Optional Stopping and the Interpretation of Bayes Factor

Jeffrey N. Rouder¹ & Julia M. Haaf²

¹ University of California, Irvine

² University of Amsterdam

Author Note

We thank Richard Morey, EJ Wagenmakers, and Peter Grünwald for insightful discussions about Bayes factors. This paper and simulation code are available at <https://osf.io/uv456/>.

Correspondence concerning this article should be addressed to Jeffrey N. Rouder, .
E-mail: jrouder@uci.edu

Abstract

In their paper *Why optional stopping is a problem for Bayesians*, de Heide and Grünwald (2020) critique the claim that those using Bayes factor can sample until there is sufficient evidence to support one or another model. de Heide and Grünwald’s main message is that unless you believe your priors to govern the data-generating mechanism, optional stopping may distort inference. We show here that the distortions are not in inference but in interpretation. Their claim is about what happens when the analyst’s models do not match the data generating models, that is, the Bayes factor for models not specified. Given that we never know the underlying data generation mechanism we ask researcher to interpret Bayes factors for what they are—the relative strength of evidence from data for two hypotheticals. We discuss how researchers should assess the robustness of their results by considering reasonable variation in model specification conditional on observed data rather than on hypothetical truths.

Keywords: Inference, Bayesian Inference, Bayes Factors, Optional Stopping

Optional Stopping and the Interpretation of Bayes Factor

We are pleased that several recent authors including de Heide and Grunwald (2020) and Tendeiro, Kiers, and van Ravenzwaaij (2020) have continued the discussion on the interpretation of Bayes factors, the role of default priors, and the effects of sampling plans on analysis. Their works are nuanced and welcomed additions that add depth to the literature. We are also pleased that interest in optional stopping remains high. Perhaps researchers intuitively know that they can safely terminate studies if a high level of evidence is reached before the target sample size.

Prior to Rouder (2014), the evaluation of Bayesian optional stopping in the psychological literature was conceptually flawed. For example, Yu, Sprenger, Thomas, and Dougherty (2014) and Sanborn and Hills (2014) simulated data from ground truth and assessed the distributions of Bayes factors under fixed and optional stopping. Rouder (2014) observed that this simulation process treats the Bayes factor as a frequentist test statistic rather than a Bayesian measure of the strength of evidence for competing models conditional on data. In Bayesian simulations, one should evaluate the probability that the data came from models conditional on data. Rouder (2014) showed how to perform these simulations. The main goal therein was to assess the Bayesian properties in light of optional stopping.

It is straightforward to show either by proof (Deng, Lu, & Chen, 2016; Hendriksen, de Heide, & Grünwald, 2020; Tendeiro et al., 2020) or simulation (de Heide & Grunwald, 2020; Rouder, 2014) that Bayesian inference is well calibrated regardless of the stopping plan as noted by Berger and Wolpert (1988), Edwards, Lindman, and Savage (1963), and Lindley (1957). In fact, our colleague Richard Morey was so singularly unimpressed with the simulations that he quipped to Rouder at the time, “So, you simulated the Law of Conditional Probability.” Indeed, as Rouder (2014) confessed that the simulations were rhetorical devices as there was no doubt on the outcomes.

Surprisingly, de Heide and Grünwald (2020) use this simulation method—the one Rouder employed as a rhetorical device—to argue that under certain conditions Bayesian inference depends on the stopping plan. de Heide and Grünwald’s main message is about subjectivity and belief—that unless you *believe* your priors, optional stopping can distort inference. They then consider default priors of the type we have previously advocated. Their logic is that because no one could possibly believe default priors, they should not be used in optional stopping designs.

Our response is that de Heide and Grünwald’s critiques are not about optional stopping as much as they are about the proper Bayesian interpretation of Bayes factors. What de Heide and Grünwald desire is that Bayes factors not only describe the comparison of the specified models, but also describe comparisons among a broader class of models that were not specified. For example, if we are comparing a null model to an alternative, de Heide and Grünwald would like the same Bayes factor to apply to different priors on the alternative models. In general, Bayes factors can not provide these comparisons—they are limited to the specified models. Moreover, this limitation applies not only for optional-stopping designs but for fixed ones as well. And, expecting otherwise in general strikes us as magical thinking. And the distortions are not distortions but artifacts of frequentist definitions of misspecification.

