UvA-DARE (Digital Academic Repository)

A tutorial on bridge sampling

Gronau, Q.F.; Sarafoglou, A.; Matzke, D.; Ly, A.; Boehm, U.; Marsman, M.; Leslie, D.S.; Forster, J.J.; Wagenmakers, E.-M.; Steingroever, H.

DOI

[10.1016/j.jmp.2017.09.005](https://doi.org/10.1016/j.jmp.2017.09.005)

Publication date

2017

Document Version

Final published version

Published in

Journal of Mathematical Psychology

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., Leslie, D. S., Forster, J. J., Wagenmakers, E.-M., & Steingroever, H. (2017). A tutorial on bridge sampling. *Journal of Mathematical Psychology*, 81, 80-97. <https://doi.org/10.1016/j.jmp.2017.09.005>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Tutorial

A tutorial on bridge sampling

Quentin F. Gronau^{a,*}, Alexandra Sarafoglou^a, Dora Matzke^a, Alexander Ly^a,
Udo Boehm^a, Maarten Marsman^a, David S. Leslie^b, Jonathan J. Forster^c,
Eric-Jan Wagenmakers^a, Helen Steingroever^a

^a Department of Psychology, University of Amsterdam, The Netherlands

^b Department Mathematics and Statistics, Lancaster University, UK

^c Mathematical Science, University of Southampton, UK

HIGHLIGHTS

- We provide a tutorial on bridge sampling for estimating marginal likelihoods.
- We use the beta-binomial model as a running example.
- We estimate the marginal likelihood for the Expectancy Valence (EV) model.
- We obtain accurate results for individual-level and hierarchical EV model versions.

ARTICLE INFO

Article history:

Received 17 March 2017

Received in revised form 31 August 2017

Available online 23 October 2017

Keywords:

Hierarchical model

Normalizing constant

Marginal likelihood

Bayes factor

Predictive accuracy

Reinforcement learning

ABSTRACT

The marginal likelihood plays an important role in many areas of Bayesian statistics such as parameter estimation, model comparison, and model averaging. In most applications, however, the marginal likelihood is not analytically tractable and must be approximated using numerical methods. Here we provide a tutorial on bridge sampling (Bennett, 1976; Meng & Wong, 1996), a reliable and relatively straightforward sampling method that allows researchers to obtain the marginal likelihood for models of varying complexity. First, we introduce bridge sampling and three related sampling methods using the beta-binomial model as a running example. We then apply bridge sampling to estimate the marginal likelihood for the Expectancy Valence (EV) model—a popular model for reinforcement learning. Our results indicate that bridge sampling provides accurate estimates for both a single participant and a hierarchical version of the EV model. We conclude that bridge sampling is an attractive method for mathematical psychologists who typically aim to approximate the marginal likelihood for a limited set of possibly high-dimensional models.

© 2017 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Bayesian statistics has become increasingly popular in mathematical psychology (Andrews & Baguley, 2013; Bayarri, Benjamin, Berger, & Sellke, 2016; Poirier, 2006; Vanpaemel, 2016; Verhagen, Levy, Millsap, & Fox, 2015; Wetzels, Tutschkow, Dolan, van der Sluis, Dutilh, & Wagenmakers, 2016). The Bayesian approach is conceptually simple, theoretically coherent, and easily applied to relatively complex problems. These problems include, for instance, hierarchical modeling (Matzke, Dolan, Batchelder, & Wagenmakers, 2015; Matzke & Wagenmakers, 2009; Rouder & Lu, 2005; Rouder, Lu, Speckman, Sun, & Jiang, 2005; Rouder, Lu, Sun, Speckman, Morey, & Naveh-Benjamin, 2007) or the comparison of non-nested models (Lee, 2008; Pitt, Myung, & Zhang, 2002;

Shiffrin, Lee, Kim, & Wagenmakers, 2008). Three major applications of Bayesian statistics concern parameter estimation, model comparison, and Bayesian model averaging. In all three areas, the marginal likelihood – that is, the probability of the observed data given the model of interest – plays a central role (see also Gelman & Meng, 1998).

First, in parameter estimation, we consider a single model and aim to quantify the uncertainty for a parameter of interest θ after having observed the data y . This is realized by means of a posterior distribution that can be obtained using Bayes theorem:

$$p(\theta | y) = \frac{p(y | \theta) p(\theta)}{\int p(y | \theta') p(\theta') d\theta'} = \frac{\overbrace{p(y | \theta)}^{\text{likelihood}} \overbrace{p(\theta)}^{\text{prior}}}{\underbrace{p(y)}_{\text{marginal likelihood}}} \quad (1)$$

* Correspondence to: Department of Psychology, PO Box 15906, 1001 NK Amsterdam, The Netherlands.

E-mail address: Quentin.F.Gronau@gmail.com (Q.F. Gronau).

Here, the marginal likelihood of the data $p(y)$ ensures that the posterior distribution is a proper probability density function (PDF) in the sense that it integrates to 1. This illustrates why in parameter estimation the marginal likelihood is referred to as a normalizing constant.

Second, in model comparison, we consider m ($m \in \mathbb{N}$) competing models, and are interested in the relative plausibility of a particular model $\mathcal{M}_i$ ($i \in \{1, 2, \dots, m\}$) given the prior model probability and the evidence from the data y (see three special issues on this topic in the *Journal of Mathematical Psychology*: Mulder & Wagenmakers, 2016; Myung, Forster, & Browne, 2000; Wagenmakers & Waldorp, 2006). This relative plausibility is quantified by the so-called posterior model probability $p(\mathcal{M}_i | y)$ of model $\mathcal{M}_i$ given the data y (Berger & Molina, 2005):

$$p(\mathcal{M}_i | y) = \frac{p(y | \mathcal{M}_i) p(\mathcal{M}_i)}{\sum_{j=1}^m p(y | \mathcal{M}_j) p(\mathcal{M}_j)}, \quad (2)$$

where the denominator is the sum of the marginal likelihood times the prior model probability of all m models. In model comparison, the marginal likelihood for a specific model is also referred to as the model evidence (Didelot, Everitt, Johansen, & Lawson, 2011), the integrated likelihood (Kass & Raftery, 1995), the predictive likelihood of the model (Gamerman & Lopes, 2006, chapter 7), the predictive probability of the data (Kass & Raftery, 1995), or the prior predictive density (Ntzoufras, 2009). Note that conceptually the marginal likelihood of Eq. (2) is the same as the marginal likelihood of Eq. (1). However, for the latter equation we dropped the model index because in parameter estimation we consider only one model.

If only two models $\mathcal{M}_1$ and $\mathcal{M}_2$ are considered, Eq. (2) can be used to quantify the relative posterior model plausibility of model $\mathcal{M}_1$ compared to model $\mathcal{M}_2$. This relative plausibility is given by the ratio of the posterior probabilities of both models, and is referred to as the posterior model odds:

$$\frac{p(\mathcal{M}_1 | y)}{p(\mathcal{M}_2 | y)} = \underbrace{\frac{p(\mathcal{M}_1)}{p(\mathcal{M}_2)}}_{\text{prior odds}} \times \underbrace{\frac{p(y | \mathcal{M}_1)}{p(y | \mathcal{M}_2)}}_{\text{Bayes factor}}. \quad (3)$$

Eq. (3) illustrates that the posterior model odds are the product of two factors: The first factor is the ratio of the prior probabilities of both models—the prior model odds. The second factor is the ratio of the marginal likelihoods of both models—the so-called Bayes factor (Etz & Wagenmakers, in press; Jeffreys, 1961; Ly, Verhagen, & Wagenmakers, 2016a, b; Robert, 2016). The Bayes factor plays an important role in model comparison and is referred to as the “standard Bayesian solution to the hypothesis testing and model selection problems” (Lewis & Raftery, 1997, p. 648) and “the primary tool used in Bayesian inference for hypothesis testing and model selection” (Berger, 2006, p. 378).

Third, the marginal likelihood plays an important role in Bayesian model averaging (BMA; Hoeting, Madigan, Raftery, & Volinsky, 1999) where aspects of parameter estimation and model comparison are combined. As in model comparison, BMA considers several models; however, it does not aim to identify a single best model. Instead it fully acknowledges model uncertainty. Model averaged parameter inference can be obtained by combining, across all models, the posterior distribution of the parameter of interest weighted by each model’s posterior model probability, and as such depends on the marginal likelihood of the models. This procedure assumes that the parameter of interest has identical interpretation across the different models. Model averaged predictions can be obtained in a similar manner.

A problem that arises in all three areas – parameter estimation, model comparison, and BMA – is that an analytical expression of the marginal likelihood can be obtained only for certain restricted

examples. This is a pressing problem in Bayesian modeling, and in particular in mathematical psychology where models can be non-linear and equipped with a large number of parameters, especially when the models are implemented in a hierarchical framework. Such a framework incorporates both commonalities and differences between participants of one group by assuming that the model parameters of each participant are drawn from a group-level distribution (for advantages of the Bayesian hierarchical framework see Ahn, Krawitz, Kim, Bussemeyer, & Brown, 2011; Navarro, Griffiths, Steyvers, & Lee, 2006; Rouder & Lu, 2005; Rouder et al., 2005; Rouder, Lu, Morey, Sun, & Speckman, 2008; Scheibehenne & Pachur, 2015; Shiffrin et al., 2008; Wetzels, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2010). For instance, consider a four-parameter Bayesian hierarchical model with four group-level distributions each characterized by two parameters and a group size of 30 participants; this then results in 30×4 individual-level parameters and 2×4 group-level parameters for a total of 128 parameters. In sum, even simple models quickly become complex once hierarchical aspects are introduced and this frustrates the derivation of the marginal likelihood.

To overcome this problem, several Monte Carlo sampling methods have been proposed to approximate the marginal likelihood. In this tutorial we focus on four such methods: the bridge sampling estimator (Bennett, 1976, Chapter 5 of Chen, Shao, & Ibrahim, 2012; Meng & Wong, 1996), and its three commonly used special cases—the naive Monte Carlo estimator, the importance sampling estimator, and the generalized harmonic mean estimator (for alternative methods see Gamerman & Lopes, 2006, Chapter 7; and for alternative approximation methods relevant to model comparison and BMA see Carlin & Chib, 1995; Green, 1995).¹ As we will illustrate throughout this tutorial, the bridge sampler is accurate, efficient, and relatively straightforward to implement (e.g., DiCiccio, Kass, Raftery, & Wasserman, 1997; Frühwirth-Schnatter, 2004; Meng & Wong, 1996).

The goal of this tutorial is to bring the bridge sampling estimator to the attention of mathematical psychologists. We aim to explain this estimator and facilitate its application by suggesting a step-by-step implementation scheme. To this end, we first show how bridge sampling and the three special cases can be used to approximate the marginal likelihood in a simple beta-binomial model. We begin with the naive Monte Carlo estimator and progressively work our way up – via the importance sampling estimator and the generalized harmonic mean estimator – to the most general case considered: the bridge sampling estimator. This order was chosen such that key concepts are introduced gradually and estimators are of increasing complexity and sophistication. The first three estimators are included in this tutorial with the sole purpose of facilitating the reader’s understanding of bridge sampling. In the second part of this tutorial, we outline how the bridge sampling estimator can be used to derive the marginal likelihood for the Expectancy Valence (EV; Bussemeyer & Stout, 2002) model—a popular, yet relatively complex reinforcement-learning model for the Iowa gambling task (Bechara, Damasio, Damasio, & Anderson, 1994). We apply bridge sampling to both an individual-level and a hierarchical implementation of the EV model.

Throughout the article, we use the software package R to implement the bridge sampling estimator for the various models. The interested reader is invited to reproduce our results by downloading the code and all relevant materials from our Open Science Framework folder at osf.io/f9cq4.

¹ The Appendix provides a derivation showing that the first three estimators are indeed special cases of the bridge sampler.

Fig. 1. Prior and posterior distribution for the rate parameter θ from the beta-binomial model. The Beta(1, 1) prior on the rate parameter θ is represented by the dotted line; the Beta(3, 9) posterior distribution is represented by the solid line and was obtained after having observed 2 correct responses out of 10 trials. Available at <https://tinyurl.com/yc8bw98v> under CC license <https://creativecommons.org/licenses/by/2.0/>.

1. Four sampling methods to approximate the marginal likelihood

In this section we outline four standard methods to approximate the marginal likelihood. For more detailed explanations and derivations, we recommend [Ntzoufras \(2009, Chapter 11\)](#) and [Gamerman and Lopes \(2006, Chapter 7\)](#); a comparative review of the different sampling methods is presented in [DiCiccio et al. \(1997\)](#). The marginal likelihood is the probability of the observed data y given a specific model of interest $\mathcal{M}$, and is defined as the integral of the likelihood over the prior:

$$\underbrace{p(y | \mathcal{M})}_{\text{marginal likelihood}} = \int \underbrace{p(y | \theta, \mathcal{M})}_{\text{likelihood}} \underbrace{p(\theta | \mathcal{M})}_{\text{prior}} d\theta, \quad (4)$$

with θ a vector containing the model parameters. Eq. (4) illustrates that the marginal likelihood can be interpreted as a weighted average of the likelihood of the data given a specific value for θ where the weight is the a priori plausibility of that specific value. Eq. (4) can therefore be written as an expected value:

$$p(y | \mathcal{M}) = \mathbb{E}_{\text{prior}} [p(y | \theta, \mathcal{M})],$$

where the expectation is taken with respect to the prior distribution. This idea is central to the four sampling methods that we discuss in this tutorial.

1.1. Introduction of the running example: the beta-binomial model

Our running example focuses on estimating the marginal likelihood for a binomial model assuming a uniform prior on the rate parameter θ (i.e., the beta-binomial model). Consider a single participant who answered $k = 2$ out of $n = 10$ true/false questions correctly. Assume that the number of correct answers follows a binomial distribution, that is, $k \sim \text{Binomial}(n, \theta)$ with $\theta \in (0, 1)$, where θ represents the latent probability for answering any one question correctly. The probability mass function (PMF) of the binomial distribution is given by:

$$\text{Binomial}(k | n, \theta) = \binom{n}{k} \theta^k (1 - \theta)^{n-k}, \quad (5)$$

where $k, n \in \mathbb{Z}_{\geq 0}$, and $k \leq n$. The PMF of the binomial distribution serves as the likelihood function in our running example.

In the Bayesian framework, we also have to specify the prior distribution of the model parameters; the prior distribution expresses our knowledge about the parameters before the data have been

observed. In our running example, we assume that all values of θ are equally likely a priori. This prior belief is captured by a uniform distribution across the range of θ , that is, $\theta \sim \text{Uniform}(0, 1)$ which can equivalently be written in terms of a beta distribution $\theta \sim \text{Beta}(1, 1)$. This prior distribution is represented by the dotted line in Fig. 1. It is evident that the density of the prior distribution equals 1 for all values of θ . One advantage of expressing the prior distribution by a beta distribution is that its two parameters (i.e., in its general form the shape parameters α and β) can be thought of as counts of “prior successes” and “prior failures”, respectively. In its general form, the PDF of a Beta(α, β) distribution ($\alpha, \beta > 0$) is given by:

$$\text{Beta}(\theta; \alpha, \beta) = \frac{\theta^{\alpha-1} (1 - \theta)^{\beta-1}}{B(\alpha, \beta)},$$

where $B(\alpha, \beta)$ is the beta function that is defined as: $B(\alpha, \beta) = \int_0^1 t^{\alpha-1} (1 - t)^{\beta-1} dt = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$, and $\Gamma(n) = (n - 1)!$ for $n \in \mathbb{N}$.

