Robustly estimating the marginal likelihood for cognitive models via importance sampling

M.-N. Tran¹ · M. Scharth¹ · D. Gunawan² · R. Kohn³ · S. D. Brown⁴ · G. E. Hawkins⁴

Published online: 1 October 2020
© The Psychonomic Society, Inc. 2020

Abstract

Recent advances in Markov chain Monte Carlo (MCMC) extend the scope of Bayesian inference to models for which the likelihood function is intractable. Although these developments allow us to estimate model parameters, other basic problems such as estimating the marginal likelihood, a fundamental tool in Bayesian model selection, remain challenging. This is an important scientific limitation because testing psychological hypotheses with hierarchical models has proven difficult with current model selection methods. We propose an efficient method for estimating the marginal likelihood for models where the likelihood is intractable, but can be estimated unbiasedly. It is based on first running a sampling method such as MCMC to obtain samples for the model parameters, and then using these samples to construct the proposal density in an importance sampling (IS) framework with an unbiased estimate of the likelihood. Our method has several attractive properties: it generates an unbiased estimate of the marginal likelihood, it is robust to the quality and target of the sampling method used to form the IS proposals, and it is computationally cheap to estimate the variance of the marginal likelihood estimator. We also obtain the convergence properties of the method and provide guidelines on maximizing computational efficiency. The method is illustrated in two challenging cases involving hierarchical models: identifying the form of individual differences in an applied choice scenario, and evaluating the best parameterization of a cognitive model in a speeded decision making context. Freely available code to implement the methods is provided. Extensions to posterior moment estimation and parallelization are also discussed.

Keywords Bayesian inference · Hierarchical LBA model · Model selection · Parallel computation · Standard error · Unbiased likelihood estimate

Introduction

Many psychologically interesting research questions involve comparing competing theories: Does sleep deprivation cause attentional lapses? Does alcohol impair the speed of information processing or reduce cautiousness, or both? Does the forgetting curve follow a power or expo-

nential function? In many cases, the competing theories can be represented as a set of quantitative models that are applied (“fitted”) to the observed data. We can then estimate a metric that quantifies the degree to which each model accounts for the patterns observed in data balanced against its flexibility. Model flexibility is often defined as the range of data patterns that a model can predict, which includes patterns that were observed as well as patterns that were not observed. Flexibility is an important consideration in model choice as it assesses the degree to which a theory is suitably constrained to predict the observed data and that it does not simply predict many possible patterns that were not observed; the latter form of model is theoretically non-informative as it “predicts” almost any pattern that could be observed. Many methods exist that attempt to quantitatively account for the flexibility of models, and hence inform model choice, including likelihood ratio tests, various information criteria (e.g., Akaike, Bayesian and Deviance Information Criteria; AIC, BIC, and DIC, respectively), minimum description length, cross validation

✉ G. E. Hawkins
guy.e.hawkins@gmail.com

¹ The University of Sydney Business School, Sydney, Australia

² School of Mathematics and Statistics, University of Wollongong, Wollongong, Australia

³ UNSW Business School, University of New South Wales, Kensington, Australia

⁴ School of Psychology, University of Newcastle, Callaghan, Australia

and others, with varying degrees of success. The emerging gold standard in quantitative psychology is the marginal likelihood; the commonly cited Bayes factor is the ratio of the marginal likelihoods of two models.

The marginal likelihood of a model is the average likelihood of the model given the data, where the average is taken over the prior distribution of the model parameters. By integrating the likelihood across the prior predictive space of a model, the marginal likelihood has the attractive property that it inherently accounts for the flexibility of the model. This removes the need for post-hoc complexity corrections to goodness of fit measures, as implemented in the “information criteria” metrics (e.g., AIC, BIC, DIC), because it provides a guarantee that a model which accounts for the observed patterns in data and few other possible though unobserved patterns will be quantitatively favored over a competing model which accounts for the observed data but also provides the capacity to account for many other patterns that were not observed. This occurs because the latter model would have low likelihood across the (many) regions of the prior space where data were not observed, which lowers the marginal likelihood.

Although theoretically attractive, estimating the marginal likelihood for psychologically interesting models poses a number of practical challenges. The likelihood function is the density function of the data with the participant-level parameters (also called random effects) integrated out, when viewed as a function of the group-level parameters. It is usually analytically or computationally intractable for psychological and statistical models because it is an integral that cannot be evaluated. This can be the case even in conceptually very simple models. For example, in generalized linear mixed models where random effects (e.g., model parameters for individual participants) account for the dependence between observations from the same individual (Fitzmaurice, Laird, & Ware, 2011), the likelihood is often intractable because it is an integral over the random effects. Recent advances in particle Markov chain Monte Carlo (P-MCMC) methods extend the scope of the applications of Bayesian inference to cases where the likelihood is intractable (see Andrieu and Roberts (2009), Andrieu et al. (2010), Chopin et al. (2013), and Tran et al. (2016)). However, some basic problems such as estimating the marginal likelihood and the standard error of this estimator remain challenging, because the marginal likelihood is the density of the data with both the group-level parameters and random effects integrated out. “[Marginal likelihood estimation for hierarchical models](#)” discusses some existing approaches for addressing this challenge.