To show what we mean about calibration with specified models, we quickly review how the simulations work. Figure 1A shows the results of the t -test example, and it is our replication of Example 2 in de Heide and Grünwald (2020). Here 10 observations are generated from one of two models, in this case the two models are

$$M_0 : Y_i \sim \text{Normal}(0, 1),$$

and

$$M_1 : Y_i | \mu \sim \text{Normal}(\mu, 1), \quad \mu \sim \text{Cauchy}(0, b),$$

where b is a scale factor we set to .5 that reflects a medium effect size. One characteristic of Bayesian models is that they provide a marginal probability distributions on data, and this probability distribution follows directly from the laws of probability (Jaynes, 2003). With this insight, we may draw data from these marginals each of two models, and compute the resulting posterior model odds. These are called the *nominal odds*, and are plotted on the x -axis. For each value of nominal odds, we may ask how many of the simulations came from one or the other model, and this proportion is then converted to odds, called the *observed odds*. These are plotted on the y -axis. In Figure 1A, the simulations are generated from Model M_0 and M_1 above, and the nominal posterior odds are the comparison of these two models for each simulated data set. Here the posterior odds are well calibrated in that nominal odds match the observed odds. It is this calibration that underlies the statement, *Bayes factors are indeed the comparison of models specified*. de Heide and Grunwald (2020) call this a “sanity check,” and we agree. It would be insane to use Bayes factors without this calibration property holding.

Perhaps surprisingly, this calibration holds regardless of sample plan. Figure 1B shows the case where sampling stops when the Bayes factor exceed 10 for either model up to 25 samples. Importantly, nominal and observed odds match. In fact, so long as the data are generated from the marginal distributions of the models, this calibration is guaranteed regardless of sampling plan. This guarantee should be viewed favorably by researchers who do not want to waste time with collecting extra samples after sufficient evidence has been gathered.

Now, suppose the data are not from the specified models. Instead, they are from models we did not specify. de Heide and Grunwald (2020), in their Example 2 used a point prior, a specific value of μ , to generate data from the alternative. Figure 2 shows the case for

a few different point values of μ . The best calibrated is $\mu = .5$, a medium effect (green circles). The calibration, while certainly not perfect, is pretty good. It is in this sense that de Heide and Grunwald (2020) claim the Cauchy prior is robust to this mismatch. But this choice of $\mu = .5$ is quite fortuitous. Also shown are two other choices, $\mu = 1.5$ and $\mu = .17$. Both of these point values are certainly compatible with the Cauchy model, and the value of $.17$ is even closer to median of the prior than the value of $.5$. Yet, there is dramatic miscalibration as the observed odds do not match the nominal odds. In general, Bayes factors are specific to the specified models. They do not describe evidence for models that are not specified.

The following facts are not in dispute: I. When the models of analysis are used to generate data, calibration holds and does so regardless of stopping plan. We see this behavior in de Heide and Grunwald (2020) Figures 1 through 3. II. For all cases where calibration fails, the analyst's models necessarily do not match the models for data generation. We see these failures in de Heide and Grunwald (2020) Figures 4 through 7, and this failure holds alike for fixed and optional stopping.

de Heide and Grunwald (2020) main claim is that the degree of failure of calibration may be smaller with fixed than optional stopping. We can see an example of this claim in Figure 2B which is highly similar to de Heide and Grunwald (2020) Figure 4. Here, the green circles are the results of optional stopping with $\mu = .5$. The miscalibration is more severe than the corresponding line in Panel A for the same simulation with fixed sampling.

We have two concerns. The first is that the claim is not in general true. Figure 2B also shows a counter example. It is the results of optional stopping with $\mu = 1.5$, the blue squares. The miscalibration here is less severe than the corresponding line in Panel A for the same simulation with fixed sampling. Our second concern is that focus on the difference in calibration between fixed and optional designs misses the the proverbial forest for the proverbial trees. If we should be concerned with miscalibration from wrong models, then the

difference between fixed and optional designs seems ancillary compared to the radical miscalibration in both types of designs. If the de Heide and Grunwald (2020) miscalibration is an indicator of misspecification or a violation of sanity, then, indeed, all analysis with Bayes factors looks insanely brittle.

