From theories to models to predictions: A Bayesian model comparison approach

Jeffrey N. Rouder^{1,2}, Julia M. Haaf², & Frederik Aust³

¹ University of California, Irvine

² University of Missouri

³ University of Cologne

Version 4, Final Before Proofs, 10/2017

Author Note

Jeffrey N. Rouder, Department of Cognitive Sciences, University of California, Irvine and the Department of Psychological Sciences, University of Missouri, Columbia. Julia M. Haaf, Department of Psychological Sciences, University of Missouri, Columbia. Frederik Aust, Department of Psychology, University of Cologne, Germany.

Correspondence concerning this article should be addressed to Jeffrey N. Rouder, 210 McAlester Hall, Columbia, MO, USA, 65203. E-mail: jrouder@uci.edu

Abstract

A key goal in research is to use data to assess competing hypotheses or theories. An alternative to the conventional significance testing is Bayesian model comparison. The main idea is that competing theories are represented by statistical models. In the Bayesian framework, these models then yield predictions about data even before the data are seen. How well the data match the predictions under competing models may be calculated, and the ratio of these matches—the Bayes factor—is used to assess the evidence for one model compared to another. We illustrate the process of going from theories to models and to predictions in the context of two hypothetical examples about how exposure to media affects attitudes toward refugees.

Keywords: Bayesian Analysis, Bayesian Model Comparison, Bayes factors, Bayesian ANOVA

From theories to models to predictions: A Bayesian model comparison approach

In the last ten years we have witnessed profound changes in the methods by which researchers produce and evaluate experimental evidence. The awareness that many analyses are underpowered (Rossi, 1990; Sedlmeier & Gigerenzer, 1989; Szucs & Ioannidis, 2017), based on questionable statistical practices (John, Loewenstein, & Prelec, 2012; Simmons, Nelson, & Simonsohn, 2011), seem implausible in aggregate (Francis, 2012; Ioannidis, 2005) or do not replicate (Open Science Collaboration, 2015) has generated a host of beneficial recommendations. Many of these recommendations are to provide greater transparency in research. One example is preregistration of experiments and analyses so that readers can see what is planned and what is not (Greve, Bröder, & Erdfelder, 2013; Wagenmakers, Wetzels, Borsboom, van der Maas, & Kievit, 2012). A second example is a new focus on open data and analyses where the raw data and the analysis code are readily available (Morey et al., 2016; Rouder, 2016).

One of the unresolved questions from the reform movement is whether researchers should continue to use p -values and the associated $p < .05$ rule for statistical inference. One concern is that the $p < .05$ rule does not take into account prior information, the role of alternatives, or context (Gallistel, 2009; Gigerenzer, 1998; Rindskopf, 1997; Rouder, Morey, Verhagen, Swagman, & Wagenmakers, In press; Sellke, Bayarri, & Berger, 2001; Wagenmakers, 2007). Some researchers find these problems so vexing that they recommend avoiding p -values entirely (Edwards, Lindman, & Savage, 1963; Gallistel, 2009; Rouder, Speckman, Sun, Morey, & Iverson, 2009; Wagenmakers, 2007).

The argument, perhaps stated most clearly by Sellke et al. (2001), goes as follows: Suppose a researcher is assessing whether there is a true effect in a data set. The observed effect size is $d = 0.10$ for 400 observations. The p -value is significant: $p \approx .046$. Now, suppose that given the context of the experiment, we expect true effects, if they exist, to be larger than .25. With this additional context, we might ask how likely is the observed effect size, $d = 0.10$, under both the null hypothesis (true effect size of 0) vs. a reasonable alternative

(true effect size of .25). Because the observed value, .10, is less than half way between the hypotheses, it is more plausible under the null than the alternative hypothesis. Even though the null is more plausible given this context, it is nonetheless ruled out by p -values.

One of the more striking methodological developments over the past 20 years in the social sciences is the popularity of Bayesian methods. This popularity is all the more impressive given that the vast majority of social scientists were not taught Bayesian methods as part of their formal training. There are many reasons for this increased popularity, and perhaps the most relevant has to do with tractability. With the advent of large models that span thousands of parameters and large data sets that span millions of observations, Bayesian models have become popular because many models that are intractable in classical frameworks are possible in the Bayesian framework.

This increased tractability, while important, is not the only advantage of Bayesian methods that strike us. One of the more under-appreciated aspects of Bayesian analysis is that it forces the analyst to make more assumptions about the data than classical methods. Most researchers would not find this advantageous; the usual wisdom is that the fewer assumptions brought to data analysis, the better. In this regard, we may say Bayesian analysis *costs* more than classical analysis, where the cost is measured in assumptions. But, as we shall show, although the Bayesian methods cost more, the payoff is well worth it. Some of the practical advantages are: (1) analysts can make principled statements of evidence for null hypotheses of no effects; (2) analysts can characterize the evidence for hypotheses in small samples without being criticized for being underpowered; and (3) analysts can interpret results that are almost-significant or just-significant by p -values in a principled way.