Our article proposes an importance sampling (IS) approach for estimating the marginal likelihood when the

likelihood is intractable but can be estimated unbiasedly. The method is implemented by first running a sampling scheme such as MCMC, whose draws are used to form an importance distribution for the fixed parameters. We then estimate the marginal likelihood using an unbiased estimate of the likelihood combined with this importance distribution. We call this approach Importance Sampling Squared (IS^2), as it is itself an importance sampling (IS) procedure and we can often estimate the likelihood by IS. We claim that it has several advantages over competing methods for estimating the marginal likelihood. The three most important ones are that a) it produces an unbiased estimate of the marginal likelihood; b) the method is robust against potential issues with the samples used to form the proposals; and c) it provides an easily computable and cheap estimate of the variability of the marginal likelihood estimator. Unbiasedness means that we can approach the true value and easily lower the variance of the estimator by getting more samples. Robustness means that the marginal likelihood estimator is simulation consistent even when the MCMC sampler has not yet converged to the posterior distribution of interest. This can mean that the MCMC may have a slightly perturbed posterior as a target or that the MCMC has the posterior as a target but it may not as yet have converged. The reason that (IS^2) works robustly is that it uses importance weights to correct for bad importance samples. Estimating the variability of the estimator directly means that we do not have to replicate the algorithm to obtain reliable estimates of the standard error of the estimator. These observations in turn suggest that it may be desirable to run the initial MCMC in parallel, and without rigorously requiring each MCMC to converge before forming the proposals for IS. Sections “[Marginal likelihood estimation for hierarchical models](#)”, “[General discussion and future work](#)” and Appendix A of the appendix expand on these points.

The rest of the paper is organized as follows. “[Marginal likelihood estimation for hierarchical models](#)” introduces the general model under consideration, reviews several competing approaches for estimating the marginal likelihood, and briefly compares them to our approach. “[Estimating the marginal likelihood by \$IS^2\$](#) ” gives a detailed description of the IS^2 method. “[Applying the \$IS^2\$ method to data](#)” illustrates the IS^2 method in two applications, which show that the method gives accurate estimates of the marginal likelihood and its standard error in reasonable computing time. “[General discussion and future work](#)” concludes and discusses future work on speeding up the computation. There are five appendices. Appendix A outlines how our results on estimating the marginal likelihood extend to estimating posterior expectations. Appendix B studies the effect

on IS of using an estimated likelihood. Appendix C provides further estimation details for the two applications in “Applying the IS² method to data” and Appendix D provides additional applications of the method. Finally, Appendix E contains the proofs of all the results in the paper.

Marginal likelihood estimation for hierarchical models

This section defines the class of hierarchical models for which the density function of the observations (individual participant data), conditional on the group-level parameters θ and a vector of individual level parameters α (random effects) is available, but the likelihood is intractable because it is an integral that cannot be computed over the individual random effects. We also briefly review competing approaches.

Let y_j be the vector of observations (responses) for the j^{th} subject and define $y = y_{1:S} = (y_1, \dots, y_S)$ as the vector of observations for all S subjects. Let $\alpha_j \in \mathbb{R}^{d_\alpha}$, where $\mathbb{R}^d$ denotes d -dimensional Euclidean space, be the vector of random effects (subject-level parameters) for subject j and define $p(\alpha_j|\theta)$ as its density. We define $\alpha = \alpha_{1:S} = (\alpha_1, \dots, \alpha_S)$ as the vector of all random effects in the model (i.e., all subject-level parameters). Let $\theta \in \mathbb{R}^{d_\theta}$ be the vector of unknown group-level parameters and let $p(\theta)$ be the prior for θ . Assuming that the $y_j, \alpha_j, j = 1, \dots, S$ are independent given θ , the joint density of the random effects and the observations is

$$p(y, \alpha|\theta) = \prod_{j=1}^S p(y_j, \alpha_j|\theta) = \prod_{j=1}^S p(y_j|\alpha_j, \theta) p(\alpha_j|\theta). \quad (1)$$

We assume the densities $p(y|\alpha, \theta)$, $p(\alpha|\theta)$, and $p(\theta)$ are available analytically or numerically (e.g., by numerical integration). However, even when these densities are separately available, the likelihood of the hierarchical model is typically analytically intractable because it is an integral over the random effects:

$$p(y|\theta) = \prod_{j=1}^S p(y_j|\theta) \quad \text{with} \quad p(y_j|\theta) = \int p(y_j|\alpha_j, \theta) p(\alpha_j|\theta) d\alpha_j. \quad (2)$$

By Bayes' rule, we can express the joint posterior density of θ and α as

$$\pi(\theta, \alpha) := p(y|\alpha, \theta) p(\alpha|\theta) p(\theta) / p(y), \quad (3)$$

where

$$p(y) = \int \int p(y|\alpha, \theta) p(\alpha|\theta) p(\theta) d\alpha d\theta \quad (4)$$

is the marginal likelihood.