The calibration of observed (simulated) and nominal odds cannot be used to understand misspecification. Consider the following example: Suppose we generate data from an alternative in the Cauchy setup with $\mu = .001$. This value is nearly null, but, logically, it is not. We simulate repeated runs and compute sample means for each run. Most of these sample means are near zero in value, and for these sample means, the Bayes factor between the null and the Cauchy alternative tends to favor the null. A small proportion of these sample means, however, are larger in magnitude, and the corresponding Bayes factors will favor the alternative. These are the nominal values, and they are, well, exactly what they should be. They are the relative strength of evidence from those sample means, whatever they are and whatever generated them, for models M_0 and M_1 . And this statement holds true *always*.

Now consider the observed odds. Sample means generated by the null, $\mu = 0$, and the alternative, $\mu = .001$, are highly similar in distribution. Consequently, so are the distributions of Bayes factors. When we condition on any particular interval, we find that the observed odds are 1.0. The resulting calibration curve is the horizontal dotted line in Figure 2A. This is radical miscalibration; in fact, there is no correlation between the observed and nominal odds. The problem is with the observed odds. They correspond to data generation mechanisms that are essentially the same, but not to the models under consideration. It is unclear to us how this mismatch implicates robustness to misspecification. Thinking that this mismatch should not occur implies an unreasonable, illogical expectation.

What are models? It is critical to understand a philosophical difference between Bayesians and frequentists. The frequentist analog to M_1 is $Y_i(\mu) \sim \text{Normal}(\mu, 1)$, that is

the random variables are functions of fixed and unknowable parameters. Frequentists play a “what if” game; they say, “what if I were God and knew the true value of μ in nature.” And then they characterize the error rates of various test procedures for hypothetical μ values, values that in reality they cannot know.

We tend to follow de Finetti (1974) who emphasizes that models are abstractions, and, consequently, parameters are abstractions too. There is no “true value” of a parameter in reality, because parameters exists in our head rather than in the world. What exists in the world are data, the Y_i in our examples. Unlike frequentist models, Bayesian models connect directly with Y_i because they specify the marginal distribution over data. And it is in this way they connect directly to reality without having to posit “true” parameters or even data generation processes. For Bayesians, the model is the combination of the likelihood and the prior, neither having preferred ontological status, and in concert defining a probability distribution on data.

In this views, models are not true or false data generation mechanisms. Instead they are hypothetical, platonic tools that help us query theoretically relevant structure in data (Box, 1976; Morey, Romeijn, & Rouder, 2013). As analysts, we start with competing theoretical positions. These positions are instantiated as models, and once these models are specified, we may perform inference, and then reformulate our inferential results about models as statements about theoretical positions (Rouder, Morey, & Wagenmakers, 2016). The inferential step may be performed by using Bayes factors to state relative evidence for the models. Relative evidence for models can then be interpreted as relative evidence for the theoretical positions so long as models are good instantiations of the positions.

What is robustness to misspecification? Robustness in de Heide and Grünwald—what happens if truth is different than the analysts’ models—is not helpful in the Bayesian framework. It is a frequentist notion, and we think it is irrelevant for evaluating Bayesian model comparison. The reason is that we never know the truth; hence, as Bayesians, we

should not condition on it. We prefer to explore robustness as the stability of the Bayes factor across reasonable variations in model specification. For example, when using default priors in linear models, there is a critical setting in the prior that sets the scale on expected effect sizes. For each application, there are values that are obviously too small (perhaps in psychology having effect sizes that are below .1) and there are values that are obviously too large (perhaps in psychology having effect sizes above 1). A researcher adds value by specifying a range of reasonable settings and then by assessing the stability of the Bayes factors across this range. If the same conclusion holds across the range, then the researcher has shown robustness to reasonable variations in model specification. If not, the data simply do not provide the resolution to adjudicate among the positions.

