta-analysis, we focused on the dependent variable *felt power* which was measured in all replication studies in the present issue except for the study by Jackson et al., which was therefore not considered in the analysis. We investigated the question whether felt power was higher in the high-power condition than in the low power condition.

Analysis of Individual Studies

When considering a single study, the power pose effect can be tested using a standard one-sided, independent-samples t-test. Hence, the first step in our analysis was to compute one-sided Bayesian t-tests (Rouder, Speckman, Sun, Morey, & Iverson, 2009; Ly, Verhagen, & Wagenmakers, 2016; Gronau, Ly, & Wagenmakers, 2017). This allowed us (1) to estimate for each study the posterior distribution of the standardized effect size that represents our beliefs about the effect size after having observed the data of that study and (2) to quantify the evidence that each study provides in favor of the hypothesis that the power pose effect is positive (H_+) versus the null hypothesis that the effect is zero (H_0).

To quantify the evidence that the data provide for or against H_+ we computed the Bayes factor (Jeffreys, 1961; Kass & Raftery, 1995) which is the predictive updating factor that quantifies how much the data have changed the relative plausibility of the competing models. The Bayes factor has an intuitive interpretation: when $BF_{+0} = 10$ this indicates that the data are ten times more likely under H_+ than under H_0; when $BF_{+0} = 1/5$ this indicates that the data are five times more likely under H_0 than under H_+.

Meta-Analysis

The next step in our analysis was to combine the studies with the help of a Bayesian meta-analysis (e.g., Marsman, Schönbrodt, Morey, Yao, Gelman, & Wagenmakers, 2017) to obtain an estimate of the overall effect size and to quantify the evidence for an effect that takes into account all studies simultaneously. In a classical meta-analysis the analyst has to make a choice between a fixed-effect and a random-effect model. A fixed-effect model makes the assumption that there is one underlying effect size so that the true effect in each study is identical; differences in the observed effect sizes are solely due to normally distributed sampling error. This can be formalized as follows: we assume that $y_i \sim N(\delta_{fixed}, SE_i^2)$, where y_i, $i = 1,2,...,n$ denotes the observed effect size in the i-th of n studies, SE_i denotes the corresponding standard error which is commonly assumed to be known, and δ_{fixed} corresponds to the common true effect size.

In contrast, a random-effect model allows for idiosyncratic study effects, that is, we no longer impose the constraint that there exists one common true effect size for all studies. The random study effects are usually assumed to follow a normal distribution with a mean equal to the overall effect size that we are interested in and a standard deviation that corresponds to the between-study heterogeneity. Note that analogously to the fixed-effect model, the model still incorporates random sampling error so that the observed effect size for a given study is not necessarily identical to the true effect size for that study. These assumptions yield a model with a hierarchical structure which can be formalized as follows: let δ_{random} denote the mean of the normal distribution of the study effects (i.e., the quantity that we are interested in), τ denote the standard deviation of that normal distribution (i.e., between-study heterogeneity), and θ_i denote the true study effect for the i-th study. Then, $\theta_i \sim N(\delta_{random}, \tau^2)$ and $y_i | \theta_i \sim N(\theta_i, SE_i^2)$. The structure of the model allows one to analytically integrate out the random study effects so that the model can equivalently be written as $y_i \sim N(\delta_{random}, \tau^2 + SE_i^2)$ which can be more convenient from a computational perspective.

Bayesian Model-Averaging

The choice of a fixed-effect or random-effect model commonly relies on a test for heterogeneity or on *a priori* considerations. Final inference is then based on either the fixed-effect or random-effect model. When the number of studies is small, this choice may be difficult; and in certain cases, the choice may be consequential. The Bayesian approach, however, allows a compromise solution: instead of selecting either a fixed-effect or random-effect model, we can use Bayesian model-averaging (e.g., Haldane, 1932; Hoeting, Madigan, Raftery, & Volinsky, 1999) and retain all models for final inference. Conclusions are then based on a combination of all models where the results of each model are taken into account according to the model's plausibility in light of the observed data. Concretely, Bayesian model-averaging allows us to obtain a model-averaged estimate for the meta-analytic effect size (Sutton & Abrams, 2001) and to quantify the overall evidence for an effect that considers both the fixed-effect and random-effect model (Scheibehenne, Gronau, Jamil, & Wagenmakers, 2017).

With respect to hypothesis testing, for the current analysis we entertained four models of interest, shown in Table 1: (1) the fixed-effect model H_+; (2) the fixed-effect model H_0 (i.e., δ_{fixed} = 0); (3) the random-effect model H_+; (4) the random-effect model H_0 (i.e., δ_{random} = 0). The fixed-effect meta-analytic Bayes factor was obtained by comparing case (1) to case (2); the random-effect meta-analytic Bayes factor pitched case (3) against case (4). To compute the model-averaged Bayes factor, we contrasted the summed posterior model probabilities (i.e., the probability of a model given the data) for cases (1) and (3) against the summed posterior model probabilities for cases (2) and (4). This assumes that all four models are equally likely *a priori*, a common assumption in model-averaging scenarios. In case the prior model probabilities were not identical, the ratio of the summed posterior model probabilities for cases (1) and (3) over (2) and (4) would need to be divided by a ratio obtained in a similar fashion but this time based on the prior model probabilities.

With respect to parameter estimation, we computed a model-averaged effect size estimate based on the four model versions described above, except that we no longer imposed the constraint that the effect size has to be positive. In other words, consistent with standard practice, we imposed a directional constraint for testing but not for estimation (cf. Jeffreys, 1961, who also used different priors for estimation and testing). This reflects the fact that the estimation framework is generally more exploratory in nature, and this mindset is inconsistent with the use of hard boundaries. The combined estimate was obtained by combining the estimates of models (1) and (3) --but without the order-constraints-- according to their posterior model probabilities. To conduct the model-averaged Bayesian meta-analysis, we used the R package *metaBMA* (Heck & Gronau, 2017) available from https://github.com/danheck/metaBMA.