1.1.1. Analytical derivation of the marginal likelihood

As we will see in this section, the beta-binomial model constitutes one of the rare examples where the marginal likelihood is analytic. Assuming a general k and n , we obtain the marginal likelihood as:

$$\begin{aligned} p(k | n) &\stackrel{\text{Eq. (4)}}{=} \int_0^1 p(k | n, \theta) p(\theta) d\theta = \int_0^1 \binom{n}{k} \theta^k (1 - \theta)^{n-k} 1 d\theta \\ &= \binom{n}{k} B(k + 1, n - k + 1) = \frac{1}{n + 1}, \end{aligned} \quad (6)$$

where we suppress the “model” in the conditioning part of the probability statements because we focus on a single model in this running example. Using $k = 2$ and $n = 10$ of our example, we obtain: $p(k = 2 | n = 10) = 1/11 \approx 0.0909$. This value will be estimated in the remainder of the running example using the naive Monte Carlo estimator, the importance sampling estimator, the generalized harmonic mean estimator, and finally the bridge sampling estimator.

As we will see below, the importance sampling, generalized harmonic mean estimator, and bridge sampling estimator require samples from the posterior distribution. These samples can be obtained using computer software such as WinBUGS ([Lunn, Thomas, Best, & Spiegelhalter, 2000](#)), JAGS ([Plummer, 2003](#)), or Stan ([Stan Development Team, 2016](#)), even when the marginal likelihood that functions here as a normalizing constant is not known (Eq. (1)). However, in our running example MCMC samples are not required because we can derive an analytical expression of the posterior distribution for θ after having observed the data. Using the analytic expression of the marginal likelihood (Eq. (6)) and Bayes theorem, we obtain:

$$\begin{aligned} p(\theta | k, n) &= \frac{p(k | n, \theta) p(\theta)}{p(k | n)} = \frac{\binom{n}{k} \theta^k (1 - \theta)^{n-k} 1}{\binom{n}{k} B(k + 1, n - k + 1)} \\ &= \frac{\theta^k (1 - \theta)^{n-k}}{B(k + 1, n - k + 1)}, \end{aligned}$$

which we recognize as the PDF of the Beta($k + 1, n - k + 1$) distribution. Thus, if we assume a uniform prior on θ and observe $k = 2$ correct responses out of $n = 10$ trials, we obtain a Beta(3, 9) distribution as posterior distribution. This distribution is represented by the solid line in Fig. 1. In general, if $k | n, \theta \sim \text{Binomial}(n, \theta)$ and $\theta \sim \text{Beta}(1, 1)$, then $\theta | n, k \sim \text{Beta}(k + 1, n - k + 1)$.

1.2. Method 1: the naive Monte Carlo estimator of the marginal likelihood

The simplest method to approximate the marginal likelihood is provided by the naive Monte Carlo estimator ([Hammersley &](#)

Fig. 2. Illustration of the naive Monte Carlo estimator for the beta-binomial example. The dotted line represents the prior distribution and the solid line represents the posterior distribution that was obtained after having observed 2 correct responses out of 10 trials. The gray dots represent the 12 samples $\{\tilde{\theta}_1, \tilde{\theta}_2, \dots, \tilde{\theta}_{12}\}$ randomly drawn from the $\text{Beta}(1, 1)$ prior distribution. Available at <https://tinyurl.com/y8uf6t8f> under CC license <https://creativecommons.org/licenses/by/2.0/>.

Handscamb, 1964; Raftery & Banfield, 1991). This method uses the standard definition of the marginal likelihood (Eq. (4)), and relies on the central idea that the marginal likelihood can be written as an expected value with respect to the prior distribution, that is, $p(y) = \mathbb{E}_{\text{prior}}[p(y | \theta)]$. This expected value of the likelihood of the data with respect to the prior can be approximated by evaluating the likelihood in N samples from the prior distribution for θ and averaging the resulting values. This yields the naive Monte Carlo estimator $\hat{p}_1(y)$:

$$\hat{p}_1(y) = \frac{1}{N} \sum_{i=1}^N \underbrace{p(y | \tilde{\theta}_i)}_{\text{average likelihood}}, \quad \underbrace{\tilde{\theta}_i \sim p(\theta)}_{\text{samples from the prior distribution}}. \quad (7)$$

1.2.1. Running example

To obtain the naive Monte Carlo estimate of the marginal likelihood in our running example, we need N samples from the $\text{Beta}(1, 1)$ prior distribution for θ . For illustrative purposes, we limit the number of samples to 12 whereas in practice one should take N to be very large. We obtain the following samples:

$$\{\tilde{\theta}_1, \tilde{\theta}_2, \dots, \tilde{\theta}_{12}\} = \{0.58, 0.76, 0.03, 0.93, 0.27, 0.97, 0.45, 0.46, 0.18, 0.64, 0.06, 0.15\},$$

where we use the tilde symbol to emphasize that we refer to a sampled value. All sampled values are represented by the gray dots in Fig. 2.

Following Eq. (7), the next step is to calculate the likelihood (Eq. (5)) for each $\tilde{\theta}_i$, and then to average all obtained likelihood values. This yields the naive Monte Carlo estimate of the marginal likelihood:

$$\begin{aligned} \hat{p}_1(k = 2 | n = 10) &= \frac{1}{12} \sum_{i=1}^{12} p(k = 2 | n = 10, \tilde{\theta}_i) \\ &= \frac{1}{12} \sum_{i=1}^{12} \binom{n}{k} (\tilde{\theta}_i)^k (1 - \tilde{\theta}_i)^{n-k} \\ &= \frac{1}{12} \binom{10}{2} (0.58^2 (1 - 0.58)^8 + \dots \\ &\quad + 0.15^2 (1 - 0.15)^8) \\ &= 0.0945. \end{aligned}$$

1.3. Method 2: the importance sampling estimator of the marginal likelihood

The naive Monte Carlo estimator introduced in the last section performs well if the prior and posterior distribution have a similar shape and strong overlap. However, the estimator is unstable if the posterior distribution is peaked relative to the prior (e.g., Gaman & Lopes, 2006; Ntzoufras, 2009). In such a situation, most of the sampled values for θ result in likelihood values close to zero and contribute only minimally to the estimate. This means that those few samples that result in high likelihood values dominate estimates of the marginal likelihood. Consequently, the variance of the estimator is increased (Newton & Raftery, 1994; Pajor, 2016).²

The importance sampling estimator, on the other hand, overcomes this shortcoming by boosting sampled values in regions of the parameter space where the integrand of Eq. (4) is large. This is realized by using samples from a so-called importance density $g_{IS}(\theta)$ instead of the prior distribution. The advantage of sampling from an importance density is that values for θ that result in high likelihood values are sampled most frequently, whereas values for θ with low likelihood values are sampled only rarely.

To derive the importance sampling estimator, Eq. (4) is used as starting point which is then extended by the importance density $g_{IS}(\theta)$:

$$\begin{aligned} p(y) &= \int p(y | \theta) p(\theta) d\theta = \int p(y | \theta) p(\theta) \frac{g_{IS}(\theta)}{g_{IS}(\theta)} d\theta \\ &= \int \frac{p(y | \theta) p(\theta)}{g_{IS}(\theta)} g_{IS}(\theta) d\theta \\ &= \mathbb{E}_{g_{IS}(\theta)} \left(\frac{p(y | \theta) p(\theta)}{g_{IS}(\theta)} \right). \end{aligned}$$

This yields the importance sampling estimator $\hat{p}_2(y)$:

$$\hat{p}_2(y) = \frac{1}{N} \sum_{i=1}^N \underbrace{\frac{p(y | \tilde{\theta}_i) p(\tilde{\theta}_i)}{g_{IS}(\tilde{\theta}_i)}}_{\text{average adjusted likelihood}}, \quad \underbrace{\tilde{\theta}_i \sim g_{IS}(\theta)}_{\text{samples from the importance density}}. \quad (8)$$

A suitable importance density should (1) be easy to evaluate; (2) have the same domain as the posterior distribution; (3) closely resemble the posterior distribution; and (4) have fatter tails than the posterior distribution (Neal, 2001; Vandekerckhove, Matzke, & Wagenmakers, 2015). The latter criterion ensures that values in the tails of the distribution cannot misleadingly dominate the estimate (Neal, 2001).³

1.3.1. Running example

To obtain the importance sampling estimate of the marginal likelihood in our running example, we first need to choose an importance density $g_{IS}(\theta)$. An importance density that fulfills the four above mentioned desiderata is a mixture between a beta density that provides the best fit to the posterior distribution and a uniform density across the range of θ (Vandekerckhove et al.,

² The interested reader is referred to Pajor (2016) for a recent improvement on the calculation of the naive Monte Carlo estimator. The proposed improvement involves trimming the prior distribution in such a way that regions with low likelihood values are eliminated, thereby increasing the accuracy and efficiency of the estimator.

³ To illustrate the need for an importance density with fatter tails than the posterior distribution, imagine you sample from the tail region of an importance density with thinner tails. In this case, the numerator in Eq. (8) would be substantially larger than the denominator resulting in a very large ratio. Since this specific ratio is only one component of the sum displayed in Eq. (8), this component would dominate the importance sampling estimate. Hence, thinner tails of the importance density run the risk of producing unstable estimates across repeated computations. In fact, the estimator may have infinite variance (e.g., Ionides, 2008; Owen & Zhou, 2000).

2015). The relative impact of the uniform density is quantified by a mixture weight γ that ranges between 0 and 1. The larger γ , the higher the influence of the uniform density resulting in a less peaked distribution with thick tails. If $\gamma = 1$, the beta mixture density simplifies to the uniform distribution on $[0, 1]$,⁴ and if $\gamma = 0$, the beta mixture density simplifies to the beta density that provides the best fit to the posterior distribution.

In our specific example, we already know that the Beta(3, 9) density is the beta density that provides the best fit to the posterior distribution because this is the analytic expression of the posterior distribution. However, to demonstrate the general case, we show how we can find the beta distribution with the best fit to the posterior distribution using the method of moments. This particular method works as follows. First, we draw samples from our Beta(3, 9) posterior distribution and obtain:⁵

$$\{\theta_1^*, \theta_2^*, \dots, \theta_{12}^*\} = \{0.22, 0.16, 0.09, 0.35, 0.06, 0.27, 0.26, 0.41, 0.20, 0.43, 0.21, 0.12\}.$$

Note that here we use θ_i^* to refer to the i th sample from the posterior distribution to distinguish it from the previously used $\tilde{\theta}_i$ —the i th sample from a distribution other than the posterior distribution, such as a prior distribution or an importance density. Second, we compute the mean and variance of these posterior samples. We obtain a mean of $\bar{\theta}^* = 0.232$ and a variance of $s_{\theta^*}^2 = 0.014$.

Third, knowing that, if $X \sim \text{Beta}(\alpha, \beta)$, then $\mathbb{E}(X) = \alpha/(\alpha + \beta)$ and $V(X) = \alpha\beta/[(\alpha + \beta)^2(\alpha + \beta + 1)]$, we obtain the following method of moment estimates for α and β :

$$\hat{\alpha} = \bar{\theta}^* \left(\frac{\bar{\theta}^*(1 - \bar{\theta}^*)}{s_{\theta^*}^2} - 1 \right) = 0.232 \left(\frac{0.232(1 - 0.232)}{0.014} - 1 \right) = 2.721,$$

$$\hat{\beta} = (1 - \bar{\theta}^*) \left(\frac{\bar{\theta}^*(1 - \bar{\theta}^*)}{s_{\theta^*}^2} - 1 \right) = (1 - 0.232) \left(\frac{0.232(1 - 0.232)}{0.014} - 1 \right) = 9.006.$$

Using a mixture weight on the uniform component of $\gamma = 0.30$ – a choice that was made to ensure that, visually, the tails of the importance density are clearly thicker than the tails of the posterior distribution – we obtain the following importance density: $\gamma \times \text{Beta}(\theta; 1, 1) + (1 - \gamma) \times \text{Beta}(\theta; \hat{\alpha}, \hat{\beta}) = .3 + .7 \text{Beta}(\theta; 2.721, 9.006)$. This importance density is represented by the dashed line in Fig. 3. The figure also shows the posterior distribution (solid line). As is evident from the figure, the beta mixture importance density resembles the posterior distribution, but has fatter tails.

In general, it is advised to choose the mixture weight on the uniform component γ small enough to make the estimator efficient, yet large enough to produce fat tails to stabilize the estimator. A suitable mixture weight can be realized by gradually minimizing

Fig. 3. Illustration of the importance sampling estimator for the beta-binomial model. The dashed line represents our beta mixture importance density and the solid gray line represents the posterior distribution that was obtained after having observed 2 correct responses out of 10 trials. The gray dots represent the 12 samples $\{\tilde{\theta}_1, \tilde{\theta}_2, \dots, \tilde{\theta}_{12}\}$ randomly drawn from our beta mixture importance density. Available at <https://tinyurl.com/yc7ho7hr> under CC license <https://creativecommons.org/licenses/by/2.0/>.

the mixture weight and investigating whether stability is still guaranteed (i.e., robustness analysis).

Drawing $N = 12$ samples for θ from our beta mixture importance density results in:

$$\{\tilde{\theta}_1, \tilde{\theta}_2, \dots, \tilde{\theta}_{12}\} = \{0.11, 0.07, 0.32, 0.25, 0.41, 0.39, 0.25, 0.13, 0.64, 0.26, 0.74, 0.92\}.$$

These samples are represented by the gray dots in Fig. 3.

The final step is to compute the average adjusted likelihood for the 12 samples using Eq. (8). This yields the importance sampling estimate of the marginal likelihood as:

$$\begin{aligned} \hat{p}_2(k = 2 | n = 10) &= \frac{1}{12} \sum_{i=1}^{12} \frac{p(k = 2 | n = 10, \tilde{\theta}_i) p(\tilde{\theta}_i)}{.3 + .7 \text{Beta}(\tilde{\theta}_i; 2.721, 9.006)} \\ &= \frac{1}{12} \left(\frac{\binom{10}{2} 0.11^2 (1 - 0.11)^8 \times 1}{.3 + .7 \text{Beta}(0.11; 2.721, 9.006)} \right. \\ &\quad \left. + \dots + \frac{\binom{10}{2} 0.92^2 (1 - 0.92)^8 \times 1}{.3 + .7 \text{Beta}(0.92; 2.721, 9.006)} \right) \\ &= \frac{1}{12} \binom{10}{2} (0.0021 + \dots + 7.3 \times 10^{-9}) \\ &= 0.0827. \end{aligned}$$

1.4. Method 3: the generalized harmonic mean estimator of the marginal likelihood

Just as the importance sampling estimator, the generalized harmonic mean estimator focuses on regions of the parameter space where the integrand of Eq. (4) is large by using an importance density $g_{IS}(\theta)$ (Gelfand & Dey, 1994).⁶ However, in contrast to the importance sampling estimator, the generalized harmonic mean estimator requires an importance density with thinner tails for an analogous reason as in importance sampling.

⁴ In our running example, the importance sampling estimator then reduces to the naive Monte Carlo estimator.

⁵ Note that, when the analytical expression of the posterior distribution is not known, posterior samples can be obtained using computer software such as WinBUGS, JAGS, or Stan, even when the marginal likelihood that functions here as a normalizing constant is not known (Eq. (1)).

⁶ Note that the generalized harmonic mean estimator is a more stable version of the harmonic mean estimator (Newton & Raftery, 1994). A problem of the harmonic mean estimator is that it is dominated by the samples that have small likelihood values.