We noted above that classical analysis by p -values require good judgment which is outside the statistical theory. The Bayesian approach also requires good judgment, but this judgment enters as part of the needed assumptions. Researchers must make these assumptions, *a priori*, that is, before the data are collected. Because these assumptions are made *a priori* they do not depend on the outcome of the study, and this lack of dependence

adds to the transparency of social science. In fact, once the critical *a priori* assumptions are made, the remaining mechanisms of Bayesian inference are strictly lawlike and are applied in an automatic fashion without any further inputs or judgments. Here is how it works:

The main idea is that important questions are answered by positing competing theoretical positions. These positions may be instantiated as statistical models. Model specification is the key step, and it represents the art of the science. It is in model specification where researchers add the needed assumptions, and it is this specification where good judgment is imparted. Once models are specified, the automatic component commences. Bayesian models predict data. Once the data are observed, the researcher may assess how well each model predicted them. Then, by comparing how well each model did, evidence may be stated for each model and, by proxy, on the corresponding theoretical positions. As it will be shown, this focus on predictions is different than in conventional analysis.

In what follows we ask readers to read for the big picture. We describe an overview of Bayesian analysis and contrast it with conventional analysis. We skim on some of the details, and we do not provide a step-by-step guide. After reading this article, it is likely that a reader new to Bayesian analysis will not be able to do them without help. We view this paper as a down payment or a promissory note. We endeavor to provide an account of the logic, appeal, and beauty of Bayesian analysis. Hopefully, with this big-picture view, the reader will be motivated to pursue this line of reasoning about data further.

We illustrate the Bayesian approach within the context of the following research problem: With the rise of the Syrian Civil War, public opinion about the plight of refugees is of practical importance. We ask about the effect of media stories on attitudes toward Syrian refugees. Suppose researchers show media stories about refugees to participants and then assess the participants' favorability ratings toward refugees with a questionnaire. To simplify the situation, we assume the questionnaire leads to a reliable, single measure of favorability. The media stories themselves are either sympathetic stories about the plight and ordeals faced by Syrian refugees displaced in Jordan or control stories about food trends.

In addition to story type, researchers record the political affiliation of the participants. For simplicity, we are going to assume that people vary through two levels for political affiliation: liberal or conservative. We do not recommend having such coarse measures in practice, but they are useful to simplify the example.

We have generated hypothetical data. Cell means are shown in Figure 1, and the standard inferential approach is to use ANOVA to assess the main effects of story type and political affiliation as well as the interaction between them. The results are shown in Table 1. Moreover, sometimes researchers with this design will subdivide the data further, for example, asking if there are separate media effects for liberals and for conservatives. When faced with a 2-by-2 factorial design, researchers often opt for this approach regardless of the context or theoretical questions.

To us, one of the problems of significance testing is that it often does not correspond to the most useful theoretical questions. We suspect this statement may seem alien if only because researchers are familiar with testing for main effects and interactions, and alternatives are not typically presented. Here, we present an alternative where the testing regime is matched to competing theoretical positions of concern to the researcher. We comment more after reviewing the theoretically relevant questions for the example.

Theoretically Meaningful Questions

One of the main benefits of Bayesian analysis is that theoretical questions intimately guide analysis. We consider the case of two different research teams with different research questions. Team A is running the experiment to see if media exposure raises refugee favorability. They are most interested in the main effect of the story type, and have no particular theoretical stance about the interaction of story type and political affiliation.

Team B considers three main theoretical positions:

1. The first position is that the plight of refugees generates an empathic response which results in an increase in favorability for conservative and liberals alike.

2. The second position is informed by classic work by Tversky & Kahneman (1974) on the confirmation bias. In a strong version considered here, people may ignore media stories that are incompatible with their world view. In the current climate, many conservatives endorse a law-and-order view where Syrian refugees present a direct threat to their welfare and security. Hence, they may ignore media stories about the plight of refugees. Accordingly, only liberals should be affected by these stories.
3. The third position is motivated by the recent work of Cook & Lewandowsky (2016) who studied the effects of presenting scientific information on beliefs about climate change. They document a strange phenomenon called *belief-polarization*. When presented with this scientific information, liberals and conservatives had markedly different responses. The information led liberals to believe more strongly in climate change and led conservatives to disbelieve more strongly in it. The increased disbelief of conservatives when presented with valid data supporting climate change is paradoxical. Their attitudes are moving in the opposite direction of the information. Cook and Lewandowsky show that this phenomenon may be explained by considering how people treat the validity of information they do not agree with. In this case, the conservative respondents see the scientific information as fraudulent, and the fact that researchers must resort to supposed fraudulent information to support global warming is viewed as further evidence against its existence. Team B wishes to assess whether the same belief-polarization dynamics are present in the refugee scenario. They suspect that if so, the positive stories might polarize attitudes. Liberals might see refugees more favorably after viewing positive stories while conservatives might actually see them more negatively.

We see here that Team B is asking targeted questions about three different theories. These questions cannot be answered by assessing main effects and interactions. Instead, they require statistical models that match the theories. We show how these models are specified, and then we show how Bayesian analyses works. We start first with Team A, below, who

asks a relatively straightforward question and then progress to Team B, who asks these more theoretically-inspired questions.

Team A's Analysis

Capturing Positions With Models:

Team A considers two competing positions. The first is that there is no effect of story type whatsoever. The second is that there is a positive effect. We wish to construct two models that capture these competing positions. The best way of doing so is to use formal notation. Let Y_{ij} be the favorability score. Subscript i indexes the story condition with $i = 1$ indicating a control story and $i = 2$ indicating a refugee story. Subscript j indexes the participant number, i.e., the j th participant in the i th story condition.