Table 1. The four meta-analysis models included in the Bayesian model-averaging for hypothesis testing.

Hypotheses	Fixed-Effect Meta-Analysis	Random-Effect Meta-Analysis
H_0: No effect	Fixed overall effect size $\delta_{fixed} = 0$	Mean overall effect size $\delta_{random} = 0$ Study heterogeneity Study effect size θ_i ($i = 1,2,...,n$)
H_+: Positive effect	Fixed overall effect size δ_{fixed}	Mean overall effect size δ_{random} Study heterogeneity Study effect size θ_i ($i = 1,2,...,n$)

Prior Distributions

In the Bayesian approach, model parameters are assigned prior distributions that reflect the knowledge, uncertainty, or beliefs for the parameters before seeing the data. Using Bayes' theorem, these prior distributions are then updated by the data to yield posterior distributions, which reflect the uncertainty for the parameters after the data have been observed. Consequently, in order to conduct our Bayesian analyses, prior distributions were required for all model parameters.

For the standardized effect size, we considered two different prior choices. First, we used what has now become the default choice in the field of psychology, that is, a zero-centered Cauchy distribution with scale parameter equal to $1/\sqrt{2}$ (Morey & Rouder, 2015). Second, we considered the informed prior distribution reported in Gronau et al. (2017): a *t* distribution with location 0.350, scale 0.102, and three degrees of freedom, which is displayed in Figure 2. This prior distribution was elicited from Dr. Oosterwijk, a social psychologist at the University of Amsterdam, for a reanalysis of the Registered Replication Report on the facial feedback hypothesis (Wagenmakers et al., 2016). We believe this prior distribution is generally plausible for a wide range of small-to-medium effects in social psychology (i.e., for effects whose presence needs to be ascertained by statistical analysis). One could elicit a "power pose prior", but we believe the resulting distribution would be highly similar to the Oosterwijk prior, and therefore yield highly similar inferences. Researchers interested in using a specific "power pose prior" are invited to explore this option using the R code provided online (https://osf.io/r2cds/).

For the one-sided hypothesis tests, the priors were truncated at zero, that is, the model encoded the *a priori* assumption that negative effect sizes are impossible. For estimating the effect size, however, we removed this truncation. The informed and default priors are depicted in Figure 2. The informed prior expresses the belief that the effect size is positive but most likely small to medium in size. The default prior on the other hand is more spread out (i.e., less informative) and it is centered on zero. Figure 2 also illustrates how the priors were truncated at zero for testing whereas for estimation, this truncation was removed.
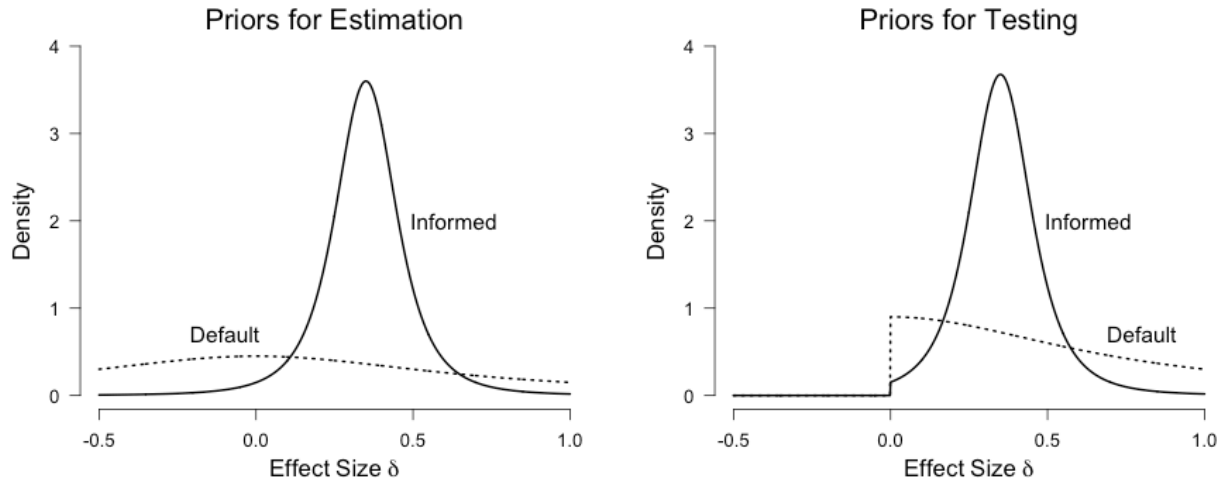
Figure 2: Depiction of the default and informed prior distribution for the standardized effect size. The default prior is a Cauchy distribution with scale $1/\sqrt{2}$, the informed prior is a *t* distribution with location 0.350, scale 0.102, and three degrees of freedom. Figure available at http://tinyurl.com/j9dthma under CC license https://creativecommons.org/licenses/by/2.0/.