To derive the generalized harmonic mean estimator, also known as reciprocal importance sampling estimator (Frühwirth-Schnatter, 2004), we use the following identity:

$$\begin{aligned} \frac{1}{p(y)} &= \int \frac{1}{p(y)} g_{IS}(\theta) d\theta = \int \frac{p(\theta | y)}{p(y | \theta)p(\theta)} g_{IS}(\theta) d\theta \\ &= \int \frac{g_{IS}(\theta)}{p(y | \theta)p(\theta)} p(\theta | y) d\theta \\ &= \mathbb{E}_{\text{post}} \left(\frac{g_{IS}(\theta)}{p(y | \theta)p(\theta)} \right). \end{aligned}$$

Rewriting results in:

$$p(y) = \left(\mathbb{E}_{\text{post}} \left(\frac{g_{IS}(\theta)}{p(y | \theta)p(\theta)} \right) \right)^{-1},$$

which is used to define the generalized harmonic mean estimator $\hat{p}_3(y)$ (Gelfand & Dey, 1994) as follows:

$$\hat{p}_3(y) = \left(\frac{1}{N} \sum_{j=1}^N \frac{\overbrace{g_{IS}(\theta_j^*)}^{\text{importance density}}}{\underbrace{p(y | \theta_j^*)}_{\text{likelihood}} \underbrace{p(\theta_j^*)}_{\text{prior}}} \right)^{-1}, \quad (9)$$

$\theta_j^* \sim p(\theta | y)$
samples from the posterior distribution

Note that the generalized harmonic mean estimator – in contrast to the importance sampling estimator – evaluates samples from the posterior distribution. In addition, note that the ratio in Eq. (9) is the reciprocal of the ratio in Eq. (8); this explains why the importance density for the generalized harmonic mean estimator should have thinner tails than the posterior distribution in order to avoid inflation of the ratios that are part of the summation displayed in Eq. (9). Thus, in the case of the generalized harmonic mean estimator, a suitable importance density should (1) have thinner tails than the posterior distribution (DiCiccio et al., 1997; Newton & Raftery, 1994), and as in importance sampling, it should (2) be easy to evaluate; (3) have the same domain as the posterior distribution; and (4) closely resemble the posterior distribution.

1.4.1. Running example

To obtain the generalized harmonic mean estimate of the marginal likelihood in our running example, we need to choose a suitable importance density. In our running example, an importance density that fulfills the four above mentioned desiderata can be obtained by following four steps: First, we draw $N = 12$ samples from the posterior distribution. Reusing the samples from the last section, we obtain:

$$\{\theta_1^*, \theta_2^*, \dots, \theta_{12}^*\} = \{0.22, 0.16, 0.09, 0.35, 0.06, 0.27, 0.26, 0.41, 0.20, 0.43, 0.21, 0.12\}.$$

Second, we probit-transform all posterior samples (i.e., $\xi_j^* = \Phi^{-1}(\theta_j^*)$, with $j \in \{1, 2, \dots, 12\}$).⁷ The result of this transformation is that the samples range across the entire real line instead of the (0, 1) interval only. We obtain:

$$\{\xi_1^*, \xi_2^*, \dots, \xi_{12}^*\} = \{-0.77, -0.99, -1.34, -0.39, -1.55, -0.61, -0.64, -0.23, -0.84, -0.18, -0.81, -1.17\}.$$

These probit-transformed samples are represented by the gray dots in Fig. 4.

Third, we search for the normal distribution that provides the best fit to the probit-transformed posterior samples ξ_j^* . Using the

method of moments, we obtain as estimates $\hat{\mu} = -0.793$ and $\hat{\sigma} = 0.423$. Note that the choice of a normal importance density justifies step 2; the probit transformation (or an equivalent transformation) was required to match the range of the posterior distribution to the one of the normal distribution.

Finally, as importance density we choose a normal distribution with mean $\hat{\mu} = -0.793$ and standard deviation $\hat{\sigma} = 0.423/1.5$. This additional division by 1.5 is to ensure thinner tails of the importance density than of the probit-transformed posterior distribution (for a discussion of alternative importance densities see DiCiccio et al., 1997). We decided to divide $\hat{\sigma}$ by 1.5 for illustrative purposes only. Our importance density is displayed in Fig. 4 (dashed line) together with the probit-transformed posterior distribution (solid line).

The generalized harmonic mean estimate can now be obtained using either the original posterior samples θ_j^* or the probit-transformed samples ξ_j^* . Here we use the latter ones (see also Overstall & Forster, 2010). Incorporating our specific importance density and a correction for having used the probit-transformation, Eq. (9) becomes:⁸

$$\hat{p}_3(y) = \left(\frac{1}{N} \sum_{j=1}^N \frac{\overbrace{\frac{1}{\hat{\sigma}} \phi \left(\frac{\xi_j^* - \hat{\mu}}{\hat{\sigma}} \right)}^{\text{importance density}}}{\underbrace{p(y | \Phi(\xi_j^*))}_{\text{likelihood}} \underbrace{\phi(\xi_j^*)}_{\text{prior}}} \right)^{-1}, \quad (10)$$

$\xi_j^* = \Phi^{-1}(\theta_j^*)$ and $\theta_j^* \sim p(\theta | y)$
probit-transformed samples from the posterior distribution

For our beta-binomial model, we now obtain the generalized harmonic mean estimate of the marginal likelihood as:

$$\begin{aligned} \hat{p}_3(k=2 | n=10) &= \left(\frac{1}{12} \sum_{j=1}^{12} \frac{\frac{1}{0.423/1.5} \phi \left(\frac{\xi_j^* + 0.793}{0.423/1.5} \right)}{p(k=2 | n=10, \Phi(\xi_j^*)) \phi(\xi_j^*)} \right)^{-1} \\ &= \left(\frac{1}{12} \left(\frac{\frac{1}{0.423/1.5} \phi \left(\frac{-0.77+0.793}{0.423/1.5} \right)}{\binom{10}{2} 0.22^2 (1-0.22)^8 \phi(-0.77)} \right. \right. \\ &\quad \left. \left. + \dots + \frac{\frac{1}{0.423/1.5} \phi \left(\frac{-1.17+0.793}{0.423/1.5} \right)}{\binom{10}{2} 0.12^2 (1-0.12)^8 \phi(-1.17)} \right) \right)^{-1} \\ &= \left(\frac{1}{12} \frac{1}{\binom{10}{2}} (716.89 + \dots + 555.50) \right)^{-1} \\ &= 0.092. \end{aligned}$$

1.5. Method 4: the bridge sampling estimator of the marginal likelihood

As became evident in the last two sections, both the importance sampling estimator and the generalized harmonic mean estimator impose strong constraints on the tail behavior of the importance density relative to the posterior distribution to guarantee a stable estimator. Such requirements can make it difficult to find a suitable importance density, especially when a high dimensional posterior is considered. The bridge sampler, on the other hand, alleviates such requirements (e.g., Frühwirth-Schnatter, 2004).

⁸ A detailed explanation is provided in the Appendix. Note that using the original posterior samples θ_j^* would involve transforming the importance density (e.g., the normal density on ξ) to the (0, 1) interval.

⁷ Other transformation are conceivable (e.g., logit transformation).

Fig. 4. Illustration of the generalized harmonic mean estimator for the beta-binomial model. The solid line represents the probit-transformed Beta(3, 9) posterior distribution that was obtained after having observed 2 correct responses out of 10 trials, and the dashed line represents the importance density $\mathcal{N}(\xi; \mu = -0.793, \sigma = 0.423/1.5)$. The gray dots represent the 12 probit-transformed samples $\{\xi_1^*, \xi_2^*, \dots, \xi_{12}^*\}$ randomly drawn from the Beta(3, 9) posterior distribution. Available at <https://tinyurl.com/yazgk8kj> under CC license <https://creativecommons.org/licenses/by/2.0/>.

Originally, bridge sampling was developed to directly estimate the Bayes factor, that is, the ratio of the marginal likelihoods of two models $\mathcal{M}_1$ and $\mathcal{M}_2$ (e.g., Jeffreys, 1961; Kass & Raftery, 1995). However, in this tutorial, we use a version of bridge sampling that allows us to approximate the marginal likelihood of a single model (for an earlier application see for example Overstall & Forster, 2010). This version is based on the following identity:

$$1 = \frac{\int p(y | \theta) p(\theta) h(\theta) g(\theta) d\theta}{\int p(y | \theta) p(\theta) h(\theta) g(\theta) d\theta}, \quad (11)$$

where $g(\theta)$ is the so-called proposal distribution and $h(\theta)$ the so-called bridge function. Multiplying both sides of Eq. (11) by the marginal likelihood $p(y)$ results in:

$$\begin{aligned} p(y) &= \frac{\int p(y | \theta) p(\theta) h(\theta) g(\theta) d\theta}{\int \frac{p(y | \theta) p(\theta)}{p(y)} h(\theta) g(\theta) d\theta} \\ &= \frac{\int p(y | \theta) p(\theta) h(\theta) \overbrace{g(\theta)}^{\text{proposal distribution}} d\theta}{\int h(\theta) g(\theta) \underbrace{p(\theta | y)}_{\text{posterior distribution}} d\theta} \\ &= \frac{\mathbb{E}_{g(\theta)} [p(y | \theta) p(\theta) h(\theta)]}{\mathbb{E}_{\text{post}} [h(\theta) g(\theta)]}. \end{aligned}$$

The marginal likelihood can now be approximated using:

$$\hat{p}(y) = \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} p(y | \tilde{\theta}_i) p(\tilde{\theta}_i) h(\tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} h(\theta_j^*) g(\theta_j^*)}, \quad \begin{aligned} &\tilde{\theta}_i \sim g(\theta) \\ &\theta_j^* \sim p(\theta | y) \end{aligned} \quad (12)$$

samples from the proposal distribution samples from the posterior distribution

Eq. (12) illustrates that we need samples from both the proposal distribution and the posterior distribution to obtain the bridge sampling estimate for the marginal likelihood. However, before we can apply Eq. (12) to our running example, we have to discuss how we can obtain a suitable proposal distribution and bridge function. Conceptually, the proposal distribution is similar to an importance density, should resemble the posterior distribution, and should have sufficient overlap with the posterior distribution. According to Overstall and Forster (2010), a convenient proposal distribution

is often a normal distribution with its first two moments chosen to match those of the posterior distribution. In our experience, this choice for the proposal distribution works well for a wide range of scenarios. However, this proposal distribution might produce unstable estimates in case of high-dimensional posterior distributions that clearly do not follow a multivariate normal distribution. In such a situation, it might be advisable to consider more sophisticated versions of bridge sampling (e.g., Frühwirth-Schnatter, 2004; Meng & Schilling, 2002; Wang & Meng, 2016).

1.5.1. Choosing the optimal bridge function

In this tutorial we use the bridge function defined as (Meng & Wong, 1996):

$$h(\theta) = C \cdot \frac{1}{s_1 p(y | \theta) p(\theta) + s_2 p(y) g(\theta)}, \quad (13)$$

where $s_1 = \frac{N_1}{N_2 + N_1}$, $s_2 = \frac{N_2}{N_2 + N_1}$, and C a constant; its particular value is not required because $h(\theta)$ is part of both the numerator and the denominator of Eq. (12), and therefore the constant C cancels. This particular bridge function is referred to as the “optimal bridge function” because Meng and Wong (1996, p. 837) proved that it minimizes the relative mean-squared error (Eq. (16)).

Eq. (13) shows that the optimal bridge function depends on the marginal likelihood $p(y)$ which is the very entity we want to approximate. We can resolve this issue by applying an iterative scheme that updates an initial guess of the marginal likelihood until the estimate of the marginal likelihood has converged according to a predefined tolerance level. To do so, we insert the expression for the optimal bridge function (Eq. (13)) in Eq. (12) (Meng & Wong, 1996). The formula to approximate the marginal likelihood on iteration $t + 1$ is then specified as follows:

$$\hat{p}(y)^{(t+1)} = \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{p(y | \tilde{\theta}_i) p(\tilde{\theta}_i)}{s_1 p(y | \tilde{\theta}_i) p(\tilde{\theta}_i) + s_2 \hat{p}(y)^{(t)} g(\tilde{\theta}_i)}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{g(\theta_j^*)}{s_1 p(y | \theta_j^*) p(\theta_j^*) + s_2 \hat{p}(y)^{(t)} g(\theta_j^*)}}, \quad (14)$$

samples from the proposal distribution samples from the posterior distribution

where $\hat{p}(y)^{(t)}$ denotes the estimate of the marginal likelihood on iteration t of the iterative scheme. Note that Eq. (14) illustrates why bridge sampling is robust to the tail behavior of the proposal distribution relative to the posterior distribution; the difference to the importance sampling and generalized harmonic mean estimator is that, in the case of the bridge sampling estimator, samples from the tail region cannot inflate individual summation terms and thus dominate the estimate. To illustrate this, we consider what happens to the bridge sampling estimator, the importance sampling estimator, and the generalized harmonic mean estimator in case (1) the proposal/importance distribution has fatter tails than the posterior distribution, and (2) the proposal/importance distribution has thinner tails than the posterior distribution (see also Frühwirth-Schnatter, 2004). Specifically, we look at a single term in the respective sums and consider the limit of that term as we move further and further out in the tails. This is insightful since a single term can have a lasting effect on the estimator (e.g., in case a single term in a sum is very large or even infinite).

In case (1) (i.e., the proposal/importance distribution has fatter tails than the posterior), the ratio in the importance sampling estimator (i.e., Eq. (8)) goes to zero as we move further out in the tails. Since samples in the tails may only be obtained occasionally and a zero term in the sum does not inflate the estimate this is not a reason for concern. In contrast, when we consider the ratio in the generalized harmonic mean estimator (i.e., Eq. (9)), we see that the ratio goes to infinity as we move further out in the tails. Even if this occurs only very rarely, this is an issue since the resulting value

will dominate the estimate. Consequently, the resulting estimator may have a large variance since samples from the tail regions may be obtained only occasionally across repeated applications. For the bridge sampling estimator (i.e., Eq. (14)), we need to consider the ratio in the numerator and denominator. The ratio in the numerator will go to zero and the ratio in the denominator will go to $\frac{1}{s_2 \hat{p}(y)^{(t)}}$. Hence, both of these ratios are bounded and will not inflate the two sums, hence also not the resulting estimate.

In case (2) (i.e., the proposal/importance distribution has thinner tails than the posterior), the ratio in the importance sampling estimator (i.e., Eq. (8)) goes to infinity as we move further out in the tails, inflating the estimate. In contrast, when we consider the ratio in the generalized harmonic mean estimator (i.e., Eq. (9)), we see that the ratio goes to zero. As explained above, this is not a reason for concern. These considerations explain why in importance sampling, the importance distribution should have fatter tails than the posterior whereas for the generalized harmonic mean estimator, it should have thinner tails. For the bridge sampling estimator (i.e., Eq. (14)), the ratio in the numerator will go to $1/s_1$ and the ratio in the denominator will go to zero. Again, both of these ratios are bounded making the bridge sampling estimator more robust to the tail behavior than the other two estimators. This of course assumes that not all terms in the denominator (for case (2)) and the numerator (for case (1)) will be zero, that is, the proposal and the posterior distribution have sufficient overlap. In the extreme scenario of no overlap the bridge sampling estimate is not defined because both sums of Eq. (14) would be zero.