We start with a null or no-effects model for Team A which we call $\mathcal{M}_{A_0}$. Under this model, ratings in each condition have the same distribution:

$$\begin{aligned}\mathcal{M}_A : \quad Y_{1j} &= \mu + \text{error}, \\ Y_{2j} &= \mu + \text{error},\end{aligned}$$

where μ is the overall mean of favorability rating and the error term is random deviation with a mean of zero.

The next step is specifying the effects model, where there is an effect. The key goal is keeping μ the overall mean across conditions. Hence, whatever is added to one condition has to be subtracted from the other: We take:

$$\begin{aligned}Y_{1j} &= \mu - \frac{\alpha}{2} + \text{error}, \\ Y_{2j} &= \mu + \frac{\alpha}{2} + \text{error}.\end{aligned}$$

Favorability ratings for those in the control story condition center around $\mu - \frac{\alpha}{2}$ while those in the refugee story center around $\mu + \frac{\alpha}{2}$. In this formulation, α , the difference between conditions, serves as the effect.

If we were conducting standard testing, we would be finished with model specification. The models above underlie the t -test, and this test would be that of the null hypothesis that $\alpha = 0$. The p -value would be the discord between the data and this null. In Bayesian testing, however, we need to specify the models to the point they provide predictions. For the no-effects model, we are done. This model states that the true effect, α is zero, and this specification is indicated by the vertical arrow in Figure 2a. The model already makes predictions about the observed t -value—it should follow a t distribution. We can make this prediction even before seeing any of the data.

The specification of the effects model is more complicated. In usual testing, there are no assumptions made about the effect—it can be any value without constraint. Yet, in the Bayesian framework, this specification is too vague because it does not yield any predictions on where the data will be (other than that the data can be anywhere). In Bayesian analysis, we must make a finer-grained specification on the effect, α .

Our choice is that under the effects model, the effect α is distributed as a half-normal:

$$\mathcal{M}_{A_1} : \quad \alpha \sim \text{Positive-Half-Normal}$$

The positive-half-normal distribution¹ is shown in Figure 2a. It is exactly what it sounds like—the positive half of a normal distribution. The distribution indicates the plausibility of effects before the data have been collected. These plausibility values are higher for effects closer to zero, indicating that, if there is an effect, it is likely to be relatively small. Because we make this fine-grained specification, predictions are now possible.

¹The density of the positive-half normal with mean μ and variance b^2 is given by $f(y) = \frac{1}{\sqrt{2\pi}b^2} \frac{\exp(-\frac{(y-\mu)^2}{2b^2})}{\Phi(\frac{\mu}{b})}$, where Φ is the CDF of a standard normal. In practice, we set μ to zero and b to 0.7 make this specification appropriate for Likert scale ratings.

Comparing Models: Evidence Through Predictions

The next step is to compare the above models. Rather than providing a formal justification of the Bayesian approach, which may be found in Jeffreys (1961), Kass & Raftery (1995), and Morey, Romeijn, & Rouder (2016), we provide an informal discussion. In Bayesian statistics, evidence for models reflects how well they predict data.

Predictions are possible because the above models are finely specified. The two models under consideration are depicted in Figure 2a, and these specifications are on true or population parameters. Predictions, however, are for data, and therefore have to take sample noise into consideration. The relevant statistic that summarizes the data in this case is the observed effect, $\hat{\alpha} = \bar{Y}_2 - \bar{Y}_1$. Take the case of the null model, $\mathcal{M}_{A_0}$. The observed effect is predicted to be spread around 0 with the amount of spread decreasing with the sample size. Figure 2b shows the prediction for the sample sizes in the hypothetical data set of Figure 1. Also shown is the prediction for the effects model $\mathcal{M}_{A_1}$. Here too the effect of sample noise is to smear the form of the model.² We see that the null better predicts effects smaller than about a quarter of a point on the Likert scale; the effects model better predicts effects larger than a quarter. In practice, analysts do not need to compute these predictions; there is software for this. We show them here so a reader can get a deeper view of the mechanics of Bayesian analysis.

Once the predictions are known, model comparison is simple. All we need to do is note where the data fall. The vertical red line in Figure 2b depicts the observed effect for Team A. Where it intersects the prediction curves is a measure of the probability of the data under the models. The ratio of the height of the prediction curves at the data (red line) is the critical quantity. In the figure, it has a value of 3.6, indicating that the data are 3.6 times more probable under Model $\mathcal{M}_{A_1}$ than under Model $\mathcal{M}_{A_0}$. This ratio is known as the *Bayes factor*, and it serves as a measure of evidence for one model compared to another. The Bayes

²More technically, the predictions are the integral $\int_{\theta} f(Y|\theta)\pi(\theta)d\theta$ where $f(Y|\theta)$ is the probability density of observations conditional on parameter values and $\pi(\theta)$ is the probability density of the parameters.

factor for Team A is shown in Table 2.

The critical difference between Bayesian and conventional statistics is now clear. In Bayesian statistics one expends good judgment specifying models to a fine degree so that they make predictions. Once these specifications are made, the predictions follow, and from these model comparison follows straightforwardly. In conventional statistics, alternative models are specified coarsely, so coarsely that they do not predict data; instead, they account for data by adjusting parameter values after the data are known. As a result, inference requires good judgment to interpret, but what that judgement is remains unclear. In the Bayesian case, in contrast, all the needed judgment is in model specification; not only are these specifications made clearly in notation, they may be defended, critiqued, and discussed transparently.