In addition to the prior distribution for the effect size, the Bayesian meta-analysis required a prior distribution for the between-study heterogeneity. Here we chose an informed prior distribution for the between-study standard deviation τ. This informed prior was based on all available between-study heterogeneity estimates for mean-difference effect sizes in meta-analyses reported in *Psychological Bulletin* in the years 1990 to 2013 (van Erp, Verhagen, Grasman, & Wagenmakers, 2017, https://osf.io/preprints/psyarxiv/myu9c). The distribution of these 162 estimates is shown in Figure 3. Note that we have excluded between-study heterogeneity estimates that were exactly equal to zero, as the prior should reflect knowledge conditional on the assumption that the random-effect model is true; between-study heterogeneity estimates of exactly zero, however, suggest that the fixed-effect model was more appropriate. The distribution of the estimates in Figure 3 suggests that (1) the between-study standard deviations in the field of psychology range from 0 to 1 and (2) there are more small estimates than large ones. These two features are captured by an Inverse-Gamma(1, 0.15) distribution (depicted in Figure 3 as a solid line).[1] Note, however, that this prior distribution does not completely rule out the possibility that between-study heterogeneity is larger than 1; the distribution merely assigns values larger than 1 a relatively small prior credibility. This inverse-gamma distribution resembles the one obtained when maximum-likelihood methods are used to fit an inverse-gamma distribution to the between-study heterogeneity estimates. However, in our opinion, the maximum-likelihood inverse-gamma distribution slightly overemphasizes small

---

[1] For computational convenience, it is common practice to assign an inverse-gamma prior to the variance instead of to the standard deviation. Here we use the inverse-gamma as a convenient summary for the empirical distribution of the between-study heterogeneity estimates.

between-study heterogeneity values. In the appendix, we present the results obtained under two alternative prior choices for between-study heterogeneity: (1) the maximum-likelihood inverse-gamma distribution; and (2) a Beta(1, 2) prior distribution. The results are robust across all of these prior choices.
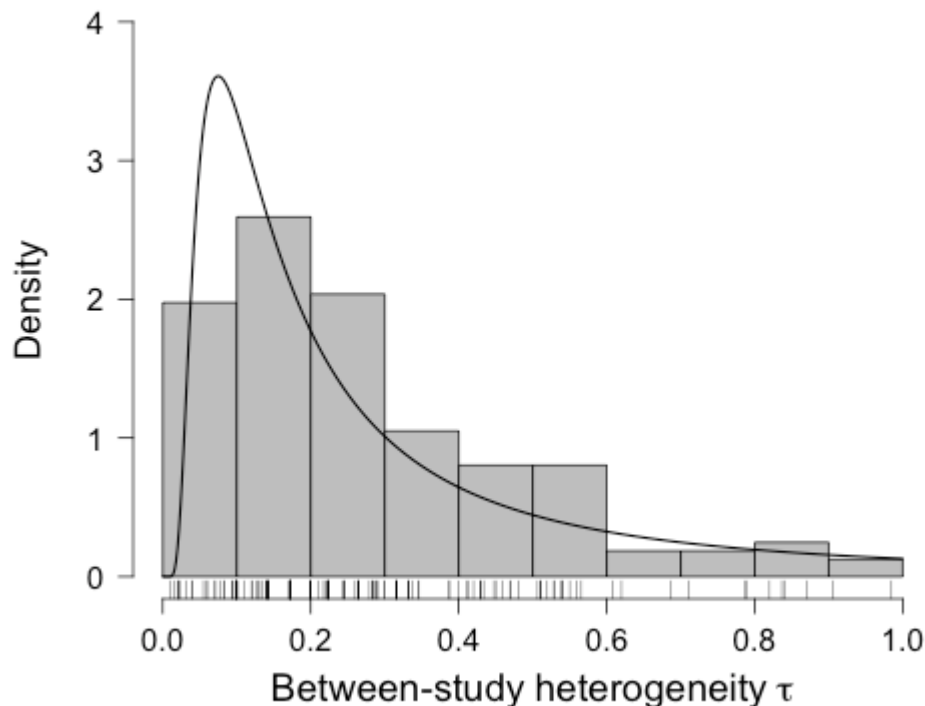

Figure 3: Distribution of the non-zero between-study standard deviations from meta-analyses reported in *Psychological Bulletin* (1990-2013; van Erp et al., 2017). The informed Inverse-Gamma(1, 0.15) prior distribution is displayed on top. Figure available at http://tinyurl.com/lwfa9rd under CC license https://creativecommons.org/licenses/by/2.0/.

Having specified the models and prior distributions, we needed to compute the probability of the data given each model under consideration. This was achieved by integrating out the model parameters with respect to their prior distributions. For the models for which this was not possible analytically, we evaluated this quantity using numerical integration as implemented in the R package *metaBMA* (Heck & Gronau, 2017). R code for reproducing all analyses can be found on the Open Science Framework: https://osf.io/r2cds/.[2]

Results

Analysis of Reported Studies: Default Prior on Effect Size
Figure 4 displays the results of the Bayesian analysis using the default effect size prior for the studies as reported in this special issue. Note that most studies did not exclude

---

[2] The R code also allows one to explore alternative prior choices easily.