Extending the numerator of the right side of Eq. (14) with $\frac{1/g(\tilde{\theta}_i)}{1/g(\tilde{\theta}_i^*)}$, and the denominator with $\frac{1/g(\theta_j^*)}{1/g(\theta_j^*)}$, and subsequently defining $l_{1,j} := \frac{p(y|\theta_j^*)p(\theta_j^*)}{g(\theta_j^*)}$ and $l_{2,i} := \frac{p(y|\tilde{\theta}_i)p(\tilde{\theta}_i)}{g(\tilde{\theta}_i)}$, we obtain the formula for the iterative scheme of the bridge sampling estimator $\hat{p}_4(y)^{(t+1)}$ at iteration $t + 1$ (Meng & Wong, 1996, p. 837).

$$\begin{aligned} \hat{p}_4(y)^{(t+1)} &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{p(y|\tilde{\theta}_i)p(\tilde{\theta}_i)}{s_1 p(y|\tilde{\theta}_i)p(\tilde{\theta}_i) + s_2 \hat{p}_4(y)^{(t)} g(\tilde{\theta}_i)} \frac{1/g(\tilde{\theta}_i)}{1/g(\tilde{\theta}_i^*)}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{g(\theta_j^*)}{s_1 p(y|\theta_j^*)p(\theta_j^*) + s_2 \hat{p}_4(y)^{(t)} g(\theta_j^*)} \frac{1/g(\theta_j^*)}{1/g(\theta_j^*)}} \\ &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{l_{2,i}}{s_1 l_{2,i} + s_2 \hat{p}_4(y)^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 l_{1,j} + s_2 \hat{p}_4(y)^{(t)}}}, \quad \tilde{\theta}_i \sim \underbrace{g(\theta)}_{\text{samples from the proposal distribution}}, \quad (15) \\ &\quad \underbrace{\theta_j^* \sim p(\theta | y)}_{\text{samples from the posterior distribution}}. \end{aligned}$$

Eq. (15) suggests that, in order to obtain the bridge sampling estimate of the marginal likelihood, a number of requirements need to be fulfilled. First, we need N_2 samples from the proposal distribution $g(\theta)$ and N_1 samples from the posterior distribution $p(\theta|y)$. Second, for all N_2 samples from the proposal distribution, we have to evaluate $l_{2,i}$. This involves obtaining the value of the unnormalized posterior (i.e., the product of the likelihood times the prior) and of the proposal distribution for all samples. Third, we evaluate $l_{1,j}$ for all N_1 samples from the posterior distribution. This is analogous to evaluating $l_{2,i}$. Fourth, we have to determine the constants s_1 and s_2 that only depend on N_1 and N_2 . Fifth, we need an initial guess of the marginal likelihood $\hat{p}_4(y)$. Since some of these five requirements can be obtained easier than others, we will point out possible challenges.

A first challenge is that using a suitable proposal distribution may involve transforming the posterior samples. Consequently, we have to determine how the transformation affects the definition of the bridge sampling estimator for the marginal likelihood (Eq. (15)).

A second challenge is how to use the N_1 samples from the posterior distribution. One option is to use all N_1 samples for both fitting the proposal distribution and for computing the bridge sampling estimate. However, Overstall and Forster (2010) showed that such a procedure may result in an underestimation of the marginal likelihood. To obtain more reliable estimates they propose to divide the posterior samples in two parts; the first part is used to obtain the best-fitting proposal distribution, and the second part is used to compute the bridge sampling estimate. Throughout this tutorial, we use two equally large parts. In the remainder we therefore state that we draw $2N_1$ samples from the posterior distribution. The first N_1 of the total of $2N_1$ samples are used for fitting the proposal distribution and the remaining N_1 samples are used in the iterative scheme (i.e., Eq. (15)).⁹

To summarize, the discussion of the requirements and challenges encountered in bridge sampling illustrated that the bridge sampling estimator imposes less strict requirements on the proposal distribution than the importance sampling and generalized harmonic mean estimator and allows for an almost automatic application due to the default choice of the bridge function.¹⁰

1.5.2. Running example

To obtain the bridge sampling estimate of the marginal likelihood in the beta-binomial example, we follow the eight steps illustrated in Fig. 5:

1. We draw $2N_1 = 24$ samples from the Beta(3, 9) posterior distribution for θ .

We obtain the following sample of 24 values:

$$\{\theta_1^*, \theta_2^*, \dots, \theta_{24}^*\} = \{0.22, 0.16, 0.09, 0.35, 0.06, 0.27, 0.26, 0.41, 0.20, 0.43, 0.21, 0.12, 0.15, 0.21, 0.24, 0.18, 0.12, 0.22, 0.15, 0.22, 0.23, 0.26, 0.29, 0.28\}.$$

Note that the first 12 samples equal the ones used in the last section, whereas the last 12 samples were obtained from drawing again 12 values from the Beta(3, 9) posterior distribution for θ .

2. We choose a proposal distribution.

Here we opt for an approach that can be easily generalized to models with multiple parameters and select a normal distribution as the proposal distribution $g(\theta)$.¹¹

3. We transform the first batch of N_1 posterior samples.

Since we use a normal proposal distribution, we have to transform the posterior samples from the rate scale to the real line so that the range of the posterior distribution matches the range of the proposal distribution. This can be achieved by probit-transforming the posterior samples, that is, $\xi_j^* = \Phi^{-1}(\theta_j^*)$ with $j \in \{1, 2, \dots, 12\}$. We obtain:

$$\{\xi_1^*, \xi_2^*, \dots, \xi_{12}^*\} = \{-0.77, -0.99, -1.34, -0.39, -1.55, -0.61, -0.64, -0.23, -0.84, -0.18, -0.81, -1.17\}.$$

4. We fit the proposal distribution to the first batch of N_1 probit-transformed posterior samples.

⁹ In case the posterior samples are obtained via MCMC sampling using multiple chains, we use the first half of the iterations per chain for fitting the proposal distribution and the second half of the iterations per chain for the iterative scheme.

¹⁰ For an explanation of where the name “bridge” comes from see <https://osf.io/9jzm3/>.

¹¹ There exist several candidates for the proposal distribution. Alternative proposal distributions are, for example, the importance density that we used for the importance sampling estimator or for the generalized harmonic mean estimator, or the analytically derived Beta(3, 9) posterior distribution.

Fig. 5. Schematic illustration of the steps involved in obtaining the bridge sampling estimate of the marginal likelihood. Available at <https://tinyurl.com/y7b2kze7> under CC license <https://creativecommons.org/licenses/by/2.0/>.

We use the method of moment estimates $\hat{\mu} = -0.793$ and $\hat{\sigma} = 0.423$ from the first batch of N_1 probit-transformed posterior samples to obtain our proposal distribution $g(\xi; \mu = -0.793, \sigma = 0.423) = \frac{1}{0.423} \phi\left(\frac{\xi + 0.793}{0.423}\right)$.

5. We draw N_2 samples from the proposal distribution. We obtain:

$$\{\tilde{\xi}_1, \tilde{\xi}_2, \dots, \tilde{\xi}_{12}\} = \{-1.11, -0.63, -1.48, -0.59, -0.48, \\ -0.69, -0.74, -0.51, -0.82, \\ -1.54, -0.76, -0.96\}.$$

6. We calculate $l_{2,i}$ for all N_2 samples from the proposal distribution.

This step involves assessing the value of the unnormalized posterior and the proposal distribution for all N_2 samples from the proposal distribution. As in the running example for the generalized harmonic mean estimator, we obtain the unnormalized posterior as: $p(k=2 | n=10, \Phi(\tilde{\xi}_i)) \phi(\tilde{\xi}_i)$, where $\phi(\tilde{\xi}_i)$ comes from using the change-of-variable method (see running example for the generalized harmonic mean estimator and the Appendix for details). Thus, as in the case of the generalized harmonic mean estimator, the uniform prior on θ translates to a standard normal prior on ξ . The values of the proposal distribution can easily be obtained (for example using the R software).

7. We transform the second batch of N_1 posterior samples. As in step 2, we use the probit transformation and obtain:

$$\{\xi_{13}^*, \xi_{14}^*, \dots, \xi_{24}^*\} = \{-1.04, -0.81, -0.71, -0.92, -1.17, \\ -0.77, -1.04, -0.77, -0.74, \\ -0.64, -0.55, -0.58\}.$$

8. We calculate $l_{1,j}$ for the second batch of N_1 probit-transformed samples from the posterior distribution. This is analogous to step 6.

9. We run the iterative scheme (Eq. (15)) until our predefined tolerance criterion is reached.

As tolerance criterion we choose $|\hat{p}_4(k=2 | n=10)^{(t+1)} - \hat{p}_4(k=2 | n=10)^{(t)}| / \hat{p}_4(k=2 | n=10)^{(t+1)} \leq 10^{-10}$.

This requires an initial guess for the marginal likelihood $\hat{p}_4(k=2 | n=10)^{(0)}$ which we set to 0.¹²

The simplicity of the beta-binomial model allows us to calculate the bridge sampling estimate by hand. To determine $\hat{p}_4(y)^{(t+1)}$ according to Eq. (15), we need to calculate the constants s_1 and s_2 . Since $N_1 = N_2 = 12$, we obtain: $s_1 = s_2 = N_2 / (N_2 + N_1) = 0.5$. In addition, we need to calculate $l_{2,i}$ ($i \in \{1, 2, \dots, 12\}$) for all samples from the proposal distribution, and $l_{1,j}$ ($j \in \{1, 2, \dots, 12\}$) for the second batch of the probit-transformed samples from the posterior distribution. Here we show how to calculate $l_{2,1}$ and $l_{1,1}$ using the first sample from the proposal distribution and the first sample of the second batch of the posterior samples, respectively:

$$l_{2,1} = \frac{p(k | n, \Phi(\tilde{\xi}_1)) \phi(\tilde{\xi}_1)}{g(\tilde{\xi}_1)} = \left(\frac{\binom{10}{2} 0.13^2 (1 - 0.13)^8 \cdot 0.22}{\frac{1}{0.423} \phi\left(\frac{-1.11 + 0.793}{0.423}\right)} \right) = 0.077,$$

$$l_{1,1} = \frac{p(k | n, \Phi(\xi_{13}^*)) \phi(\xi_{13}^*)}{g(\xi_{13}^*)} = \left(\frac{\binom{10}{2} 0.15^2 (1 - 0.15)^8 \cdot 0.23}{\frac{1}{0.423} \phi\left(\frac{-1.04 + 0.793}{0.423}\right)} \right) = 0.080.$$

For $\hat{p}_4(k=2 | n=10)^{(t+1)}$, we then get:

$$\hat{p}_4(k=2 | n=10)^{(t+1)} = \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{l_{2,i}}{s_1 l_{2,i} + s_2 \hat{p}_4(k=2 | n=10)^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 l_{1,j} + s_2 \hat{p}_4(k=2 | n=10)^{(t)}}} = \frac{\frac{1}{12} \left(\frac{0.077}{0.5 \cdot 0.077 + 0.5 \cdot \hat{p}_4(k=2 | n=10)^{(t)}} + \dots + \frac{0.084}{0.5 \cdot 0.084 + 0.5 \cdot \hat{p}_4(k=2 | n=10)^{(t)}} \right)}{\frac{1}{12} \left(\frac{1}{0.5 \cdot 0.080 + 0.5 \cdot \hat{p}_4(k=2 | n=10)^{(t)}} + \dots + \frac{1}{0.5 \cdot 0.103 + 0.5 \cdot \hat{p}_4(k=2 | n=10)^{(t)}} \right)}.$$

Using $\hat{p}(y)^{(0)} = 0$, we obtain as updated estimate of the marginal likelihood $\hat{p}_4(k=2 | n=10)^{(1)} = 0.0908$. This iterative procedure has to be repeated until our predefined tolerance criterion is reached. For our running example, this criterion is reached after

¹² A better initial guess can be obtained from, for example, the importance sampling estimator or the generalized harmonic mean estimator explained in the previous sections. In our experience, however, usually the exact choice of the initial value does not seem to influence the convergence of the bridge sampler much.

five iterations. We now obtain the bridge sampling estimate of the marginal likelihood as $\hat{p}_4(k = 2 \mid n = 10)^{(5)} = 0.0902$.

1.6. Interim summary

So far we used the beta-binomial model to illustrate the computation of four different estimators of the marginal likelihood. These four estimators were discussed in order of increasing sophistication, such that the first three estimators provided the proper context for understanding the fourth, most general estimator—the bridge sampler. This estimator is the focus in the remainder of this tutorial. The goal of the next sections is to demonstrate that bridge sampling is particularly suitable to estimate the marginal likelihood of popular models in mathematical psychology. Importantly, bridge sampling may be used to obtain accurate estimates of the marginal likelihood of hierarchical models (for a detailed comparison of bridge sampling versus its special cases see [Frühwirth-Schnatter, 2004](#); [Sinharay & Stern, 2005](#)).

1.7. Assessing the accuracy of the bridge sampling estimate

In this section we show how to quantify the accuracy of the bridge sampling estimate. A straightforward approach would be to apply the bridge sampling procedure multiple times and investigate the variability of the marginal likelihood estimate. In practice, however, this solution is often impractical due to the substantial computational burden of obtaining the posterior samples and evaluating the relevant quantities in the bridge sampling procedure.

[Frühwirth-Schnatter \(2004\)](#) proposed an alternative approach that approximates the estimator's expected relative mean-squared error:

$$RE^2 = \frac{\mathbb{E} \left[(\hat{p}_4(y) - p(y))^2 \right]}{p(y)^2}. \quad (16)$$

The derivation of this approximate relative mean-squared error by [Frühwirth-Schnatter](#) takes into account that the samples from the proposal distribution $g(\theta)$ are independent, whereas the MCMC samples from the posterior distribution $p(\theta|y)$ may be autocorrelated. The approximate relative mean-squared error is given by:

$$\widehat{RE}^2 = \frac{1}{N_2} \frac{V_{g(\theta)}(f_1(\theta))}{\mathbb{E}_{g(\theta)}^2(f_1(\theta))} + \frac{\rho_{f_2}(0)}{N_1} \frac{V_{\text{post}}(f_2(\theta))}{\mathbb{E}_{\text{post}}^2(f_2(\theta))}, \quad (17)$$

where $f_1(\theta) = \frac{p(\theta|y)}{s_1 p(\theta|y) + s_2 g(\theta)}$, $f_2(\theta) = \frac{g(\theta)}{s_1 p(\theta|y) + s_2 g(\theta)}$, $V_{g(\theta)}(f_1(\theta)) = \int (f_1(\theta) - \mathbb{E}[f_1(\theta)])^2 g(\theta) d\theta$ denotes the variance of $f_1(\theta)$ with respect to the proposal distribution $g(\theta)$ (the variance $V_{\text{post}}(f_2(\theta))$ is defined analogously), and $\rho_{f_2}(0)$ corresponds to the normalized spectral density of the autocorrelated process $f_2(\theta)$ at the frequency 0.

In practice, we approximate the unknown variances and expected values by the corresponding sample variances and means. Hence, for evaluating the variance and expected value with respect to $g(\theta)$, we use the N_2 samples for θ_i from the proposal distribution. To evaluate the variance and expected value with respect to the posterior distribution, we use the second batch of N_1 samples θ_j^* from the posterior distribution which we also use in the iterative scheme for computing the marginal likelihood. Because the posterior samples are obtained via an MCMC procedure and are hence autocorrelated, the second term in Eq. (17) is adjusted by the normalized spectral density (for details see [Frühwirth-Schnatter, 2004](#)).¹³ To evaluate the normalized posterior density

which appears in the numerator of $f_1(\theta)$ and the denominator of both $f_1(\theta)$ and $f_2(\theta)$, we use the bridge sampling estimate as normalizing constant.

Note that, under the assumption that the bridge sampling estimator $\hat{p}_4(y)$ is an unbiased estimator of the marginal likelihood $p(y)$, the square root of the expected relative mean-squared error (Eq. (16)) can be interpreted as the coefficient of variation (i.e., the ratio of the standard deviation and the mean; [Brown, 1998](#)). In the remainder of this article, we report the coefficient of variation to quantify the accuracy of the bridge sampling estimate.

2. Case study: bridge sampling for reinforcement learning models

In this section, we illustrate the computation of the marginal likelihood using bridge sampling in the context of a published data set ([Bussemeyer & Stout, 2002](#)) featuring the Expectancy Valence (EV) model—a popular reinforcement learning (RL) model for the Iowa gambling task (IGT; [Bechara et al., 1994](#)). We first introduce the task and the model, and then use bridge sampling to estimate the marginal likelihood of the EV model implemented in both an individual-level and a hierarchical Bayesian framework. For the individual-level framework, we compare estimates obtained from bridge sampling to importance sampling estimates published in [Steingroever, Wetzels, and Wagenmakers \(2016\)](#). For the hierarchical framework, we compare our results to estimates from the Savage–Dickey density ratio test ([Dickey, 1971](#); [Dickey & Lientz, 1970](#); [Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010](#); [Wetzels, Grasman, & Wagenmakers, 2010](#)).