You Can Do It Too

The case for Team A is fairly standard. There are two conditions, and we wish to know if there is a positive effect or a null effect. In the above example, we worked through the model specification and model predictions in great detail to explicate Bayesian testing.

In practice, one does not need to work out the predictions of the model. We have done so already for this case (Morey & Rouder, 2011; see Rouder et al., 2009), and we have set up software for computation. There are two main choices for users. The first is JASP (Love et al., 2015). JASP is a freely available, open-source SPSS clone that performs both conventional and Bayesian analyses. It is a wonderful resource for anyone who uses SPSS. Tutorials on how to use JASP may be found on the [JASP website](#). Another approach, one that we use³, is the BayesFactor package in R (Morey & Rouder, 2015). This package is the backend to JASP, and it adds additional flexibility. Documentation for BayesFactor comes

³Throughout this paper, we used R (3.4.0, R Core Team, 2016) and the R-packages *afex* (0.18.0, Singmann, Bolker, Westfall, & Aust, 2016), *BayesFactor* (0.9.12.2, Morey & Rouder, 2015), *cowplot* (0.8.0, Wilke, 2016), *dplyr* (0.5.0, Wickham & Francois, 2016), *ggplot2* (2.2.1, Wickham, 2009), *Hmisc* (4.0.3, Harrell Jr, Charles Dupont, & others., 2016), *msm* (1.6.4, Jackson, 2011), *mvtnorm* (1.0.6, Genz & Bretz, 2009; Wilhelm & G, 2015), *papaja* (0.1.0.9492, Aust & Barth, 2016), *shape* (1.4.2, Soetaert, 2014), *spatialfil* (0.15, Dinapoli & Gatta, 2015), and *mvtnorm* (1.4.10, Wilhelm & G, 2015) for all our analyses and figures.

with the package.

We follow on with the more advanced case for Team B. Though the data are the same for both teams, Team B considers an additional variable, political affiliation, and several theoretical models. We present it to show the power of the theory-to-models-to-predictions viewpoint, but also note that analysis is not yet fully available in JASP or in BayesFactor/R. Computing the predictions of these models requires some custom coding, and the details are provided in Haaf & Rouder (2017). Given our involvement in JASP and BayesFactor, we suspect these types of analyses will become available without custom coding shortly.

Team B's Analysis

We expand the setup⁴ to consider two effects: the effect of the story condition for the conservatives, denoted α , and the effect of the story condition for the liberals, denoted β . The key endeavor is expressing the theoretical positions as constraints on α and β .

Team B starts with a null model, that there is no effect of story for either liberals or conservatives. This model is instantiated as

$$\mathcal{M}_{B_0} : \quad \alpha = \beta = 0.$$

The left column of Figure 3 shows several models, and Figure 3a shows Model $\mathcal{M}_{B_0}$. Here, there is a spike at the origin, which is in this case indicated by the dot. No other values are possible.

⁴The notation is expanded as follows. Let u_j be an indicator of whether the j th participant is conservative or liberal, with $u_j = 1$ if a participant is conservative and $u_j = 0$ if a participant is liberal, respectively. Likewise, let w_i be an indicator of the condition with $w_i = -1/2$ and $w_i = 1/2$ for control and refugee-plight story condition, respectively. We model the data with four parameters: an overall mean, μ , an overall main effect of political orientation, η , an effect of story condition for conservatives, α , and an effect of story condition for liberals, β . With these parameters, a model on data may be expressed as:

$$\mathcal{M}_B : \quad Y_{ij} = \mu + (u_j - \frac{1}{2})\eta + u_j w_i \alpha + (1 - u_j) w_i \beta + \text{error}.$$

The first term, μ , is the grand mean; the second term, $(u_j - \frac{1}{2})\eta$, takes on values of $\eta/2$ when the person is conservative, and $-\eta/2$ when the person is liberal. The third term, $u_j w_i \alpha$, only comes into play for conservatives, and it is $-\alpha/2$ for the control story and $\alpha/2$ for the refugee story. The last term is the analogous term for liberals.

The next position is that the media effect is the same for both conservatives and liberals. This position is instantiated as $\alpha = \beta$. Once again, because we take a Bayesian approach, a fine-grained specification is needed:

$$\mathcal{M}_{B_1} : \quad \alpha = \beta \sim \text{Positive-Half-Normal}.$$

Here, we put the positive half of a normal distribution on a common effect. This specification means that if there is an effect, it is (a) equal for conservatives and liberals, (b) positive, and (c) relatively small in size. A graphical representation is provided in Figure 3b. Here, there is mass on the diagonal corresponding to the equality between the effect for conservatives and liberals.

The theoretical position that people ignore information that is incompatible with their world view is implemented by setting $\alpha = 0$, that is, no effect for conservatives, and $\beta > 0$, that is, there is an effect for liberals. The fine-grained specification is

$$\begin{aligned} \mathcal{M}_{B_2} : \quad \alpha &= 0 \\ \beta &\sim \text{Positive-Half-Normal}. \end{aligned}$$

Figure 3c provides the graphical representation of this model. The mass of this specification is on a vertical line where the effect for conservatives is zero. Most of the mass is located where liberals have a small effect because of the half-normal shape of the prior distribution on β .