participants who were familiar with the effect, for instance, from viewing the TED talk about power posing, which is currently the second most popular TED talk of all time (https://www.ted.com/playlists/171/the_most_popular_talks_of_all). This analysis is based on a total of 1071 participants. Below, we investigate how the results change when considering only those participants who indicated not to know the power pose effect. The upper part of Figure 4 displays the results of the Bayesian *t*-tests. The left-part of the figure displays for each study the median of the posterior distribution for the effect size (grey dots) and a 95% highest density interval (HDI; i.e., the shortest interval that captures 95% of the posterior mass). The right part of the figure shows the one-sided default Bayes factors in favor of H_+ and, for comparison, the (two-sided) *p*-values obtained from
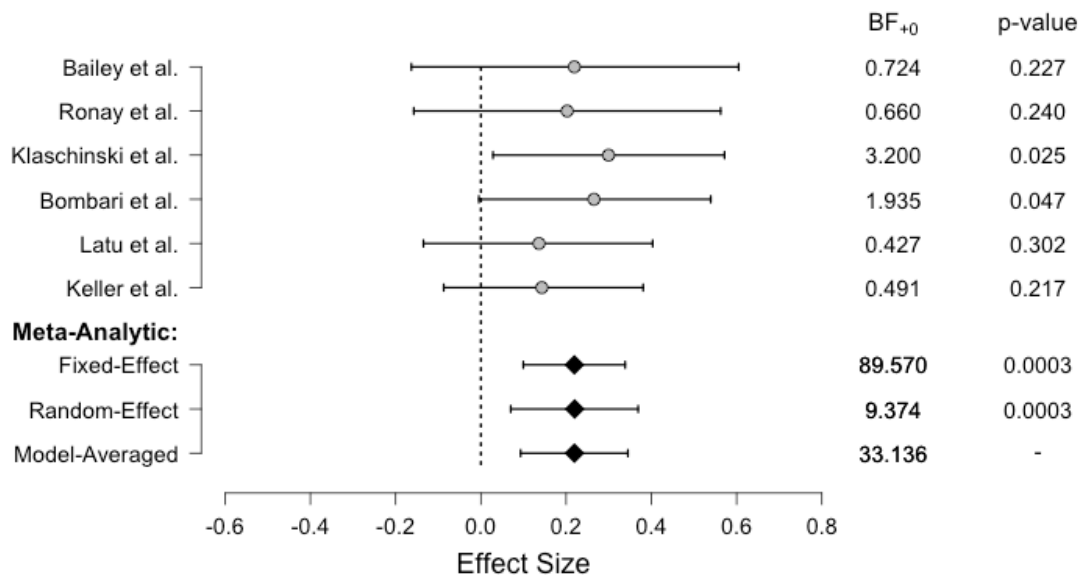

Figure 4: Bayesian model-averaged meta-analysis using the default Cauchy prior with scale $1/\sqrt{2}$ for the standardized effect size. The dots and diamonds correspond to the median of the posterior distribution for the effect size; the lines correspond to the 95% highest density intervals. The one-sided Bayes factors are displayed on the right, flanked by classical two-sided *p*-values. Figure available at http://tinyurl.com/kz2jpwb under CC license https://creativecommons.org/licenses/by/2.0/.

classical independent samples *t*-tests. Based on the posterior distributions, it appears that there might be a positive effect. However, this is hard to assess since the 95% highest density intervals are relatively wide. All Bayes factors except one are between ⅓ and 3 indicating that there is not much evidence for H_+ or H_0. Hence, when considering the individual studies separately, we cannot draw strong conclusions about whether there is an effect or not.

Each study alone does not provide much evidence in favor of either hypothesis; however, a Bayesian meta-analysis allows us to obtain an impression of the overall evidence obtained when considering all studies simultaneously. The lower part of Figure 4 displays the result of the Bayesian meta-analysis using the default Cauchy prior with scale $1/\sqrt{2}$ for the meta-analytic effect size. The black diamonds display the median of the posterior distribution of the meta-analytic effect size for the fixed-effect, random-effect, and model-averaged analysis,

and the lines correspond to the 95% highest density intervals. The model-averaged posterior distribution is obtained by combining the estimates of the fixed-effect and the random-effect model according to their plausibility in light of the data. The lower right part of Figure 4 shows the meta-analytic one-sided Bayes factors and, for the fixed-effect and the random-effect model, the two-sided *p*-value obtained by conducting classical meta-analyses. The meta-analytic fixed-effect Bayes factor equals BF_{+0} = 89.6, indicating very strong evidence in favor of an effect of power posing on felt power. The meta-analytic random-effect Bayes factor is less extreme but still indicates evidence for an effect: BF_{+0} = 9.4. The observed data support a fixed-effect model more than a random-effect model: the Bayes factor that compares case (1), fixed-effect H_+, to case (3), random-effect H_+, (not displayed) indicates that the data are 4.0 times more likely under the fixed-effect model than under the random-effect model. This is reflected in the model-averaged result: the meta-analytic model-averaged Bayes factor equals BF_{+0} = 33.1 indicating very strong evidence in favor of an effect of power posing on felt power. The median of the model-averaged meta-analytic effect size is equal to 0.22 [95% HDI: 0.09, 0.34].

To sum up, the Bayesian meta-analytic results based on the default prior for the effect size provide very strong evidence in favor of the hypothesis that power posing leads to an increase in felt power.

Analysis of Reported Studies: Informed Prior on Effect Size
Next, we consider the results based on the informed *t* prior distribution for the effect size with location 0.350, scale 0.102, and three degrees of freedom (cf. Figure 2). The results are displayed in Figure 5**.** The effect size posterior distributions for the individual studies clearly show the influence of the informed prior distribution: the posteriors are narrower and slightly shifted towards the location of the informed prior. The individual study one-sided informed Bayes factors are larger than the default ones. This can be explained by interpreting the Bayes factor as an assessment tool of the predictive success of two competing hypotheses. The informed alternative hypothesis makes much riskier predictions than the default alternative hypothesis; however, these risky predictions are rewarded because the observed effect sizes fall within the range of values predicted by the informed hypothesis. Hence, since the predictions match the observed data, the informed hypothesis yields more evidence for the presence of the power pose effect as compared to an alternative hypothesis that specifies a default prior for the effect size. Nevertheless, only two of the study-specific Bayes factors provide moderate evidence for an effect, whereas the other four provide only anecdotal evidence for H_+ or H_0.

The informed meta-analytic fixed-effect Bayes factor is BF_{+0} = 191.8 indicating extreme evidence in favor of an effect of power posing on felt power. The informed meta-analytic random-effect Bayes factor is less extreme but still indicates strong evidence for an effect: BF_{+0} = 20.7. As for the default prior, the observed data support a fixed-effect model more than a random-effect model, the Bayes factor that compares case (1), fixed-effect H_+, to case (3), random-effect H_+, (not displayed) indicates that the data are 3.9 times more likely under the fixed-effect model than under the random-effect model (not displayed). The informed meta-analytic model-averaged Bayes factor is equal to BF_{+0} = 71.4 indicating very strong evidence in favor of an effect of power posing on felt power. The median of the model-averaged meta-analytic effect size is similar to the default one and is equal to 0.26 [95% HDI: 0.14, 0.37].