2.1. The Iowa gambling task

In this section we describe the IGT (see also [Steingroever, Pachur, Šmíra, & Lee, submitted for publication](#); [Steingroever, Wetzels, Horstmann, Neumann, & Wagenmakers, 2013](#); [Steingroever, Wetzels, & Wagenmakers, 2013a](#); [Steingroever, Wetzels, & Wagenmakers, 2013b](#); [Steingroever, Wetzels, & Wagenmakers, 2014](#); [Steingroever et al., 2016](#)). Originally, [Bechara et al. \(1994\)](#) developed the IGT to distinguish decision-making strategies of patients with lesions to the ventromedial prefrontal cortex from the ones of healthy controls (see also [Bechara, Damasio, Tranel, & Anderson, 1998](#); [Bechara, Damasio, Damasio, & Lee, 1999](#); [Bechara, Tranel, & Damasio, 2000](#)). During the last decades, the scope of application of the IGT has increased tremendously covering clinical populations with, for example, pathological gambling ([Cavedini, Riboldi, Keller, D'Annuncci, & Bellodi, 2002](#)), obsessive-compulsive disorder ([Cavedini, Riboldi, D'Annuncci, Belotti, Cisima, & Bellodi, 2002](#)), psychopathic tendencies ([Blair, Colledge, & Mitchell, 2001](#)), and schizophrenia ([Bark, Dieckmann, Bogerts, & Northoff, 2005](#); [Martino, Bucay, Butman, & Allegrì, 2007](#)).

The IGT is a card game that requires participants to choose, over several rounds, cards from four different decks in order to maximize their long-term net outcome ([Bechara et al., 1994](#); [Bechara, Damasio, Tranel, & Damasio, 1997](#)). The four decks differ in their payoffs, and two of them result in negative long-term outcomes (i.e., the bad decks), whereas the remaining two decks result in positive long-term outcomes (i.e., the good decks). After each choice, participants receive feedback on the rewards and losses (if any) associated with that card, as well as their running tally of net outcomes over all trials so far. Unbeknownst to the participants, the task (typically) contains 100 trials.

A crucial aspect of the IGT is whether and to what extent participants eventually learn to prefer the good decks because only choosing from the good decks maximizes their long-term net outcome. The good decks are typically labeled as decks C and D, whereas the bad decks are labeled as decks A and B. [Table 1](#)

¹³ We estimate the spectral density at frequency zero by fitting an autoregressive model using the `spectrum0.ar()` function as implemented in the `coda` R package ([Plummer, Best, Cowles, & Vines, 2006](#)).

Table 1

Summary of the Payoff Scheme of the Traditional IGT as Developed by Bechara et al. (1994).

	Deck A Bad deck with infrequent losses	Deck B Bad deck with frequent losses	Deck C Good deck with infrequent losses	Deck D Good deck with frequent losses
Reward/trial	100	100	50	50
Number of losses/10 cards	5	1	5	1
Loss/10 cards	–1250	–1250	–250	–250
Net outcome/10 cards	–250	–250	250	250

presents a summary of the traditional payoff scheme as developed by Bechara et al. (1994). This table illustrates that decks A and B yield high constant rewards, but even higher unpredictable losses: hence, the long-term net outcome is negative. Decks C and D, on the other hand, yield low constant rewards, but even lower unpredictable losses: hence, the long-term net outcome is positive. In addition to the different payoff magnitudes, the decks also differ in the frequency of losses: decks A and C yield frequent losses, while decks B and D yield infrequent losses.

2.2. The expectancy valence model

In this section, we describe the EV model (see also Steingroever et al., 2013a; Steingroever et al., 2014 2016 submitted for publication). Originally proposed by Busemeyer and Stout (2002), the EV model is arguably the most popular model for the IGT (for references see Steingroever et al., 2013a, and for alternative IGT models see Ahn, Busemeyer, Wagenmakers, & Stout, 2008; Dai, Kerestes, Upton, Busemeyer, & Stout, 2015; Steingroever et al., 2014; Worthy & Maddox, 2014; Worthy, Pang, & Byrne, 2013). The model formalizes participants' performance on the IGT through the interaction of three model parameters that represent distinct psychological processes. The first model assumption is that after choosing a card from deck k , $k \in \{1, 2, 3, 4\}$, on trial t , participants compute a weighted mean of the experienced reward $W(t)$ and loss $L(t)$ to obtain the utility of deck k on trial t , $v_k(t)$:

$$v_k(t) = (1 - w)W(t) + wL(t).$$

The weight that participants assign to losses relative to rewards is the attention weight parameter w . A small value of w , that is, $w < .5$, is characteristic for decision makers who put more weight on the immediate rewards and can thus be described as reward-seeking, whereas a large value of w , that is, $w > .5$, is characteristic for decision makers who put more weight on the immediate losses and can thus be described as loss-averse (Ahn et al., 2008; Busemeyer & Stout, 2002).

The EV model further assumes that decision makers use the utility of deck k on trial t , $v_k(t)$, to update only the expected utility of deck k , $Ev_k(t)$; the expected utilities of the unchosen decks are left unchanged. This updating process is described by the Delta learning rule, also known as the Rescorla–Wagner rule (Rescorla & Wagner, 1972):

$$Ev_k(t) = Ev_k(t - 1) + a(v_k(t) - Ev_k(t - 1)).$$

If the experienced utility $v_k(t)$ is higher than expected, the expected utility of deck k is adjusted upward. If the experienced utility $v_k(t)$ is lower than expected, the expected utility of deck k is adjusted downward. This updating process is influenced by the second model parameter—the updating parameter a . This parameter quantifies the memory for rewards and losses. A value of a close to zero indicates slow forgetting and weak recency effects, whereas a value of a close to one indicates rapid forgetting and strong recency effects. For all models, we initialized the expectancies of all decks to zero, $Ev_k(0) = 0$ ($k \in \{1, 2, 3, 4\}$). This setting reflects neutral prior knowledge about the payoffs of the decks.

In the next step, the model assumes that the expected utilities of each deck guide participants' choices on the next trial $t + 1$. This

assumption is formalized by the softmax choice rule, also known as the ratio-of-strength choice rule (Luce, 1959):

$$Pr[S_k(t + 1)] = \frac{e^{\theta(t) \cdot Ev_k(t)}}{\sum_{j=1}^4 e^{\theta(t) \cdot Ev_j(t)}}.$$

The EV model uses this rule to compute the probability of choosing each deck on each trial. This rule contains a sensitivity parameter θ that indexes the extent to which trial-by-trial choices match the expected deck utilities. Values of θ close to zero indicate random choice behavior (i.e., strong exploration), whereas large values of θ indicate choice behavior that is strongly determined by the expected utilities (i.e., strong exploitation).

The EV model uses a trial-dependent sensitivity parameter $\theta(t)$, which also depends on the final model parameter, response consistency $c' \in [-5, 5]$:

$$\theta(t) = (t/10)^{c'}.$$

If c' is positive, successive choices become less random and more determined by the expected deck utilities; if c' is negative, successive choices become more random and less determined by the expected deck utilities, a pattern that is clearly non-optimal. We restricted the consistency parameter of the EV model to the range $[-2, 2]$ instead of the proposed range $[-5, 5]$ (Busemeyer & Stout, 2002). This modification improved the estimation of the EV model and prevented the choice rule from producing numbers that exceed machine precision (see also Steingroever et al., 2014).

In sum, the EV model has three parameters: (1) the attention weight parameter $w \in [0, 1]$, which quantifies the weight of losses over rewards; (2) the updating parameter $a \in [0, 1]$, which determines the memory for past expectancies; and (3) the response consistency parameter $c' \in [-2, 2]$, which determines the balance between exploitation and exploration.

2.3. Data

We applied bridge sampling to a data set published by Busemeyer and Stout (2002). The data set consists of 30 healthy participants each contributing $T = 100$ IGT card selections (see Busemeyer and Stout for more details on the data sets).¹⁴

2.4. Application of bridge sampling to an individual-level implementation of the EV model

In this section we describe how we use bridge sampling to estimate the marginal likelihood of an individual-level implementation of the EV model. This implementation estimates model parameters for each participant separately. Accordingly, we also obtain a marginal likelihood of the EV model for every participant.

¹⁴ Note that we excluded three participants due to incomplete choice data.

2.4.1. Schematic execution of the bridge sampler

To obtain the bridge sampling estimate of the marginal likelihood for each participant, we follow the steps outlined in Fig. 5.

For each participant $s, s \in \{1, 2, \dots, 30\}$, we proceed as follows:

1. For each parameter, we draw $2N_1$ samples from the posterior distribution.

Since Steingroever et al. (2016) already fit an individual-level implementation of the EV model separately to the data of each participant in Busemeyer and Stout (2002), we reuse their posterior samples (see Steingroever et al., 2016, for details on the prior distributions and model implementation). Note that they parameterized the model not in terms of $c' \in [-2, 2]$, but in terms of $c = (c' + 2)/4$, $c \in [0, 1]$, and in the remainder of this article, we also use this reparameterization.

For each participant, we choose $2N_1$ to match the number of samples obtained from Steingroever et al. (2016) which was at least 5000; however, whenever this number of samples was insufficient to ensure convergence of the Hamiltonian Monte Carlo (HMC) chains, Steingroever et al. (2016) repeated the fitting routine with 5000 additional samples. Steingroever et al. (2016) confirmed convergence of the HMC chains by reporting that all $\hat{R}$ statistics were below 1.05.

2. We choose a proposal distribution.

We generalize our approach from the running example and use a multivariate normal distribution as a proposal distribution.

3. We transform the first batch of N_1 posterior samples.

Since we use a multivariate normal distribution as a proposal distribution, we have to transform all posterior samples to the real line using the probit transformation, that is, $\omega_{s,j}^* = \Phi^{-1}(w_{s,j}^*)$, $\alpha_{s,j}^* = \Phi^{-1}(a_{s,j}^*)$, $\gamma_{s,j}^* = \Phi^{-1}(c_{s,j}^*)$, $j = \{1, 2, \dots, N_1\}$.

4. We fit the proposal distribution to the first batch of N_1 probit-transformed posterior samples.

We use method of moment estimates for the mean vector and the covariance matrix obtained from the first batch of N_1 probit-transformed posterior samples to specify our multivariate normal proposal distribution.

5. We draw N_2 samples from the proposal distribution.

We use the R software to randomly draw N_2 samples from the proposal distribution obtained in step 4. We obtain $(\tilde{\omega}_{s,i}, \tilde{\alpha}_{s,i}, \tilde{\gamma}_{s,i})$ with $i \in \{1, 2, \dots, N_2\}$.

6. We calculate $l_{2,i}$ for all N_2 samples from the proposal distribution.

This step involves assessing the value of the unnormalized posterior and the proposal distribution for all N_2 samples from the proposal distribution. Before we can assess the value of the unnormalized posterior (i.e., the product of the likelihood and the prior), we have to derive how our transformation in step 3 affects the unnormalized posterior. First, we derive how our transformation affects the likelihood. To evaluate the likelihood, we need to transform the probit-transformed samples back to the original parameter scales. That is, we evaluate the likelihood for $(\Phi(\tilde{\omega}_{s,i}), \Phi(\tilde{\alpha}_{s,i}), \Phi(\tilde{\gamma}_{s,i}))$. Before formalizing the likelihood of the observed choices of participant s , we define the following variables: We define $Ch_s(t)$ as a vector containing the sequence of choices made by participant s up to and including trial t , and $X_s(t)$ as a vector containing the corresponding sequence of net outcomes. We now obtain the following expression for the likelihood of the observed choices of participant s :

$$p(Ch_s(T) | \Phi(\tilde{\omega}_{s,i}), \Phi(\tilde{\alpha}_{s,i}), \Phi(\tilde{\gamma}_{s,i}), X_s(T-1)) = \prod_{t=0}^{T-1} \prod_{k=1}^4 \Pr[S_k(t+1)] \cdot \delta_k(t+1). \quad (18)$$

Here T is the total number of trials, $\Pr[S_k(t+1)]$ is the probability of choosing deck k on trial $t+1$, and $\delta_k(t+1)$ is a dummy variable which is 1 if deck k is chosen on trial $t+1$ and 0 otherwise.

Second, we have to derive how our transformation affects the priors on each EV model parameter to yield priors on the probit-transformed model parameters. Since Steingroever et al. (2016) used independent uniform priors on $[0, 1]$ we obtain standard normal priors on the probit-transformed model parameters (see beta-binomial example and Appendix D for an explanation).

7. We transform the second batch of N_1 posterior samples.

This is analogous to step 2.

8. We calculate $l_{1,j}$ for the second batch of N_1 probit-transformed samples from the posterior distribution.

This is analogous to step 6.

9. We run the iterative scheme (Eq. (15)) until our predefined tolerance criterion is reached.

We use a tolerance criterion and initialization analogous to the running example. Once convergence is reached, we receive an estimate of the marginal likelihood for each participant, and derive the coefficient of variation for each participant using Eq. (17). The largest coefficient of variation is 2.07% suggesting that the bridge sampler has low variance.¹⁵

2.4.2. Assessing the accuracy of our implementation

To assess the accuracy of our implementation, we compared the marginal likelihood estimates obtained with our bridge sampler to the estimates obtained with importance sampling (Steingroever et al., 2016). Fig. 6 shows the log marginal likelihoods for the 30 participants of Busemeyer and Stout (2002) obtained with bridge sampling (x-axis) and importance sampling reported by Steingroever et al. (2016; y-axis). The two sets of estimates correspond almost perfectly. These results indicate a successful implementation of the bridge sampler. Thus, this section emphasizes that both the importance sampler and bridge sampler can be used to estimate the marginal likelihood for the data of individual participants. However, when we want to estimate the marginal likelihood of a Bayesian hierarchical model, it may be difficult to find a suitable importance density. The bridge sampler, on the other hand, can be applied more easily and more efficiently.

2.5. Application of bridge sampling to a hierarchical implementation of the EV model

In this section we illustrate how bridge sampling can be used to estimate the marginal likelihood of a hierarchical EV model. This hierarchical implementation assumes that the parameters w , a , and c from each participant are drawn from three separate group-level distributions. This model specification hence incorporates both the differences and the similarities between participants. We illustrate this application using again the Busemeyer and Stout (2002) data set, and assume that these participants constitute one group.

2.5.1. Schematic execution of the bridge sampler

To compute the marginal likelihood, we again follow the steps outlined in Fig. 5, with a few minor modifications.

1. For each parameter, that is, all individual-level and group-level parameters, we draw $2N_1 = 60,000$ samples from the posterior distribution.

¹⁵ Note that this measure relates to the marginal likelihoods, not to the log marginal likelihoods.

Fig. 6. Comparison of the log marginal likelihoods obtained with bridge sampling (x-axis) and importance sampling reported by Steingrover et al. (2016) (y-axis). The main diagonal indicates perfect correspondence between the two methods. Available at <https://tinyurl.com/yac3o8qs> under CC license <https://creativecommons.org/licenses/by/2.0/>.

To obtain the posterior samples, we fit a hierarchical Bayesian implementation of the EV model to the Bussemeyer and Stout (2002) data set using the software JAGS (Plummer, 2003).¹⁶ We assume that, for each participant s , $s \in \{1, 2, \dots, 30\}$, each probit-transformed individual-level parameter (i.e., $\omega_s = \Phi^{-1}(w_s)$, $\alpha_s = \Phi^{-1}(a_s)$, $\gamma_s = \Phi^{-1}(c_s)$) is drawn from a group-level normal distribution characterized by a group-level mean and standard deviation parameter. For all group-level mean parameters μ_ω , μ_α , μ_γ we assume a standard normal distribution, and for all group-level standard deviation parameters σ_ω , σ_α , σ_γ a uniform distribution ranging from 0 to 1.5. For a detailed explanation of the hierarchical implementation of the EV model, see Wetzels, Vandekerckhove et al. (2010).