This view of robustness has a number of properties that are different from the frequentist view. First, robustness is dependent on data. We perform robustness conditional on the specific, observed data sets. A comparison may be robust to reasonable variation in model specification for one data set and not for another. In the frequentist setting, model robustness depends on hypothetical truth and is averaged across repeated data sets. We think in real-world cases, where researchers observe a single data set, this view of robustness conditional on observed data is most valuable. Second, our view of robustness depends on defining a reasonable range of variability in specification. This range reflects substantive concerns, say what is known about possible sizes of effects and the like. For example, suppose we are analyzing response-time data from tasks that take under one second to perform. When we check for robustness, we may check for scales that correspond to effects as small as 10 ms and as large as 100 ms (Haaf & Rouder, 2017). If results remain stable across these substantively defensible ranges, then we have shown robustness for the collected data. The key here is that the specification of ranges must reflect substantive considerations and can be defended or critiqued on this basis. Finally, these robustness checks may be used regardless of stopping plan.

We should comment on what de Heide and Grunwald (2020) call *strong calibration*. Strong is an understatement here. Perhaps it should be called *amazing calibration*. Strong/amazing calibration is when calibration holds even when a different model is used to produce the data. It comes into play in a tangential way in analysis. In default setups for linear models, there are some parameters that define the difference among the compared models and others that are common across them all. Under certain conditions, it is possible to define priors for which using the wrong model with respect to these common parameters does not affect the calibration. From our point-of-view, such a property has aesthetic appeal but does not impact our everyday use of Bayes factors. We care not if strong calibration is obeyed or not.

With these facts at hand—that the Bayes factor is well calibrated for the specified model regardless of stopping plan, and that practical robustness checks hold regardless of stopping plan—we conclude that Bayes factors do indeed *handle* optional stopping.

References

- Berger, J. O., & Wolpert, R. L. (1988). *The likelihood principle (2nd ed.)*. Hayward, CA: Institute of Mathematical Statistics.
- Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791–799.
- de Finetti, B. (1974). *Theory of probability* (Vol. 1). New York: John Wiley; Sons.
- de Heide, R., & Grunwald, P. D. (2020). Why optional stopping is a problem for bayesians. *Psychonomic Bulletin & Review*.
- Deng, A., Lu, J., & Chen, S. (2016). Continuous monitoring of a/b tests without pain: Optional stopping in bayesian testing. In *2016 ieee international conference on data science and advanced analytics (dsaa)* (pp. 243–252). IEEE.
- 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>
- Haaf, J. M., & Rouder, J. N. (2017). Developing constraint in Bayesian mixed models. *Psychological Methods*, 22(4), 779–798.
- Hendriksen, A., de Heide, R., & Grünwald, P. (2020). Optional stopping with bayes factors: A categorization and extension of folklore results, with an application to invariant situations. *arXiv Preprint arXiv:1807.09077*.
- Jaynes, E. (2003). *Probability theory: The logic of science*. Cambridge, UK: Cambridge University Press.
- Lindley, D. V. (1957). A statistical paradox. *Biometrika*, 44, 187–192.

- Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2013). The humble Bayesian: Model checking from a fully Bayesian perspective. *British Journal of Mathematical and Statistical Psychology*, *66*, 68–75. Retrieved from <http://dx.doi.org/10.1111/j.2044-8317.2012.02067.x>
- Rouder, J. N. (2014). Optional stopping: No problem for Bayesians. *Psychonomic Bulletin & Review*, *21*, 301–308. Retrieved from <http://dx.doi.org/10.3758/s13423-014-0595-4>
- 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>
- Sanborn, A. N., & Hills, T. T. (2014). The frequentist implications of optional stopping on Bayesian hypothesis tests. *Psychonomic Bulletin & Review*, *21*, 283–300.
- Tendeiro, J., Kiers, H., & van Ravenzwaaij, D. (2020). Mathematical evidence for the adequacy of bayesian optional stopping. *Psychonomic Bulletin & Review*.
- Yu, E. C., Sprenger, A. M., Thomas, R. P., & Dougherty, M. R. (2014). When decision heuristics and science collide. *Psychonomic Bulletin & Review*.

Figure 1. Observed odds as a function of nominal odds a null model vs. a two-tailed Cauchy alternative model. The match shows that the Bayes factor is well calibrated. A. Fixed design with $N = 10$. B. Optional-stopping design where sampling stops if the Bayes factor is over 10-to-1 for either model or a maximum sample size of $N = 25$.

Figure 2. Miscalibration results when data are not generated by the analyzed models. A. Three different point alternatives with a fixed sample size. B. Optional stopping for $\mu = 1.5$. Sampling stopped if the Bayes factors was 10-to-1 in favor of one of the models or a maximum sample size of $N = 25$.