To sum up, the Bayesian meta-analytic results based on the informed prior for the effect size provide very strong evidence in favor of the hypothesis that power posing leads to an increase in felt power. The informed analysis yields more evidence for an effect as compared to the default analysis indicating that the successful predictions of the informed hypothesis are rewarded.
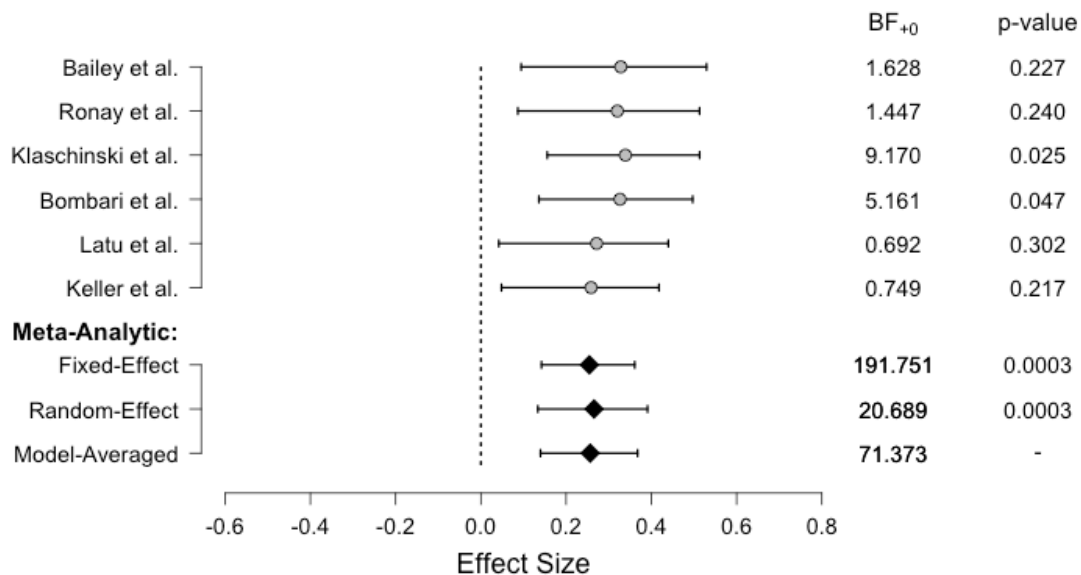

Figure 5: Bayesian model-averaged meta-analysis using the informed *t* prior with location 0.350, scale 0.102, and three degrees of freedom for the standardized effect size (depicted in Figure 2A). The dots and diamonds correspond to the median of the posterior distribution for the effect size; the lines correspond to the 95% highest density intervals. The one-sided Bayes factors are displayed on the right, flanked by classical two-sided *p*-values. Figure available at http://tinyurl.com/n8mwfsv under CC license https://creativecommons.org/licenses/by/2.0/.

Moderator Analysis: Knowledge of the Effect (Default Prior on Effect Size)

Next we investigate whether and how the results change when considering only participants who indicated to be unaware of the power posing effect. Hence, participants who could guess the goal of the study or were familiar with the power pose TED talk were excluded in all studies under consideration, leaving a total of 809 participants. Figure 6 displays the

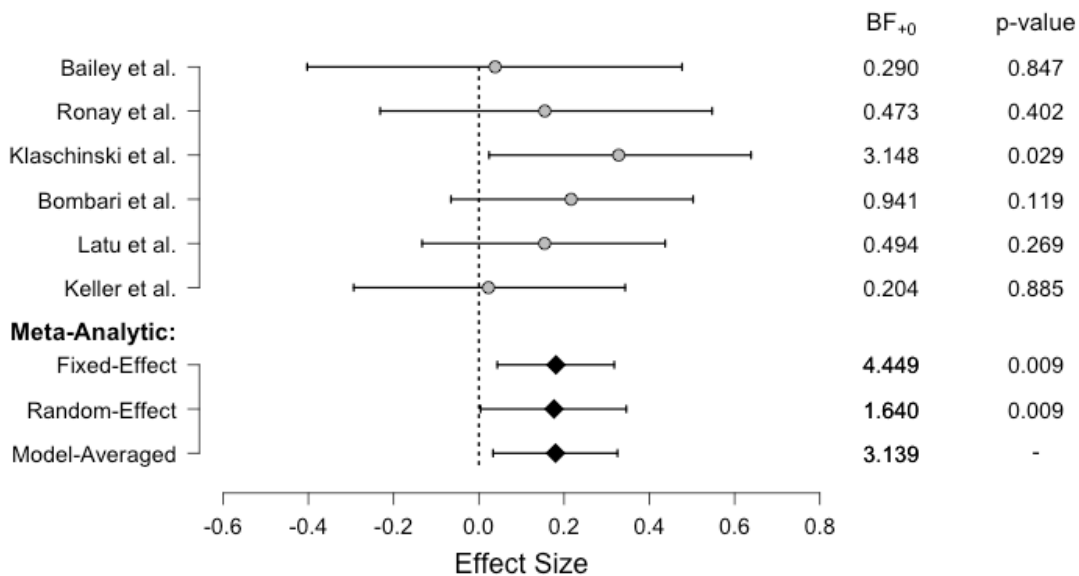results of the Bayesian analysis using the default effect size prior.

Figure 6: Bayesian model-averaged meta-analysis for the subset of participants unfamiliar with the effect using the default Cauchy prior with scale $1/\sqrt{2}$ for the standardized effect size. The dots and diamonds correspond to the median of the posterior distribution for the effect size; the lines correspond to the 95% highest density intervals. The one-sided Bayes factors are displayed on the right, flanked by classical two-sided *p*-values. Figure available at http://tinyurl.com/kmfcnhz under CC license https://creativecommons.org/licenses/by/2.0/.