The Cook and Lewandowski theory that refugee stories polarize pre-existing favorability ratings is instantiated as $\alpha < 0 < \beta$. With a fine-grained specification, the model is

$$\begin{aligned} \mathcal{M}_{B_3} : \quad \alpha &\sim \text{Negative-Half-Normal} \\ \beta &\sim \text{Positive-Half-Normal}. \end{aligned}$$

The graphical representation of this model is shown in Figure 3d. In this case, the mass is in

the upper left quadrant of the figure, and there is most mass close to zero, indicating that smaller effects are more likely for both conservatives and liberals.

Finally, much as it is wise to consider a null model, it is also wise to consider an unrestricted model:

$$\begin{aligned}\mathcal{M}_{B_4} : \quad & \alpha \sim \text{Normal}, \\ & \beta \sim \text{Normal}.\end{aligned}$$

The graphical representation of this model is shown in Figure 3e. Here, we allow the effects to vary as they may. If this model is preferred over any of the previous ones, say because the effect of media on liberals is negative, then we should stop. Rather than discriminating among the previous theories, we should be exploring new ones (and, perhaps, the possibility of a mistake in data collection).

Team B has predictions from several models. Predictions are made for sample effects: $\hat{\alpha} = \bar{Y}_{2c} - \bar{Y}_{1c}$ and $\hat{\beta} = \bar{Y}_{2\ell} - \bar{Y}_{1\ell}$, where the subscripts c and ℓ index conservative and liberal participants, respectively. The right-hand column of 3 shows the predictions for the observed effects, $\hat{\alpha}$ and $\hat{\beta}$. Like the predictions from Team A's models, the predictions from Team B's models are smeared versions of the models.

With these predictions graphed, we are ready for model comparison. The red point in the right-hand column is the observed value for conservative and liberal participants, $\hat{\alpha}$ and $\hat{\beta}$, respectively. As is shown, the density is greatest for the Model $\mathcal{M}_{B_2}$, the model where there is an effect for liberals but not for conservatives.

Bayes factors are ratios of these densities. We choose the null model as a standard and measure all ratios relative to it. These ratios are shown in Table 2. The best model, Model $\mathcal{M}_{B_2}$, is 9000 times more probable than the null model. The second-best model, Model $\mathcal{M}_{B_3}$, is 2000 times more probable than the null model. With Bayes factors, it is possible to compare any two models. For example, the Bayes factor between $\mathcal{M}_{B_2}$ and $\mathcal{M}_{B_3}$ is 4.2-to-1, indicating that the data are more than four times more probable under the former than the

latter.

It may seem counter-intuitive at first that the unrestricted model, $\mathcal{M}_{B_4}$, did so poorly. Data were more probable under the winning model, $\mathcal{M}_{B_2}$, than under this unrestricted model by a factor of 7.6-to-1. What is strange is that $\mathcal{M}_{B_2}$ is a proper restriction on $\mathcal{M}_{B_4}$, and in classical statistics, more restricted models explain fewer data patterns than more general models. This insight is unhelpful in Bayesian statistics. The graph for Team A is helpful here. If we look at the prediction graph in Figure 2b, we see the overall area under each distribution needs to be 1.0, that is, the data have to occur somewhere. Hence, if a model explains more possible data patterns, as does Model $\mathcal{M}_{A_1}$ compared to Model $\mathcal{M}_{A_0}$, it does so at a lower probability. For example, if the effect is small, then although both models account for this effect, Model $\mathcal{M}_{A_0}$ predicts it at a higher probability. These same dynamics are present for Team B. The unrestricted model can account for all patterns of α and β , but as a consequence, it has the lowest density prediction for any observed data. If the data are compatible with the other models, the unrestricted model will fare poorly. However, if the data are incompatible with the other models, say the observed effect was large and negative in value for liberals, the unrestricted model will fare well.

It is in this way that Bayes factors instantiate a principle of parsimony. Simplicity here refers to constraints on predictions. If one model makes highly constrained predictions—and therefore is easier to falsify—while another does not, and if the data are in line with these constrained predictions, the first will be favored in the Bayes factor analysis. It is this strength—that one can state positive evidence for constraint and parsimony—that makes the Bayes factor approach so attractive.

Reporting and Interpreting Bayes Factors

Reporting Bayes factors is straightforward. When there are several to-be-compared models, we opt for tables much like the second part of Table 2. When there are one or two to-be-compared models, we tend to just state the evidence as straightforwardly as we can.

For example, if we were Team A in the previous example, we may write: “We tested whether there was a positive effect of sympathetic stories against a null model with no effect using Bayes factors. The Bayes factor favored the positive-effect model by 3.6-to-1, indicating that the data are almost four times as probable under the effect model than under the no-effect model.”

How shall we interpret values such as 3.6-to-1? As discussed, Bayes factors describe the relative probability of data under competing models, or the predictive accuracy of the models for the observed data. To us, these values are odds or ratios, much like odds on elections or horse races. Interestingly, in most contexts where people use odds to convey information, the odds themselves are directly interpretable without much need for decisions or verbal labels. For example, if a horse is a 5-to-2 favorite, the value is what it is without need for judgment about whether these odds are big.