To reach convergence and reduce autocorrelation, we collect two MCMC chains, each with 120,000 samples from the posterior distributions after having excluded the first 30,000 samples as burn-in. Out of these 120,000 samples per chain, we retained every 4th value yielding 30,000 samples per chain. This setting resulted in all $\hat{R}$ statistics below 1.05 suggesting that all chains have successfully converged from their starting values to their stationary distributions.

2. We choose a proposal distribution.

We use a multivariate normal distribution as a proposal distribution.

3. We transform the first batch of N_1 posterior samples.

As before, we ensure that the range of the posterior distribution matches the range of the proposal distribution by using the probit transformation, that is, $\omega_{s,j}^* = \Phi^{-1}(w_{s,j}^*)$, $\alpha_{s,j}^* = \Phi^{-1}(a_{s,j}^*)$, $\gamma_{s,j}^* = \Phi^{-1}(c_{s,j}^*)$, $\tau_{\omega,j}^* = \Phi^{-1}((\sigma_{\omega,j}^*)/1.5)$, $\tau_{\alpha,j}^* = \Phi^{-1}((\sigma_{\alpha,j}^*)/1.5)$, and $\tau_{\gamma,j}^* = \Phi^{-1}((\sigma_{\gamma,j}^*)/1.5)$, $j = \{1, 2, \dots, N_1\}$. The group-level mean parameters do not have to be transformed because they already range across the entire real line.

4. We fit the proposal distribution to the first batch of the N_1 probit-transformed posterior samples.

We use method of moment estimates for the mean vector and the covariance matrix obtained from the first batch of N_1 probit-transformed posterior samples to specify our multivariate normal proposal distribution.

5. We draw N_2 samples from the proposal distribution.

We use the R software to randomly draw N_2 samples from the proposal distribution obtained in step 4. We obtain $(\tilde{\omega}_{s,i}, \tilde{\alpha}_{s,i}, \tilde{\gamma}_{s,i})$ and $(\tilde{\mu}_{\omega,i}, \tilde{\tau}_{\omega,i}, \tilde{\mu}_{\alpha,i}, \tilde{\tau}_{\alpha,i}, \tilde{\mu}_{\gamma,i}, \tilde{\tau}_{\gamma,i})$ with $i \in \{1, 2, \dots, N_2\}$ and $s \in \{1, 2, \dots, 30\}$.

6. We calculate $l_{2,i}$ for all N_2 samples from the proposal distribution.

This step involves assessing the value of the unnormalized posterior and the proposal distribution for all N_2 samples from the proposal distribution. The unnormalized posterior is defined as:

$\left(\prod_{s=1}^{30} p(\mathcal{C}_s(T) \mid \Phi(\tilde{\kappa}_{s,i}), X_s(T-1)) p(\tilde{\kappa}_{s,i} \mid \tilde{\xi}_i) \right) p(\tilde{\xi}_i)$, where $\mathcal{C}_s(T)$ refers to all choices of subject s , $X_s(T-1)$ to the net outcomes that subject s experienced on trials $1, 2, \dots, T-1$, $\tilde{\kappa}_{s,i} = (\tilde{\omega}_{s,i}, \tilde{\alpha}_{s,i}, \tilde{\gamma}_{s,i})$ to the i th sample from the proposal distribution for the individual-level parameters of subject s , and $\tilde{\xi}_i$ to the i th sample from the proposal distribution for all group-level parameters (e.g., $\tilde{\xi}_i = (\tilde{\mu}_{\omega,i}, \tilde{\tau}_{\omega,i}, \tilde{\mu}_{\alpha,i}, \tilde{\tau}_{\alpha,i}, \tilde{\mu}_{\gamma,i}, \tilde{\tau}_{\gamma,i})$).

The likelihood function for a given participant is the same as in the individual case. However, for each participant we now have to add besides the prior on the individual-level parameters also the prior on the group-level parameters. The product of the likelihood and the priors gives the unnormalized posterior density (see Appendix E for more details).

7. We follow steps 7–9, as outlined for the bridge sampler of the individual-level implementation of the EV model.

2.5.2. Assessing the accuracy of our implementation

To investigate the accuracy of our implementation, we compare Bayes factors obtained with bridge sampling to Bayes factors obtained from the Savage–Dickey density ratio test (Dickey, 1971; Dickey & Lientz, 1970; for a tutorial, see Wagenmakers et al., 2010). The Savage–Dickey density ratio is a simple method for computing Bayes factors for nested models. We artificially create three nested models by taking the full EV model $\mathcal{M}_f$ in which all parameters are free to vary, and then restricting one of the three group-level mean parameters, that is, μ_ω , μ_α , or μ_γ , to a predefined value. For these values we choose the intersection point of the prior and posterior distribution of each group-level mean parameter. To obtain these intersection points, we fit the full EV model and then use a nonparametric logspline density estimator (Stone, Hansen, Kooperberg, Truong, et al., 1997). The obtained values are presented in Table 2. Since we compare the full model to each restricted model, we obtain three Bayes factors.

According to the Savage–Dickey density ratio test, the Bayes factor for the full model versus a specific restricted model $\mathcal{M}_r$ can be obtained by dividing the height of the prior density at the predefined parameter value θ_0 by the height of the posterior at the same location:

$$BF_{\mathcal{M}_f, \mathcal{M}_r} = \frac{p(y \mid \mathcal{M}_f)}{p(y \mid \mathcal{M}_r)} = \frac{p(\theta = \theta_0 \mid \mathcal{M}_f)}{p(\theta = \theta_0 \mid y, \mathcal{M}_f)}. \quad (19)$$

Since we choose θ_0 to be the intersection point of the prior and posterior distribution, $BF_{\mathcal{M}_f, \mathcal{M}_r}$ equals 1. This Savage–Dickey Bayes factor of 1 indicates that the marginal likelihood under the full model equals the marginal likelihood under the restricted model. Fig. 7 illustrates the Savage–Dickey Bayes factor comparing the full model to the model assuming μ_α fixed to -0.604 .

The computation of the three bridge sampling Bayes factors, on the other hand, works as follows: First, we follow the steps outlined above to obtain the bridge sampling estimate of the full EV model. Second, we obtain the bridge sampling estimate of the marginal likelihood for the three restricted models. This requires adapting the steps outlined above to each of the three restricted models.

¹⁶ We used a model file that is an adapted version of the model file used by Ahn et al. (2011).

Fig. 7. Prior and posterior distribution of the group-level mean μ_α in the Bussemeyer and Stout (2002) data set. The figure shows the posterior distribution (solid line) and the prior distribution (dotted line). The gray dot indicates the intersection of the prior and the posterior distributions, for which the Savage–Dickey Bayes factor equals 1. Available at <https://tinyurl.com/y7cyxclq> under CC license <https://creativecommons.org/licenses/by/2.0/>.

Lastly, we use the first equality in Eq. (19) to obtain the three Bayes factors.

The Bayes factors derived from bridge sampling are reported in Table 2. It is evident that Bayes factors derived from bridge sampling closely approximate the Savage–Dickey Bayes factors of 1. These results suggest a successful implementation of the bridge sampler. This is also reflected by the close match between the log marginal likelihoods of the four models presented in the third column of Table 2.

Finally, we confirm that the bridge sampler has low variance; the coefficient of variation with respect to the marginal likelihood of the full model and the three restricted models ranges between 9.71 and 16.44%.

3. Discussion

In this tutorial, we explained how bridge sampling can be used to estimate the marginal likelihood of popular models in mathematical psychology. As a running example, we used the beta-binomial model to illustrate step-by-step the bridge sampling estimator. To facilitate the understanding of the bridge sampler, we first discussed three of its special cases—the naive Monte Carlo estimator, the importance sampling estimator, and the generalized harmonic mean estimator. Consequently, we introduced key concepts that became gradually more complicated and sophisticated. In the second part of this tutorial, we showed how bridge sampling can be used to estimate the marginal likelihood of both an individual-level and a hierarchical implementation of the Expectancy Valence (EV; Bussemeyer & Stout, 2002) model—a popular reinforcement-learning model for the Iowa gambling task (IGT; Bechara et al., 1994). The running example and the application of bridge sampling to the EV model demonstrated the positive aspects of the bridge sampling estimator, that is, its accuracy, reliability, practicality, and ease-of-implementation (DiCiccio et al., 1997; Frühwirth-Schnatter, 2004; Meng & Wong, 1996).

The bridge sampling estimator is superior to the naive Monte Carlo estimator, the importance sampling estimator, and the generalized harmonic mean estimator for several reasons. First, Meng and Wong (1996) showed that, among the four estimators discussed in this article, the bridge sampler presented in this article minimizes the mean-squared error because it uses the optimal bridge function. Second, in bridge sampling, choosing a suitable proposal distribution is much easier than choosing a suitable importance density for the importance sampling estimator or the generalized harmonic mean estimator because bridge sampling

is more robust to the tail behavior of the proposal distribution relative to the posterior distribution. This advantage facilitates the application of the bridge sampler to higher dimensional and complex models. This characteristic of the bridge sampler combined with the popularity of higher dimensional and complex models in mathematical psychology suggests that bridge sampling can advance model comparison exercises in many areas of mathematical psychology (e.g., reinforcement-learning models, response time models, multinomial processing tree models, etc.). Third, bridge sampling is relatively straightforward to implement. In particular, our step-by-step procedure can be easily applied to other models with only minor changes of the code (i.e., the unnormalized posterior and potentially the proposal function have to be adapted). In our opinion, this is one of the most attractive features of bridge sampling: It is an accurate yet very generic method. Exploiting this generic characteristic, we have implemented the bridge sampling procedure in the `bridgesampling` R package (Gronau, Singmann, & Wagenmakers, 2017) in order to maximize its accessibility.

Despite the numerous advantages of the bridge sampler, the take-home message of this tutorial is not that the bridge sampler should be used blindly. There exist a large variety of methods to approximate the marginal likelihood that differ in their efficiency.¹⁷

The most appropriate method optimizes the trade-off between accuracy and implementation effort. This trade-off depends on a number of aspects such as the complexity of the model, the number of models under consideration, the statistical experience of the researcher, and the time available. This suggests that the choice of the method should be reconsidered each time a marginal likelihood needs to be obtained. Obviously, when the marginal likelihood can be determined analytically, bridge sampling is not needed at all. If the goal is to compare (at least) two nested models, the Savage–Dickey density ratio test (Dickey, 1971; Dickey & Lientz, 1970) might be a better alternative. Note, however, that this requires an approximation of the marginal posterior density of one or more parameters which can be unstable in case the test value falls in the tail of the distribution. If only an individual-level implementation of a model is used, importance sampling may be easier to implement and may require less computational effort. This presupposes that one can find a proposal distribution with fatter tails than the posterior which may not always be trivial (even in an individual-level case). If the goal is to obtain the marginal likelihood of a large number of relatively simple models, the product space or reversible jump method (RJMCMC) might be more appropriate (Carlin & Chib, 1995; Green, 1995; Lodewyckx et al., 2011). In contrast to bridge sampling, implementations of these methods tend to be problem-specific rather than generic (but see Lunn, Best, & Whittaker, 2009). If a researcher with a limited programming background and/or little time resources wants to conduct a model comparison exercise, rough approximations of the Bayes factor, such as the Bayesian information criterion, might be more suitable (Schwarz, 1978). It should be kept in mind, however, that this approximation assumes a certain prior structure that may not respect the knowledge or information that researchers have at their disposal. On the other hand, a researcher with an extensive background in programming and mathematical statistics might consider using path sampling—a generalization of bridge sampling (Gelman & Meng, 1998).

To conclude, in this tutorial we showed that bridge sampling offers a reliable and easy-to-implement approach to estimating a model's marginal likelihood. Bridge sampling can be profitably applied to a wide range of problems in mathematical psychology involving parameter estimation, model comparison, and Bayesian model averaging.

¹⁷ In general, a large number of approaches for model selection exist which are based on MCMC posterior sampling and some of them are not based on approximating the models' marginal likelihoods (e.g., Ando, 2007; Spiegelhalter, Best, Carlin, & van der Linde, 2002). A comparison of these methods is beyond the scope of this tutorial.

Table 2

Bayes factors comparing the full EV model to the restricted EV models, log marginal likelihoods, and coefficient of variation (with respect to the marginal likelihood) expressed as a percentage.

Model	Bayes Factor	Log Marginal Likelihood	CV[%]
Full model	–	–3800.434	10.13
Restricted at $\mu_{\omega} = -0.334$	1.202	–3800.618	16.44
Restricted at $\mu_{\alpha} = -0.604$	1.052	–3800.484	9.71
Restricted at $\mu_{\gamma} = 0.92$	1.068	–3800.500	12.03

Table A.1

Summary of the bridge sampling estimators for the marginal likelihood, and its special cases: the naive Monte Carlo, importance sampling, and generalized harmonic mean estimator.

Method	Estimator	Samples	Bridge Function $h(\theta)$
Bridge sampling	$\frac{\frac{1}{N_2} \sum_{i=1}^{N_2} p(y \tilde{\theta}_i) p(\tilde{\theta}_i) h(\tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} h(\theta_j^*) g(\theta_j^*)}$	$\tilde{\theta}_i \sim g(\theta)$ $\theta_j^* \sim p(\theta y)$	$h(\theta) = \frac{C}{\frac{N_1}{N_2+N_1} p(y \theta) p(\theta) + \frac{N_2}{N_2+N_1} p(y g(\theta))}$
Naive Monte Carlo	$\frac{1}{N} \sum_{i=1}^N p(y \tilde{\theta}_i)$	$\tilde{\theta}_i \sim p(\theta)$	$h(\theta) = \frac{1}{g(\theta)}$ and $g(\theta) = p(\theta)$
Importance sampling	$\frac{1}{N} \sum_{i=1}^N \frac{p(y \tilde{\theta}_i) p(\tilde{\theta}_i)}{g_{IS}(\tilde{\theta}_i)}$	$\tilde{\theta}_i \sim g_{IS}(\theta)$	$h(\theta) = \frac{1}{g(\theta)}$ and $g(\theta) = g_{IS}(\theta)$
Generalized harmonic mean	$\left(\frac{1}{N} \sum_{i=1}^N \frac{g_{IS}(\theta_i^*)}{p(y \theta_i^*) p(\theta_i^*)} \right)^{-1}$	$\theta_i^* \sim p(\theta y)$	$h(\theta) = \frac{1}{p(y \theta) p(\theta)}$ and $g(\theta) = g_{IS}(\theta)$

Note. $p(\theta)$ is the prior distribution, $g_{IS}(\theta)$ is the importance density, $p(\theta|y)$ is the posterior distribution, $g(\theta)$ is the proposal distribution, $h(\theta)$ is the bridge function, and C is a constant. The last column shows the bridge function needed to obtain the special cases.

Acknowledgments

We thank [Busemeyer and Stout \(2002\)](#) for providing the data used in this article.

Funding for this research was provided by the Berkeley Initiative for Transparency in the Social Sciences, a program of the Center for Effective Global Action (CEGA), with support from the Laura and John Arnold Foundation. This work was supported in part by a Vici grant from the Netherlands Organisation for Scientific Research (NWO) (016.Vici.170.083). This research was furthermore supported by an NWO grant to QFG (406-16-528), UB (406-12-125), and to HS (404-10-086), a European Research Council (ERC) grant to AL and EJW (283876), and a Veni grant (451-15-010) from the NWO to DM.