Compared to Figure 4, the posterior distributions are shifted towards smaller values and the 95% highest density intervals are relatively wide (due to the reduced sample size). Three Bayes factors are between ⅓ and 3 indicating that there is little evidence for H_+ or H_0, one Bayes factor indicates moderate evidence for the alternative hypothesis, and two Bayes factors indicate moderate evidence for the null hypothesis. Hence, similar to the previous analysis, when considering the individual studies separately, we cannot draw strong conclusions about whether or not there is an effect.

The lower part of Figure 6 displays the result of the Bayesian meta-analysis using the default Cauchy prior with scale $1/\sqrt{2}$. The meta-analytic fixed-effect Bayes factor equals BF_{+0} = 4.4 indicating moderate evidence in favor of an effect of power posing on felt power. The meta-analytic random-effect Bayes factor equals BF_{+0} = 1.6 indicating only anecdotal evidence for the alternative hypothesis. The observed data support a fixed-effect model more than a random-effect model: the Bayes factor that compares case (1), fixed-effect H_+, to case (3), random-effect H_+, (not displayed) indicates that the data are 3.1 times more likely under the fixed-effect model than under the random-effect model. This is reflected in the model-averaged result: the meta-analytic model-averaged Bayes factor is equal to BF_{+0} = 3.1 indicating moderate evidence in favor of an effect of power posing on felt power. The median of the model-averaged meta-analytic effect size is equal to 0.18 [95% HDI: 0.03, 0.33].

To sum up, when considering only participants who were unaware of the effect and using the default effect size prior, we obtain only moderate evidence for an effect of power

posing on felt power. This is in contrast to the results of the previous analysis in which participants who were familiar with the effect were mostly not excluded.

Moderator Analysis: Knowledge of the Effect (Informed Prior on Effect Size)

Next we consider the results based on the informed t prior distribution for effect size with location 0.350, scale 0.102, and three degrees of freedom (depicted in Figure 2) when taking into account only participants unfamiliar with the effect. The results are displayed in Figure 7. As before, the effect size posterior distributions for the individual studies clearly show the influence of the informed prior distribution: the posteriors are narrower and slightly shifted towards the location of the informed prior. Again, the individual study one-sided informed Bayes factors are larger than the default ones. Nevertheless, only one Bayes factor provides moderate evidence for an effect, four provide anecdotal evidence for the alternative or the null hypothesis, and one provides moderate evidence for the null.

The informed meta-analytic fixed-effect Bayes factor equals $BF_{+0} = 6.8$, indicating moderate evidence in favor of an effect of power posing on felt power. The informed meta-analytic random-effect Bayes factor is $BF_{+0} = 2.6$, indicating anecdotal evidence for an effect. As for the default prior, the observed data support a fixed-effect model more than a random-effect model, the Bayes factor that compares case (1), fixed-effect H_+, to case (3), random-effect H_+, (not displayed) indicates that the data are 3.0 times more likely under the fixed-effect model than under the random-effect model. The informed meta-analytic model-averaged Bayes factor is equal to $BF_{+0} = 4.9$ indicating moderate evidence in favor of an effect of power posing on felt power. The median of the model-averaged meta-analytic effect size is equal to 0.23 [95% HDI: 0.10, 0.36].

To sum up, when considering only participants who were unaware of the effect, the results were robust with respect to using the informed or the default prior for the effect size. In both analyses, we found only moderate evidence in favor of the hypothesis that power posing leads to an increase in felt power.
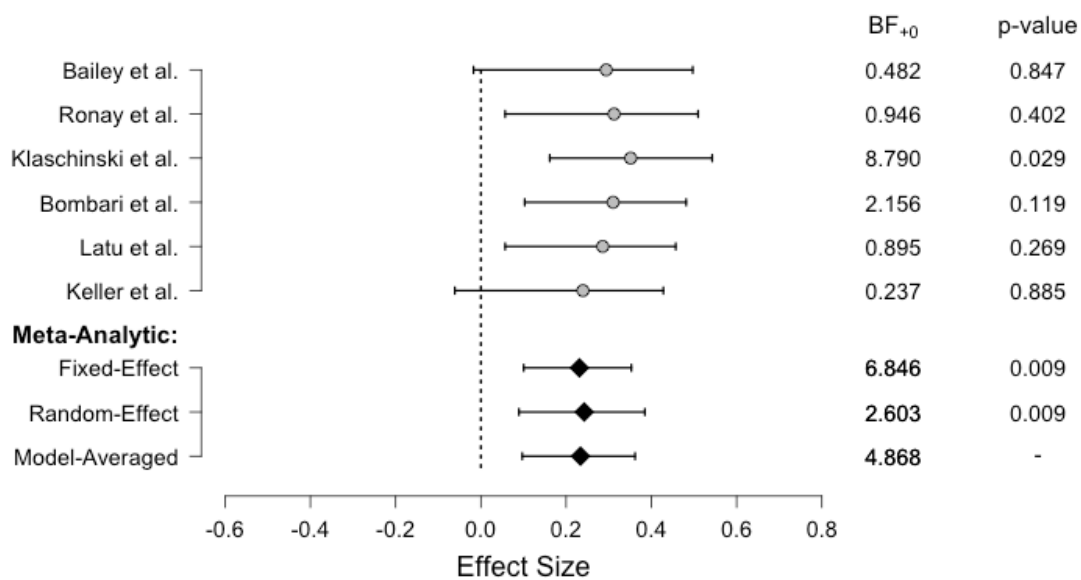
Figure 7: Bayesian model-averaged meta-analysis for the subset of participants unfamiliar with the effect using the informed *t* prior with location 0.350, scale 0.102, and three degrees of freedom for the standardized effect size. The dots and diamonds correspond to the median of the posterior distribution for the effect size; the lines correspond to the 95% highest density intervals. The one-sided Bayes factors are displayed on the right, flanked by classical two-sided *p*-values. Figure available at http://tinyurl.com/n7r4huj under CC license https://creativecommons.org/licenses/by/2.0/.