We recommend that researchers take a similar approach. They may simply report Bayes factors and not worry too much about whether they are big enough (Morey, 2015; Rouder, Morey, & Wagenmakers, 2016). The rationale here is that the Bayes factor is a continuous metric of evidence, and researchers do not need to make decisions. For example, if we were Team B, we might write: “The data provide some evidence against the target media-polarization effect when compared to a model where incompatible information is ignored ($BF = 4.2\text{-to-1}$).”

Some researchers prefer to have more prescriptive decision rules. A common convention is that 3-to-1 is a minimal value for favoring one model over another, and 10-to-1 is considered strong evidence for favoring one model over another (Jeffreys, 1961; Kass & Raftery, 1995; Wetzels & Wagenmakers, 2012). Kass & Raftery (1995) note that these prescriptive rules are unnecessary, and that researchers who use them treat them as loose rules of thumb. We would prefer that researchers use substantive context rather than such rules-of-thumb when deciding how Bayes factors support substantive conclusions in practice.

The above approach, where Bayes factor may always be reported as the strength of

evidence from data for competing positions, offers the following pragmatic advantages over classical testing: (1) Evidence for the null hypothesis may be stated. For example, the following statement is perfectly valid—"The data support the null hypothesis over an effect model by a Bayes factor value of 5-to-1." In classical testing, in contrast, the analyst can never state positive support for the null. Instead she or he notes that there is a lack of evidence for an effect. (2) The analyst need not be excessively concerned with power in small-sample designs. When there are small samples, large Bayes factor values become unlikely. But if an analyst is content on reporting Bayes factor values, then these values are reported just like in any other study. One example of such an approach comes from Engelhardt, Mazurek, Hilgard, Rouder, & Bartholow (2015), who studied whether violent media led to aggression in individuals with autism. Because it is hard to find many high-functioning autistic individuals, Engelhardt et al.'s sample size was necessarily limited. Nonetheless, they used Bayes factors to describe the evidence in their data for competing predictions. In some cases, though, data do not provide evidence for either of the models compared, and, in this case, researchers may encounter Bayes factors such as 1.2-to-1. In this case, the data may not be enough to discriminate between the two models at hand and researchers may consider collecting more data. (3) Analysts need not be so concerned with "not quite significant" or "barely significant" outcomes. Bayes factor values are what they are without further fanfare.

Conclusion

We promote the Bayes factor approach presented here because it allows researchers to state the strength of evidence from data for the theoretical positions they wish to assess. The Bayes factor approach we discussed here is not new. Laplace (1812, see Laplace, 1986 for republication) was perhaps the first to discuss odds on hypotheses, and substantial development was made throughout the early and middle parts of the 20th century. Nonetheless, it is only recently that social science researchers have the computational tools

for analysis. Perhaps the easiest software to use is JASP (Love et al., 2015). Another approach, one that we use, is the BayesFactor package in R (Morey & Rouder, 2015).

The above packages are effective for a number of common designs and theoretical positions, but not for all of them. Researchers looking to compute Bayes factors for t -tests, ANOVA, regression, contingency-table modeling are in luck. Moreover, progress is ongoing. The development of Bayes factor methods for ever sophisticated models is a trending area in statistics and quantitative social science research. Application developers are incorporating the latest statistical developments into their packages. Models for which there are no software available this year might very well be implemented by next year.

Who should adopt Bayesian methods? Using Bayesian methods, especially the way we envision it, may seem risky. Will reviewers buy it? Will they think you are cheating? Will they agree with reporting all values without recourse to a metric of significance? Our experience has been very positive. Not only have we never had a paper rejected because of Bayes factors, often our usage is seen as constructive and innovative. In fact, there are papers, such as Engelhardt et al. (2015), that may not have been published without Bayes factor analyses.

We end with one final practical note. Bayesian model comparison is well suited for collaboration between substantive researchers and statisticians. Statisticians by-and-large consider conventional techniques, say structural equation modeling and hierarchical linear models, to be solved problems. Bayes-factor analysis is not solved, and generating new solutions is timely and topical. We find that by focusing on techniques that are state-of-the-art, it is easier to engage our academic colleagues. By working together, substantive researchers and statisticians can specify tractable, novel, insightful models that capture theoretical positions of interest. Then, with this specification, it is often possible for statisticians to compute predictions. With these predictions, it is straightforward to state evidence from data for the models.

Readers interested in Bayes factors have many choices for further exploration. One

starting point is Etz, Gronau, Dablander, Edelsbrunner, & Baribault (In press). These authors have collated, reviewed, ranked, and commented on many of the leading articles and books on Bayesian methods in the social sciences.