Appendix A. The bridge sampling estimator as a general case of methods 1–3

In this section we show that the naive Monte Carlo, the importance sampling, and the generalized harmonic mean estimators are special cases of the bridge sampling estimator under specific choices of the bridge function $h(\theta)$ and the proposal distribution $g(\theta)$.¹⁸ An overview is provided in [Table A.1](#).

To prove that the bridge sampling estimator reduces to the naive Monte Carlo estimator, consider bridge sampling, choose the prior distribution as the proposal distribution (i.e., $g(\theta) = p(\theta)$), and specify the bridge function as $h(\theta) = 1/g(\theta)$. Inserting these specifications into [Eq. \(12\)](#) yields:

$$\begin{aligned} \hat{p}_4(y | h(\theta) = \frac{1}{g(\theta)}, g(\theta) = p(\theta)) &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{1}{p(\tilde{\theta}_i)} p(y | \tilde{\theta}_i) p(\tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{p(\theta_j^*)} p(\theta_j^*)} \\ &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} p(y | \tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} p(\theta_j^*)} = \frac{1}{N_2} \sum_{i=1}^{N_2} p(y | \tilde{\theta}_i), \quad \tilde{\theta}_i \sim p(\theta), \quad \theta_j^* \sim p(\theta | y) \end{aligned}$$

which is equivalent to the naive Monte Carlo estimator shown in [Eq. \(7\)](#).

To prove that the bridge sampling estimator reduces to the importance sampling estimator, consider bridge sampling, choose the importance density as the proposal distribution (i.e., $g(\theta) = g_{IS}(\theta)$), and specify the bridge function as $h(\theta) = 1/g(\theta)$. Inserting these specifications into [Eq. \(12\)](#) yields:

$$\begin{aligned} \hat{p}_4(y | h(\theta) = \frac{1}{g(\theta)}, g(\theta) = g_{IS}(\theta)) &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{1}{g_{IS}(\tilde{\theta}_i)} p(y | \tilde{\theta}_i) p(\tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{g_{IS}(\theta_j^*)} g_{IS}(\theta_j^*)} \\ &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{p(y|\tilde{\theta}_i) p(\tilde{\theta}_i)}{g_{IS}(\tilde{\theta}_i)}}{\frac{1}{N_1} \sum_{j=1}^{N_1} g_{IS}(\theta_j^*)} \\ &= \frac{1}{N_2} \sum_{i=1}^{N_2} \frac{p(y | \tilde{\theta}_i) p(\tilde{\theta}_i)}{g_{IS}(\tilde{\theta}_i)}, \quad \tilde{\theta}_i \sim g_{IS}(\theta), \quad \theta_j^* \sim p(\theta | y) \end{aligned}$$

which is equivalent to the importance sampling estimator shown in [Eq. \(8\)](#).

To prove that the bridge sampling estimator reduces to the generalized harmonic mean estimator, consider bridge sampling, choose the importance density as the proposal distribution (i.e., $g(\theta) = g_{IS}(\theta)$), and specify the bridge function as $h(\theta) = 1/(p(y | \theta) p(\theta))$. Inserting these specifications into [Eq. \(12\)](#) yields:

$$\begin{aligned} \hat{p}_4(y | h(\theta) = \frac{1}{p(y | \theta) p(\theta)}, g(\theta) = g_{IS}(\theta)) &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{1}{p(y|\tilde{\theta}_i) p(\tilde{\theta}_i)} p(y | \tilde{\theta}_i) p(\tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{p(y|\theta_j^*) p(\theta_j^*)} g_{IS}(\theta_j^*)} \\ &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} g_{IS}(\tilde{\theta}_i)}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{g_{IS}(\theta_j^*)}{p(y|\theta_j^*) p(\theta_j^*)}} = \left(\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{g_{IS}(\theta_j^*)}{p(y|\theta_j^*) p(\theta_j^*)} \right)^{-1}, \quad \tilde{\theta}_i \sim g(\theta), \quad \theta_j^* \sim p(\theta | y) \end{aligned}$$

which is equivalent to the generalized harmonic mean estimator shown in [Eq. \(9\)](#).

¹⁸ Note that bridge sampling is also a general case of the [Chib and Jeliazkov \(2001\)](#) method of estimating the marginal likelihood using the Metropolis–Hastings acceptance probability ([Meng & Schilling, 2002](#); [Mira & Nicholls, 2004](#)).

Appendix B. Bridge sampling implementation: avoiding numerical issues

In order to avoid numerical issues, we can rewrite Eq. (15) in the following way:

$$\begin{aligned}
 \hat{p}_4(y)^{(t+1)} &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{l_{2,i}}{s_1 l_{2,i} + s_2 \hat{p}_4(y)^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 l_{1,j} + s_2 \hat{p}_4(y)^{(t)}}} \\
 &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i}))}{s_1 \exp(\log(l_{2,i})) + s_2 \hat{p}_4(y)^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 \exp(\log(l_{1,j})) + s_2 \hat{p}_4(y)^{(t)}}} \\
 &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i})) \exp(-l^*)}{s_1 \exp(\log(l_{2,i})) \exp(-l^*) + s_2 \hat{p}_4(y)^{(t)} \exp(-l^*)}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{\exp(-l^*)}{s_1 \exp(\log(l_{1,j})) \exp(-l^*) + s_2 \hat{p}_4(y)^{(t)} \exp(-l^*)}} \\
 &= \frac{1}{\exp(-l^*)} \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i}) - l^*)}{s_1 \exp(\log(l_{2,i}) - l^*) + s_2 \hat{p}_4(y)^{(t)} \exp(-l^*)}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 \exp(\log(l_{1,j}) - l^*) + s_2 \hat{p}_4(y)^{(t)} \exp(-l^*)}} \\
 &= \exp(l^*) \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i}) - l^*)}{s_1 \exp(\log(l_{2,i}) - l^*) + s_2 \hat{p}_4(y)^{(t)} \exp(-l^*)}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 \exp(\log(l_{1,j}) - l^*) + s_2 \hat{p}_4(y)^{(t)} \exp(-l^*)}}.
 \end{aligned}$$

l^* is a constant which we can choose in a way that keeps the terms in the sums manageable. We used $l^* = \text{median}(\log(l_{1,j}))$. Let

$$\hat{r}^{(t)} = \hat{p}_4(y)^{(t)} \exp(-l^*),$$

so that

$$\hat{p}_4(y)^{(t)} = \hat{r}^{(t)} \exp(l^*).$$

Then we obtain

$$\begin{aligned}
 \hat{p}_4(y)^{(t+1)} &= \exp(l^*) \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i}) - l^*)}{s_1 \exp(\log(l_{2,i}) - l^*) + s_2 \hat{r}^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 \exp(\log(l_{1,j}) - l^*) + s_2 \hat{r}^{(t)}}} \\
 \hat{p}_4(y)^{(t+1)} \exp(-l^*) &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i}) - l^*)}{s_1 \exp(\log(l_{2,i}) - l^*) + s_2 \hat{r}^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 \exp(\log(l_{1,j}) - l^*) + s_2 \hat{r}^{(t)}}} \\
 \hat{r}^{(t+1)} &= \frac{\frac{1}{N_2} \sum_{i=1}^{N_2} \frac{\exp(\log(l_{2,i}) - l^*)}{s_1 \exp(\log(l_{2,i}) - l^*) + s_2 \hat{r}^{(t)}}}{\frac{1}{N_1} \sum_{j=1}^{N_1} \frac{1}{s_1 \exp(\log(l_{1,j}) - l^*) + s_2 \hat{r}^{(t)}}}.
 \end{aligned}$$

Hence, we can run the iterative scheme with respect to $\hat{r}$ which is more convenient because it keeps the terms in the sums manageable and multiply the result by $\exp(l^*)$ to obtain the estimate of the marginal likelihood or, equivalently, we can take the logarithm of the result and add l^* to obtain an estimate of the logarithm of the marginal likelihood.

Appendix C. Correcting for the probit transformation

In this section we describe how the probit transformation affects our expression of the generalized harmonic mean estimator (Eq. (9)) to yield Eq. (10). Recall that we derived the generalized harmonic mean estimator using the following equality:

$$\frac{1}{p(y)} = \int \frac{g_{IS}(\theta)}{p(y | \theta)p(\theta)} p(\theta | y) d\theta. \quad (C.1)$$

For practical reasons, in the running example, we used a normal distribution on ξ as importance density. This ξ was defined as the probit transform of θ (i.e., $\xi = \Phi^{-1}(\theta)$). In particular, the normal importance density was given by $\frac{1}{\hat{\sigma}} \phi\left(\frac{\xi - \hat{\mu}}{\hat{\sigma}}\right)$. Note that this importance density is a function of ξ , whereas the general importance density g_{IS} in Eq. (C.1) is specified in terms of θ . Therefore, to include our specific importance density into Eq. (C.1), we need to write it in terms of θ . This yields $\frac{1}{\hat{\sigma}} \phi\left(\frac{\Phi^{-1}(\theta) - \hat{\mu}}{\hat{\sigma}}\right) \frac{1}{\phi(\Phi^{-1}(\theta))}$, where the latter factor comes from applying the change-of-variable method. Replacing $g_{IS}(\theta)$ in Eq. (C.1) by this expression, results in:

$$\begin{aligned}
 \frac{1}{p(y)} &= \int \frac{\frac{1}{\hat{\sigma}} \phi\left(\frac{\Phi^{-1}(\theta) - \hat{\mu}}{\hat{\sigma}}\right) \frac{1}{\phi(\Phi^{-1}(\theta))}}{p(y | \theta)p(\theta)} p(\theta | y) d\theta \\
 &= \mathbb{E}_{\text{post}} \left(\frac{\frac{1}{\hat{\sigma}} \phi\left(\frac{\Phi^{-1}(\theta) - \hat{\mu}}{\hat{\sigma}}\right) \frac{1}{\phi(\Phi^{-1}(\theta))}}{p(y | \theta)p(\theta)} \right). \quad (C.2)
 \end{aligned}$$

Rewriting results in:

$$p(y) = \left(\mathbb{E}_{\text{post}} \left(\frac{\frac{1}{\hat{\sigma}} \phi\left(\frac{\Phi^{-1}(\theta) - \hat{\mu}}{\hat{\sigma}}\right) \frac{1}{\phi(\Phi^{-1}(\theta))}}{p(y | \theta)p(\theta)} \right) \right)^{-1},$$

which can be approximated as:

$$\begin{aligned}
 \hat{p}_3(y) &= \left(\frac{1}{N} \sum_{j=1}^N \frac{\overbrace{\frac{1}{\hat{\sigma}} \phi\left(\frac{\Phi^{-1}(\theta_j^*) - \hat{\mu}}{\hat{\sigma}}\right) \frac{1}{\phi(\Phi^{-1}(\theta_j^*))}}^{\text{importance density}}}{\underbrace{p(y | \theta_j^*)}_{\text{likelihood}} \underbrace{p(\theta_j^*)}_{\text{prior}}} \right)^{-1}, \\
 &\quad \underbrace{\theta_j^* \sim p(\theta | y)}_{\text{samples from the posterior distribution}} \\
 &= \left(\frac{1}{N} \sum_{j=1}^N \frac{\overbrace{\frac{1}{\hat{\sigma}} \phi\left(\frac{\xi_j^* - \hat{\mu}}{\hat{\sigma}}\right)}^{\text{importance density}}}{\underbrace{p(y | \Phi(\xi_j^*))}_{\text{likelihood}} \underbrace{p(\Phi(\xi_j^*)) \phi(\xi_j^*)}_{\text{prior}}} \right)^{-1}, \\
 &\quad \underbrace{\xi_j^* = \Phi^{-1}(\theta_j^*) \text{ and } \theta_j^* \sim p(\theta | y)}_{\text{probit-transformed samples from the posterior distribution}} \quad (C.3)
 \end{aligned}$$

which shows that the generalized harmonic estimate can be obtained using the samples from the posterior distribution, or the probit-transformed ones. In the online-provided code, we use the latter approach (see also Overstall & Forster, 2010). Note that in our running example, $\forall \xi_j^* : p(\Phi(\xi_j^*)) = 1$.

Appendix D. Details on the application of bridge sampling to the individual-level EV model

In this section, we provide more details on how we obtained the unnormalized posterior distribution for a specific participant s , $s \in \{1, 2, \dots, 30\}$. Since we focus on one specific participant, we drop the subscript s in the remainder of this section. As explained in Appendix B, we run the iterative scheme with respect to $\hat{r}$ to avoid numerical issues. Consequently, we have to compute $\log(l_{1,j})$ and $\log(l_{2,i})$. Using $\tilde{\kappa}_i = (\tilde{\omega}_i, \tilde{\alpha}_i, \tilde{\gamma}_i)$ for the i th sample from the proposal

distribution, we get for $\log(l_{2,i})$ ($\log(l_{1,j})$ works analogously):

$$\log(l_{2,i}) = \log \left(\frac{p(\text{Ch}(T) | \Phi(\tilde{\kappa}_i), X(T-1)) p(\Phi(\tilde{\kappa}_i)) \phi(\tilde{\kappa}_i)}{g(\tilde{\kappa}_i)} \right).$$

Therefore, instead of computing the unnormalized posterior distribution directly, we compute the logarithm of the unnormalized posterior distribution:

$$\begin{aligned} \log(p(\text{Ch}(T) | \Phi(\tilde{\kappa}_i), X(T-1)) p(\Phi(\tilde{\kappa}_i)) \phi(\tilde{\kappa}_i)) \\ = \log(p(\text{Ch}(T) | \Phi(\tilde{\kappa}_i), X(T-1))) + \\ \log(\phi(\tilde{\omega}_i)) + \log(\phi(\tilde{\alpha}_i)) + \log(\phi(\tilde{\gamma}_i)), \end{aligned}$$

because we assumed independent priors on each model parameter w, a, c . $\log(p(\Phi(\tilde{\kappa}_i))) = 0$ because p refers to the uniform prior on $[0, 1]$.

Appendix E. Details on the application of bridge sampling to the hierarchical EV model

Analogous to the last section, we explain here how we obtained the logarithm of the unnormalized posterior for the hierarchical implementation of the EV model. Using $\tilde{\kappa}_{s,i} = (\tilde{\omega}_{s,i}, \tilde{\alpha}_{s,i}, \tilde{\gamma}_{s,i})$ for the i th sample from the proposal distribution for the individual-level parameters of subject s , and $\tilde{\zeta}_i$ for the i th sample from the proposal distribution for all group-level parameters (i.e., $\tilde{\zeta}_i = (\tilde{\mu}_{\omega,i}, \tilde{\tau}_{\omega,i}, \tilde{\mu}_{\alpha,i}, \tilde{\tau}_{\alpha,i}, \tilde{\mu}_{\gamma,i}, \tilde{\tau}_{\gamma,i})$), we get:

$$\begin{aligned} \log \left(\left(\prod_{s=1}^{30} p(\text{Ch}_s(T) | \Phi(\tilde{\kappa}_{s,i}), X_s(T-1)) p(\tilde{\kappa}_{s,i} | \tilde{\zeta}_i) \right) p(\tilde{\zeta}_i) \right) \\ = \sum_{s=1}^N [\log(p(\text{Ch}_s(T) | \Phi(\tilde{\kappa}_{s,i}), X_s(T-1))) + \\ \log \left(\frac{1}{1.5\Phi(\tilde{\tau}_{\omega,i})} \phi \left(\frac{\tilde{\omega}_{s,i} - \tilde{\mu}_{\omega,i}}{1.5\Phi(\tilde{\tau}_{\omega,i})} \right) \right) \\ + \log \left(\frac{1}{1.5\Phi(\tilde{\tau}_{\alpha,i})} \phi \left(\frac{\tilde{\alpha}_{s,i} - \tilde{\mu}_{\alpha,i}}{1.5\Phi(\tilde{\tau}_{\alpha,i})} \right) \right) + \\ \log \left(\frac{1}{1.5\Phi(\tilde{\tau}_{\gamma,i})} \phi \left(\frac{\tilde{\gamma}_{s,i} - \tilde{\mu}_{\gamma,i}}{1.5\Phi(\tilde{\tau}_{\gamma,i})} \right) \right)] + \\ \log(\phi(\tilde{\mu}_{\omega,i})) + \log(\phi(\tilde{\mu}_{\alpha,i})) + \log(\phi(\tilde{\mu}_{\gamma,i})) + \\ \log(\phi(\tilde{\tau}_{\omega,i})) + \log(\phi(\tilde{\tau}_{\alpha,i})) + \log(\phi(\tilde{\tau}_{\gamma,i})). \end{aligned}$$

References

- Ahn, W.-Y., Busemeyer, J. R., Wagenmakers, E.-J., & Stout, J. C. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive Science*, 32, 1376–1402.
- Ahn, W.-Y., Krawitz, A., Kim, W., Busemeyer, J. R., & Brown, J. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience Psychology and Economics*, 4, 95–110.
- Ando, T. (2007). Bayesian predictive information criterion for the evaluation of hierarchical Bayesian and empirical Bayes models. *Biometrika*, 94, 443–458.
- Andrews, M., & Baguley, T. (2013). Prior approval: The growth of Bayesian methods in psychology. *The British Journal of Mathematical and Statistical Psychology*, 66, 1–7.
- Bark, R., Dieckmann, S., Bogerts, B., & Northoff, G. (2005). Deficit in decision making in catatonic schizophrenia: An exploratory study. *Psychiatry Research*, 134, 131–141.
- Bayarri, M., Benjamin, D. J., Berger, J. O., & Sellke, T. M. (2016). Rejection odds and rejection ratios: A proposal for statistical practice in testing hypotheses. *Journal of Mathematical Psychology*, 72, 90–103.
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50, 7–15.
- Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *Journal of Neuroscience*, 19, 5473–5481.