Discussion

Six preregistered studies in this special issue were subjected to a Bayesian meta-analysis of the effect of power posing on self-reported felt power. The Bayesian approach enabled us to fully acknowledge uncertainty with respect to the choice of a fixed-effect or a random-effect model, and allowed us to incorporate prior information about between-study heterogeneity and plausible effect sizes in the field of psychology. The informed prior distribution for between-study heterogeneity was based on an extensive literature review, and we believe it may serve as an informed default in the field of psychology more generally (cf. Rhodes et al., 2015, for a similar approach in medicine).

When considering the studies as reported (i.e., most studies did not exclude participants who were familiar with the effect), we obtained very strong evidence that adopting high-power poses increases subjective feelings of power; this was the case for both the analysis based on a default prior and an informed prior for the effect size. However, when considering only participants unfamiliar with the effect, we obtained only moderate evidence for an effect for both the default and informed effect size prior analysis. This suggests that knowledge of the effect might play a role with respect to the size of the effect of power posing on felt power, although a formal assessment of this possibility requires a different statistical analysis (e.g., Gelman & Stern, 2006; Nieuwenhuis, Forstmann, & Wagenmakers, 2011), the development of which is

beyond the scope of this paper. Future studies might investigate this potential moderating effect and explore the extent to which the felt power effect is a demand characteristic. Note that the Bayesian approach allows us to seamlessly update the evidence as more studies become available (e.g., Scheibehenne et al., 2017).

　　　Our meta-analysis focused on the effect of power posing on feelings of subjective power and did not consider behavioral or hormonal measures. Nevertheless, we would like to emphasize that given a set of preregistered studies that include the behavioral and hormonal measures of interest, our methodology can readily be applied to quantify evidence in a coherent Bayesian way for those measures as well.

References

Bailey, A. H., LaFrance, M., & Dovidio, J. F. (this issue). Could a woman be superman? Gender and the embodiment of power postures. *Comprehensive Results in Social Psychology*.

Bombari, D., Schmid Mast, M., & Pulfrey, C. (this issue). Real and imagined power poses: Is the physical experience necessary after all? *Comprehensive Results in Social Psychology*.

Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2010). Power posing: Brief nonverbal displays affect neuroendocrine levels and risk tolerance. *Psychological Science*, 21, 1363–1368.

Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2015). Review and summary of research on the embodied effects of expansive (vs. contractive) nonverbal displays. *Psychological Science*, 26, 657–663.

Garrison, K. E., Tang, D., & Schmeichel, B. J. (2016). Embodying power: A preregistered replication and extension of the power pose effect. *Social Psychological and Personality Science*, 7, 623– 630.

Gelman, A. & Stern, H. (2006). The difference between "significant" and "not significant" is not itself statistically significant. *The American Statistician, 60*, 328–331.

Gronau, Q. F., Ly, A., & Wagenmakers, E.-J. (2017). Informed Bayesian *t*-tests. Manuscript submitted for publication. https://arxiv.org/abs/1704.02479

Haldane, J. (1932). A Note on Inverse Probability. *Mathematical Proceedings of the Cambridge Philosophical Society*, 28, 55–61.

Heck, D. W. & Gronau, Q. F. (2017). metaBMA: Bayesian Model Averaging for Random and Fixed Effects Meta-Analysis. https://github.com/danheck/metaBMA

Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science*, 14, 382–417.

Jackson, B., Nault, K., Smart Richman, L., LaBelle, O., & Rohleder, N. (this issue). Does that pose becomes you? Testing the effect of body postures on self-concept. *Comprehensive Results in Social Psychology*.

Jeffreys, H. (1961). Theory of probability (3rd ed.). Oxford, UK: Oxford University Press.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795.

Keller, V. N., Johnson, D. J., & Harder, J. A. (this issue). Meeting your inner super(wo)man: Are power poses effective when taught? *Comprehensive Results in Social Psychology*.

Klaschinski, L., Schröder-Abé, M., & Schnabel, K. (this issue). Benefits of power posing: Effects on dominance and social sensitivity. *Comprehensive Results in Social Psychology*.

Latu, I. M., Duffy, S., Pardal, V., & Alger, M. (this issue). Power vs. persuasion: Can open body postures embody openness to persuasion? *Comprehensive Results in Social Psychology*.

Ly, A., Verhagen, A. J., & Wagenmakers, E.-J. (2016). Harold Jeffreys's default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72, 19–32.

Marsman, M., Schönbrodt, F. D., Morey, R. D., Yao, Y., Gelman, A., & Wagenmakers, E.-J. (2017). A Bayesian bird's eye view of "replications of important results in social psychology". *Royal Society Open Science*, 4, 160426.

Morey, R. D., & Rouder, J. N. (2015). BayesFactor: Computation of Bayes Factors for Common Designs. R package version 0.9.12-2. https://CRAN.R-project.org/package=BayesFactor.

Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E.-J. (2011). Erroneous analyses of interactions in neuroscience: A problem of significance. *Nature Neuroscience, 14*, 1105-1107.

Ranehill, E., Dreber, A., Johannesson, M., Leiberg, S., Sul, S., & Weber, R. A. (2015). Assessing the robustness of power posing: No effect on hormones and risk tolerance in a large sample of men and women. *Psychological Science*, 26, 653–656.