References

- Aust, F., & Barth, M. (2016). *Papaja: Create APA manuscripts with R Markdown*. Retrieved from <https://github.com/crsh/papaja>
- Cook, J., & Lewandowsky, S. (2016). Rational irrationality: Modeling climate change belief polarization using Bayesian networks. *Topics in Cognitive Science*, 8(1), 160–179. Retrieved from <https://dx.doi.org/10.1111/tops.12186>
- Dinapoli, N., & Gatta, R. (2015). *Spatialfil: Application of 2D convolution kernel filters to matrices or 3D arrays*. Retrieved from <https://CRAN.R-project.org/package=spatialfil>
- Edwards, W., Lindman, H., & Savage, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, 70, 193–242. Retrieved from <http://dx.doi.org/10.1037/h0044139>
- Engelhardt, C. R., Mazurek, M. O., Hilgard, J., Rouder, J. N., & Bartholow, B. D. (2015). Effects of violent-video-game exposure on aggressive behavior, aggressive-thought accessibility, and aggressive affect among adults with and without autism spectrum disorder. *Psychological Science*, 26, 1187–1200. Retrieved from <https://dx.doi.org/10.1177/0956797615583038>
- Etz, A., Gronau, Q. F., Dablander, F., Edelsbrunner, P., & Baribault, B. (In press). How to become a Bayesian in eight easy steps: An annotated reading list. *Psychonomic Bulletin & Review*.
- Francis, G. (2012). Too good to be true: Publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review*, 19, 151–156. Retrieved from <https://dx.doi.org/10.3758/s13423-012-0227-9>
- Gallistel, C. R. (2009). The importance of proving the null. *Psychological Review*, 116,

- 439–453. Retrieved from <http://psycnet.apa.org/doi/10.1037/a0015251>
- Genz, A., & Bretz, F. (2009). *Computation of multivariate normal and t probabilities*. Heidelberg: Springer-Verlag.
- Gigerenzer, G. (1998). We need statistical thinking, not statistical rituals. *Behavioral and Brain Sciences*, 21, 199–200. Retrieved from <https://doi.org/10.1017/S0140525X98281167>
- Greve, W., Bröder, A., & Erdfelder, E. (2013). Result-blind peer reviews and editorial decisions: A missing pillar of scientific culture. *European Psychologist*, 18(286-294). Retrieved from <https://doi.org/10.1027/1016-9040/a000144>
- Haaf, J. M., & Rouder, J. N. (2017). Developing constraint in Bayesian mixed models. *Psychological Methods*.
- Harrell Jr, F. E., Charles Dupont, & others. (2016). *Hmisc: Harrell miscellaneous*. Retrieved from <https://CRAN.R-project.org/package=Hmisc>
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2, 0696–0701. Retrieved from <https://doi.org/10.1371/journal.pmed.0020124>
- Jackson, C. H. (2011). Multi-state models for panel data: The msm package for R. *Journal of Statistical Software*, 38(8), 1–29. Retrieved from <http://www.jstatsoft.org/v38/i08/>
- Jeffreys, H. (1961). *Theory of probability (3rd edition)*. New York: Oxford University Press.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23(5), 524–532. Retrieved from <http://pss.sagepub.com/content/23/5/524.abstract>
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795. Retrieved from

<http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1995.10476572>

Laplace, P. S. (1986). Memoir on the probability of the causes of events. *Statistical Science*, 1(3), 364–378. Retrieved from <http://www.jstor.org/stable/2245476>

Love, J., Selker, R., Verhagen, J., Smira, M., Wild, A., Marsman, M., . . . Wagenmakers, E.-J. (2015). Software to sharpen your stats. *APS Observer*, 28, 27–29.

Morey, R. D. (2015, january). On verbal categories for the interpretation of bayes factors. Blog entry in BayesFactor. Retrieved from <http://bayesfactor.blogspot.com/2015/01/on-verbal-categories-for-interpretation.html>

Morey, R. D., & Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychological Methods*, 16, 406–419. Retrieved from <http://dx.doi.org/10.1037/a0024377>

Morey, R. D., & Rouder, J. N. (2015). *BayesFactor: Computation of bayes factors for common designs*. Retrieved from <https://CRAN.R-project.org/package=BayesFactor>

Morey, R. D., Chambers, C. D., Etchells, P. J., Harris, C. R., Hoekstra, R., Lakens, D., & Zwaan, R. A. . . . and. (2016). The peer reviewers openness initiative: Incentivizing open research practices through peer review. *Royal Society Open Science*, 3(1). Retrieved from <http://rsos.royalsocietypublishing.org/content/3/1/150547>

Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of Bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, –. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0022249615000723>

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6521), 943. Retrieved from dx.doi.org/10.1126/science.aac4716

R Core Team. (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from

<https://www.R-project.org/>

- Rindskopf, D. (1997). Testing 'small', not null, hypotheses: Classical and Bayesian approaches. In L. Harlow, S. Mulaik, & J. Steiger (Eds.), *What if there were no significance tests?* (pp. 221–257). Mahwah, New Jersey: Erlbaum.
- Rossi, J. S. (1990). Statistical power of psychological research: What have we gained in 20 years? *Journal of Consulting and Clinical Psychology*, 58(5), 646–656. Retrieved from <http://search.ebscohost.com.proxy-ub.rug.nl/login.aspx?direct=true&db=psyh&AN=1991-03239-001&site=ehost-live&scope=site>
- Rouder, J. N. (2016). The what, why, and how of born-open data. *Behavioral Research Methods*, 48, 1062–1069. Retrieved from [10.3758/s13428-015-0630-z](https://doi.org/10.3758/s13428-015-0630-z)
- Rouder, J. N., Morey, R. D., & Wagenmakers, E.-J. (2016). The interplay between subjectivity, statistical practice, and psychological science. *Collabra*, 2, 6. Retrieved from [http://doi.org/10.1525/collabra.28](https://doi.org/10.1525/collabra.28)
- Rouder, J. N., Morey, R. D., Verhagen, A. J., Swagman, A. R., & Wagenmakers, E.-J. (In press). Bayesian analysis of factorial designs. *Psychological Methods*.
- 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 and Review*, 16, 225–237. Retrieved from [http://dx.doi.org/10.3758/PBR.16.2.225](https://doi.org/10.3758/PBR.16.2.225)
- Sedlmeier, P., & Gigerenzer, G. (1989). Do studies of statistical power have an effect on the power of studies? *Psychological Bulletin*, 105, 309–316. Retrieved from [dx.doi.org/10.1037/0033-2909.105.2.309](https://doi.org/10.1037/0033-2909.105.2.309)
- Sellke, T., Bayarri, M. J., & Berger, J. O. (2001). Calibration of *p* values for testing precise null hypotheses. *American Statistician*, 55, 62–71. Retrieved from