- Bechara, A., Damasio, H., Tranel, D., & Anderson, S. W. (1998). Dissociation of working memory from decision making within the human prefrontal cortex. *Journal of Neuroscience*, 18, 428–437.
- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, 275, 1293–1295.
- Bechara, A., Tranel, D., & Damasio, H. (2000). Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions. *Brain*, 123, 2189–2202.
- Bennett, C. H. (1976). Efficient estimation of free energy differences from Monte Carlo data. *Journal of Computational Physics*, 22, 245–268.
- Berger, J. O. (2006). Bayes factors. In S. Kotz, N. Balakrishnan, C. Read, B. Vidakovic, & N. L. Johnson (Eds.), *Encyclopedia of statistical sciences*, Vol. 1. (2nd ed.). (pp. 378–386). Hoboken, NJ: Wiley.
- Berger, J. O., & Molina, G. (2005). Posterior model probabilities via path-based pairwise priors. *Statistica Neerlandica*, 59, 3–15.
- Blair, R. J. R., Colledge, E., & Mitchell, D. G. V. (2001). Somatic markers and response reversal: Is there orbitofrontal cortex dysfunction in boys with psychopathic tendencies? *Journal of Abnormal Child Psychology*, 29, 499–511.
- Brown, C. E. (1998). Coefficient of variation. In *Applied multivariate statistics in geohydrology and related sciences* (pp. 155–157). Springer.
- Busemeyer, J. R., & Stout, J. C. (2002). A Contribution of cognitive decision models to clinical assessment: Decomposing performance on the bechara gambling task. *Psychological Assessment*, 14, 253–262.
- Carlin, B. P., & Chib, S. (1995). Bayesian model choice via Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 473–484.
- Cavedini, P., Riboldi, G., D'Annunzi, A., Belotti, P., Cisima, M., & Bellodi, L. (2002). Decision-making heterogeneity in obsessive-compulsive disorder: Ventromedial prefrontal cortex function predicts different treatment outcomes. *Neuropsychologia*, 40, 205–211.
- Cavedini, P., Riboldi, G., Keller, R., D'Annunzi, A., & Bellodi, L. (2002). Frontal lobe dysfunction in pathological gambling patients. *Biological Psychiatry*, 51, 334–341.
- Chen, M.-H., Shao, Q.-M., & Ibrahim, J. G. (2012). *Monte Carlo methods in Bayesian computation*. Springer Science & Business Media.
- Chib, S., & Jeliazkov, I. (2001). Marginal likelihood from the Metropolis–Hastings output. *Journal of the American Statistical Association*, 96, 270–281.
- Dai, J., Kerestes, R., Upton, D. J., Busemeyer, J. R., & Stout, J. C. (2015). An improved cognitive model of the Iowa and Soochow Gambling Tasks with regard to model fitting performance and tests of parameter consistency. *Frontiers in Psychology*, 6, 229.
- DiCiccio, T. J., Kass, R. E., Raftery, A., & Wasserman, L. (1997). Computing Bayes factors by combining simulation and asymptotic approximations. *Journal of the American Statistical Association*, 92, 903–915.
- Dickey, J. M. (1971). The weighted likelihood ratio, linear hypotheses on normal location parameters. *The Annals of Mathematical Statistics*, 204–223.
- Dickey, J. M., & Lientz, B. (1970). The weighted likelihood ratio, sharp hypotheses about chances, the order of a Markov chain. *The Annals of Mathematical Statistics*, 41, 214–226.
- Didelot, X., Everitt, R. G., Johansen, A. M., & Lawson, D. J. (2011). Likelihood-free estimation of model evidence. *Bayesian Analysis*, 6(1), 49–76.
- Etz, A., & Wagenmakers, E.-J. (in press). J. B. S. Haldane's contribution to the Bayes factor hypothesis test. *Statist. Sci.*
- Frühwirth-Schnatter, S. (2004). Estimating marginal likelihoods for mixture and Markov switching models using bridge sampling techniques. *The Econometrics Journal*, 7, 143–167.
- Gamerman, D., & Lopes, H. F. (2006). *Markov chain Monte Carlo: Stochastic simulation for Bayesian inference*. (pp. 237–286). CRC Press.
- Gelfand, A. E., & Dey, D. K. (1994). Bayesian model choice: asymptotics and exact calculations. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 501–514.
- Gelman, A., & Meng, X.-L. (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Science*, 163–185.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82, 711–732.
- Gronau, Q. F., Singmann, H., & Wagenmakers, E.-J. (2017). Bridgesampling: Bridge sampling for marginal likelihoods and Bayes factors. Retrieved from <https://github.com/quentingronau/bridgesampling> (R package version 0.2-2).
- Hammersley, J. M., & Handscomb, D. C. (1964). *Monte Carlo methods*. London: Methuen.
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science*, 382–401.

- Ionides, E. L. (2008). Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 17, 295–311.
- Jeffreys, H. (1961). *Theory of probability*. (3rd ed.). Oxford, England: Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795.
- Lee, M. D. (2008). Three case studies in the Bayesian analysis of cognitive models. *Psychonomic Bulletin & Review*, 15, 1–15.
- Lewis, S. M., & Raftery, A. E. (1997). Estimating Bayes factors via posterior simulation with the Laplace-Metropolis estimator. *Journal of the American Statistical Association*, 92, 648–655.
- Lodewyckx, T., Kim, W., Lee, M. D., Tuerlinckx, F., Kuppens, P., & Wagenmakers, E.-J. (2011). A tutorial on Bayes factor estimation with the product space method. *Journal of Mathematical Psychology*, 55, 331–347.
- Luce, R. (1959). *Individual choice behavior*. New York: Wiley.
- Lunn, D. J., Best, N., & Whittaker, J. C. (2009). Generic reversible jump MCMC using graphical models. *Statistics and Computing*, 19, 395–408.
- Lunn, D. J., Thomas, A., Best, N., & Spiegelhalter, D. (2000). WinBUGS—a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, 10, 325–337.
- Ly, A., Verhagen, A., & Wagenmakers, E.-J. (2016a). An evaluation of alternative methods for testing hypotheses, from the perspective of Harold Jeffreys. *Journal of Mathematical Psychology*, 72, 43–55.
- Ly, A., Verhagen, J., & Wagenmakers, E.-J. (2016b). Harold Jeffreys's default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72, 19–32.
- Martino, D. J., Bucay, D., Butman, J. T., & Allegri, R. F. (2007). Neuropsychological frontal impairments and negative symptoms in schizophrenia. *Psychiatry Research*, 152, 121–128.
- Matzke, D., Dolan, C. V., Batchelder, W. H., & Wagenmakers, E.-J. (2015). Bayesian estimation of multinomial processing tree models with heterogeneity in participants and items. *Psychometrika*, 80, 205–235.
- Matzke, D., & Wagenmakers, E.-J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review*, 16, 798–817.
- Meng, X.-L., & Schilling, S. (2002). Warp bridge sampling. *Journal of Computational and Graphical Statistics*, 11, 552–586.
- Meng, X.-L., & Wong, W. H. (1996). Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statistica Sinica*, 831–860.
- Mira, A., & Nicholls, G. (2004). Bridge estimation of the probability density at a point. *Statistica Sinica*, 14, 603–612.
- Mulder, J., & Wagenmakers, E.-J. (2016). Editors' introduction to the special issue Bayes factors for testing hypotheses in psychological research: Practical relevance and new developments. *Journal of Mathematical Psychology*, 72, 1–5.
- Myung, I. J., Forster, M. R., & Browne, M. W. (2000). Guest editors' introduction: Special issue on model selection. *Journal of Mathematical Psychology*, 44, 1–2.
- Navarro, D. J., Griffiths, T. L., Steyvers, M., & Lee, M. D. (2006). Modeling individual differences using Dirichlet processes. *Journal of Mathematical Psychology*, 50, 101–122.
- Neal, R. M. (2001). Annealed importance sampling. *Statistics and Computing*, 11, 125–139.
- Newton, M. A., & Raftery, A. E. (1994). Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 3–48.
- Ntzoufras, I. (2009). Bayesian model and variable evaluation. In *Bayesian modeling using WinBUGS* (pp. 389–433). John Wiley & Sons.
- Overstall, A. M., & Forster, J. J. (2010). Default Bayesian model determination methods for generalised linear mixed models. *Computational Statistics & Data Analysis*, 54, 3269–3288.
- Owen, A., & Zhou, Y. (2000). Safe and effective importance sampling. *Journal of the American Statistical Association*, 95, 135–143.
- Pajor, A. (2016). Estimating the marginal likelihood using the arithmetic mean identity. *Bayesian Analysis*, 1–27.
- Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, 109(3), 472–491.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International workshop on distributed statistical computing*, Vol. 124, (pp. 1–8).
- Plummer, M., Best, N., Cowles, K., & Vines, K. (2006). CODA: Convergence diagnosis and output analysis for MCMC. *R News*, 6, 7–11.
- Poirier, D. J. (2006). The growth of Bayesian methods in statistics and economics since 1970. *Bayesian Analysis*, 1, 969–979.
- Raftery, A. E., & Banfield, J. D. (1991). Stopping the gibbs sampler, the use of morphology, and other issues in spatial statistics (Bayesian image restoration, with two applications in spatial statistics)–(discussion). *Annals of the Institute of Statistical Mathematics*, 43, 32–43.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy, (Eds.), *Classical conditioning II: Current research and theory*, New York, (pp. 64–99).
- Robert, C. (2016). The expected demise of the Bayes factor. *Journal of Mathematical Psychology*, 72, 33–37.
- Rouder, J. N., & Lu, J. (2005). An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychonomic Bulletin & Review*, 12, 573–604.
- Rouder, J. N., Lu, J., Morey, R. D., Sun, D., & Speckman, P. L. (2008). A hierarchical process-dissociation model. *Journal of Experimental Psychology: General*, 137, 370–389.
- Rouder, J. N., Lu, J., Speckman, P., Sun, D., & Jiang, Y. (2005). A hierarchical model for estimating response time distributions. *Psychonomic Bulletin & Review*, 12, 195–223.
- Rouder, J. N., Lu, J., Sun, D., Speckman, P., Morey, R., & Naveh-Benjamin, M. (2007). Signal detection models with random participant and item effects. *Psychometrika*, 72, 621–642.
- Scheibehenne, B., & Pachur, T. (2015). Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychonomic Bulletin & Review*, 22, 391–407.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6, 461–464.
- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32, 1248–1284.
- Sinharay, S., & Stern, H. S. (2005). An empirical comparison of methods for computing Bayes factors in generalized linear mixed models. *Journal of Computational and Graphical Statistics*, 14, 415–435.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society. Series B*, 64, 583–639.
- Stan Development Team. (2016). RStan: the R interface to Stan. Retrieved from <http://mc-stan.org/> (R package version 2.14.1).
- Steingroever, H., Pachur, T., Šmíra, M., & Lee, M. D. (submitted for publication). Bayesian techniques for analyzing group differences in the Iowa gambling task: A case study of intuitive and deliberate decision makers.
- Steingroever, H., Wetzels, R., Horstmann, A., Neumann, J., & Wagenmakers, E.-J. (2013). Performance of healthy participants on the Iowa gambling task. *Psychological Assessment*, 25, 180–193.
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2013a). A Comparison of reinforcement-learning models for the Iowa gambling task using parameter space partitioning. *The Journal of Problem Solving*, 5, Article 2.
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2013b). Validating the PVL-Delta model for the Iowa gambling task. *Frontiers in Psychology*, 4, 898.
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2014). Absolute performance of reinforcement-learning models for the Iowa gambling task. *Decision*, 1, 161–183.
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2016). Bayes factors for reinforcement-learning models of the Iowa gambling task. *Decision*, 3, 115–131.
- Stone, C. J., Hansen, M. H., Kooperberg, C., Truong, Y. K., et al. (1997). Polynomial splines and their tensor products in extended linear modeling: 1994 Wald memorial lecture. *The Annals of Statistics*, 25, 1371–1470.
- Vandekerckhove, J., Matzke, D., & Wagenmakers, E.-J. (2015). Model comparison and the principle of parsimony. In J. Busemeyer, J. Townsend, Z. J. Wang, & A. Eidels (Eds.), *Oxford handbook of computational and mathematical psychology*. Oxford: Oxford University Press.
- Vanpaemel, W. (2016). Prototypes, exemplars and the response scaling parameter: A Bayes factor perspective. *Journal of Mathematical Psychology*, 72, 183–190.
- Verhagen, J., Levy, R., Millsap, R. E., & Fox, J.-P. (2015). Evaluating evidence for invariant items: A Bayes factor applied to testing measurement invariance in IRT models. *Journal of Mathematical Psychology*, 72, 171–182.
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology*, 60, 158–189.
- Wagenmakers, E.-J., & Waldorp, L. (2006). Editors' introduction. *Journal of Mathematical Psychology*, 50, 99–100.
- Wang, L., & Meng, X.-L. (2016). Warp bridge sampling: The next generation. arXiv preprint arXiv:1609.07690.
- Wetzels, R., Grasman, R. P., & Wagenmakers, E.-J. (2010a). An encompassing prior generalization of the Savage–Dickey density ratio. *Computational Statistics & Data Analysis*, 54, 2094–2102.
- Wetzels, R., Tutschko, D., Dolan, C., van der Sluis, S., Dutilh, G., & Wagenmakers, E.-J. (2016). A Bayesian test for the hot hand phenomenon. *Journal of Mathematical Psychology*, 72, 200–209.
- Wetzels, R., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E.-J. (2010). Bayesian parameter estimation in the Expectancy Valence model of the Iowa gambling task. *Journal of Mathematical Psychology*, 54, 14–27.
- Worthy, D. A., & Maddox, W. T. (2014). A comparison model of reinforcement-learning and win-stay-lose-shift decision-making processes: A tribute to W. K. Estes. *Journal of Mathematical Psychology*, 59, 41–49.
- Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in Psychology*, 4.