Rhodes, K. M., Turner, R. M., & Higgins, J. P. T. (2015). Predictive distributions were developed for the extent of heterogeneity in meta-analyses of continuous outcome data. *Journal of Clinical Epidemiology, 68*, 52–60. Doi: 10.1016/j.jclinepi.2014.08.012

Ronay, R., Tybur, J. M., van Huijstee, D., & Morssinkhof, M. (this issue). Embodied power, testosterone, and overconfidence as a causal pathway to risk taking. *Comprehensive Results in Social Psychology*.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225–237.

Scheibehenne, B., Gronau, Q. F., Jamil, T., & Wagenmakers, E.-J. (2017). Fixed or random? A resolution through model-averaging. Manuscript submitted for publication.

Sutton, A. J., & Abrams, K. R. (2001). Bayesian methods in meta-analysis and evidence synthesis. *Statistical methods in medical research*, 10, 277–303.

van Erp, S., Verhagen, J., Grasman, R., & Wagenmakers, E.-J. (2017). Estimates of between-study heterogeneity for 705 meta-analyses reported in Psychological Bulletin from 1990-2013. Manuscript submitted for publication. https://osf.io/preprints/psyarxiv/myu9c

Wagenmakers, E.-J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams, R., . . . Zwaan, R. A. (2016). Registered Replication Report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, 11, 917–928.

Appendix

Here we investigate whether and how the analyses results change under different priors for the between-study heterogeneity. Specifically, we explore two alternative prior choices to the Inverse-Gamma(1, 0.15) prior: (1) the maximum-likelihood inverse-gamma distribution (depicted as a dashed line in Figure 8); and (2) a Beta(1,2) prior distribution (depicted as a dotted line in Figure 8).
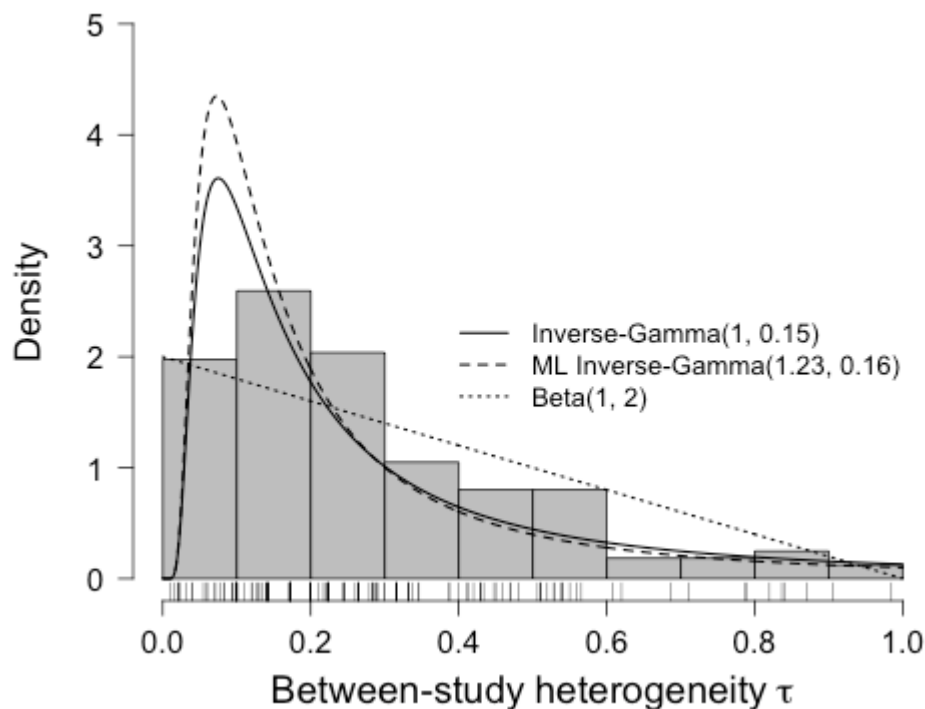
Figure 8: Distribution of the non-zero between-study standard deviations from meta-analyses reported in *Psychological Bulletin* (1990-2013; van Erp et al., 2017). The informed Inverse-Gamma(1, 0.15) prior distribution is displayed on top as a solid line, the maximum-likelihood inverse-gamma distribution is depicted as a dashed line, and the Beta(1, 2) distribution is depicted as a dotted line. Figure available at http://tinyurl.com/k6yyz6b under CC license https://creativecommons.org/licenses/by/2.0/.

Table 2 displays the results for the reported data and Table 3 displays the results for the data of the subset of participants who were unfamiliar with the power pose effect: for all three prior choices for the between-study heterogeneity the results are highly similar.

Table 2. Meta-analytic Bayes factors (BF_{+0}) for different prior choices for the between-study heterogeneity (reported data).

	Inverse-Gamma(1, 0.15)	ML Inverse-Gamma	Beta(1, 2)
meta-analytic fixed-effect Bayes factor	89.6	89.6	89.6
informed meta-analytic fixed-effect Bayes factor	191.8	191.8	191.8
meta-analytic random-effect Bayes factor	9.4	10.0	9.2
informed meta-analytic random-effect Bayes factor	20.7	22.0	20.2
meta-analytic model-averaged Bayes factor	33.1	32.1	35.1
informed meta-analytic model-averaged Bayes factor	71.4	69.1	75.5

Table 3. Meta-analytic Bayes factors (BF_{+0}) for different prior choices for the between-study heterogeneity (unfamiliar participants).

	Inverse-Gamma(1, 0.15)	ML Inverse-Gamma	Beta(1, 2)
meta-analytic fixed-effect Bayes factor	4.4	4.4	4.4
informed meta-analytic fixed-effect Bayes factor	6.8	6.8	6.8
meta-analytic random-effect Bayes factor	1.6	1.7	1.7
informed meta-analytic random-effect Bayes factor	2.6	2.7	2.7
meta-analytic model-averaged Bayes factor	3.1	3.1	3.3
informed meta-analytic model-averaged Bayes factor	4.9	4.8	5.1