<http://dx.doi.org/10.1198/000313001300339950>

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology:

Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359–1366. Retrieved from

<https://doi.org/10.1177/0956797611417632>

Singmann, H., Bolker, B., Westfall, J., & Aust, F. (2016). *Afex: Analysis of factorial experiments*. Retrieved from <https://CRAN.R-project.org/package=afex>

Soetaert, K. (2014). *Shape: Functions for plotting graphical shapes, colors*. Retrieved from

<https://CRAN.R-project.org/package=shape>

Szucs, D., & Ioannidis, J. P. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biology*, 15(3), e2000797. Retrieved from <https://doi.org/10.1371/journal.pbio.2000797>

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. Retrieved from

<https://doi.org/10.1126/science.185.4157.1124>

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problem of p values.

Psychonomic Bulletin and Review, 14, 779–804. Retrieved from

<https://doi.org/10.3758/BF03194105>

Wagenmakers, E.-J., Wetzels, R., Borsboom, D., van der Maas, H. L. J., & Kievit, R. A.

(2012). An agenda for purely confirmatory research. *Perspectives on Psychological Science*, 7, 627–633. Retrieved from <https://doi.org/10.1177/1745691612463078>

Wetzels, R., & Wagenmakers, E. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review*, 19, 1057–1064. Retrieved

from <https://dx.doi.org/10.3758%2Fs13423-012-0295-x>

Wickham, H. (2009). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York.

Retrieved from <http://ggplot2.org>

Wickham, H., & Francois, R. (2016). *Dplyr: A grammar of data manipulation*. Retrieved

from <https://CRAN.R-project.org/package=dplyr>

Wilhelm, S., & G, M. B. (2015). *tmvtnorm: Truncated multivariate normal and student t*

distribution. Retrieved from <http://CRAN.R-project.org/package=tmvtnorm>

Wilke, C. O. (2016). *Cowplot: Streamlined plot theme and plot annotations for 'ggplot2'*.

Retrieved from <https://CRAN.R-project.org/package=cowplot>

Table 1

Results of the standard ANOVA analyses and separate follow-up test of the effect of the media story for conservatives and liberals.

Effect	F	df_1	df_2	MSE	p
Omnibus ANOVA					
Story	5.88	1	116	0.73	.017
Political affiliation	17.10	1	116	0.73	< .001
Story $\times$ Political affiliation	4.89	1	116	0.73	.029
Follow-up tests					
Story _{Conservatives}	0.02	1	58	0.95	.895
Story _{Liberals}	15.20	1	58	0.52	< .001

Table 2

Bayes factors for each alternative model against the corresponding null model. The first line contains the Bayes factor for Team A. The other lines contain Bayes factors for Team B's models. For Team B, the model with an effect of story only for liberals is preferred over all other models.

Models	Description	BF ₁₀
Team A		
$\mathcal{M}_{A_0}$	No effect	
$\mathcal{M}_{A_1}$	Story effect	3.61
Team B		
$\mathcal{M}_{B_0}$	No effect	
$\mathcal{M}_{B_1}$	Empathic response	470.06
$\mathcal{M}_{B_2}$	Confirmation bias	9,094.61
$\mathcal{M}_{B_3}$	Belief polarization	2,150.03
$\mathcal{M}_{B_4}$	Unconstrained effects	1,203.01

Figure 1. Simulated ratings of participants' favorability towards immigrants; higher values represent more favorable ratings. Points and error bars denote the sample mean and 95% confidence intervals.

Figure 2. Model and predictions for Team A. a. The dashed arrow shows the null model (no effect), the solid line shows the prior distribution for the alternative model of a positive effect of media story on favorability ratings. Higher values represent higher plausibility of the favorability effect before the data have been collected. b. The dashed curve shows the prediction of the no-effect model, the solid curve shows the prediction of the positive-effect model. Predictions take into account sampling noise and are smeared versions of the models themselves. The solid red line represents the observed favorability effect. The data are better predicted by the positive-effect model than by the no-effect model; the Bayes factor favors the positive-effect model 3.61-to-1.

Figure 3. Models (left) and predictions (right) for Team B. Darker areas represent higher plausibility of the favorability effects for Conservatives and Liberals before the data have been collected. Predictions take into account sampling noise and are smeared versions of the models themselves. The red point represents the observed favorability effect. The data are best predicted by the confirmation-bias model.