The main goal of the article is to develop an efficient method for estimating the marginal likelihood, given samples from the posterior distribution over the group-level parameters and individual-level random effects. These samples will have been obtained through generic sampling routines (e.g., JAGS, STAN) or custom (e.g., differential evolution-MCMC (DE-MCMC) in Turner, Sederberg, Brown, & Steyvers, 2013). Even custom samplers still suffer from the problems of high autocorrelation between successive iterates of samples from the posterior, and of slow or uncertain convergence in some problems; for example, when the vector of random effects is high dimensional. These problems can lead to unreliable or biased posterior samples, in which case model selection metrics derived from those posterior samples can also be dramatically wrong. This issue has the potential to influence the reliability of some modern developments in estimating the marginal likelihood for cognitive models, including bridge sampling (Gronau et al., 2017; Gronau, Wagenmakers, Heck, & Matzke, 2019) and thermodynamic integration (Evans & Annis, 2019). Both of those approaches have the same structure: they begin with posterior samples, and post-process these samples to estimate the marginal likelihood. Both bridge sampling and thermodynamic integration are therefore sensitive to the accuracy of the sampling method – when the posterior samples are biased in some way, the estimated marginal likelihood will also be biased. An alternative approach is to bypass posterior samples altogether. For example, Evans and Brown (2018) proposed estimating the marginal likelihood directly by Monte Carlo integration over the prior. This method works well in small problems, but it can be highly inefficient. This “direct” method quickly becomes computationally infeasible in high-dimensional problems, or with hierarchically-specified models.

Our paper proposes a new approach. The basic idea is that we first obtain samples of the model parameters (θ ; group level), but not the random effects (α ; individual participants). These samples are used to form efficient proposal densities for an importance sampling algorithm on the model parameters. The IS² method therefore does not suffer from the same potential drawbacks as the bridge sampling and thermodynamic integration methods, because the samples are only used to form the proposals for θ . This means that unreliable samples of θ may lead to inefficient proposal densities, which can decrease the efficiency of the IS² method, but will not lead to bias in the marginal likelihood. In fact, the IS² method makes it possible to obtain unbiased and simulation consistent estimators of the marginal likelihood without ensuring that the underlying

Markov chain(s) have in fact converged, although they need to produce reasonable estimates of the posterior. Section “[General discussion and future work](#)” discusses the robustness property of the IS² estimator and how to use it to speed up the estimation of the marginal likelihood and posterior moments.

Estimating the marginal likelihood by IS²

The likelihood of hierarchical models of the form in Eq. 2 can be analytically intractable, though it can be estimated unbiasedly using importance sampling (other hierarchical models such as state space models may require a particle filter). Let $\{m_j(\alpha_j|\theta, y_j); j = 1, \dots, S\}$ be a family of proposal densities that we use to approximate the conditional posterior densities $\{\pi(\alpha_j|\theta); j = 1, \dots, S\}$. Let $g_{IS}(\theta)$ be a proposal density for the group-level parameters, θ . We note that many different methods can be used to obtain efficient and reliable proposals $g_{IS}(\theta)$ for the group-level parameters and $m_j(\alpha_j|\theta, y_j)$ for the random effects α_j , for each subject $j = 1, \dots, S$. A simple frequently used approach takes the proposal as a multivariate Student t distribution with a small number of degrees of freedom, whose location is a mode of the log-likelihood and whose scale matrix is the inverse of minus the Hessian matrix at this mode. Our article constructs the proposal distribution by first running MCMC, even just a short MCMC, to obtain samples from the posterior distribution. We then construct a proposal distribution in the IS² procedure by fitting a mixture of normal or Student t distributions to these samples, as outlined in “[Applying the IS² method to data](#)” and Appendix C.

The density $p(y_j|\theta)$ is estimated unbiasedly by

$$\hat{p}(y_j|\theta) = \frac{1}{N} \sum_{i=1}^N w(\alpha_j^{(i)}, \theta), \quad \text{where } w(\alpha_j^{(i)}, \theta) = \frac{p(y_j|\alpha_j^{(i)}, \theta) p(\alpha_j^{(i)}|\theta)}{m_j(\alpha_j^{(i)}|\theta, y_j)}, \quad (5)$$

where $\alpha_j^{(i)} \stackrel{iid}{\sim} m_j(\alpha_j|\theta, y_j)$ and N is the number of importance samples used in the likelihood estimate, which we will refer to as the number of “particles”. We note that for importance sampling to be effective, it is necessary for the support of $m_j(\alpha_j|\theta, y_j)$ to contain the support of its corresponding conditional density (i.e., the tails of the proposal should be fatter than the tails of the target). This condition for importance sampling is usually easily satisfied using the defensive sampling approach in Hesterberg (1995). Appendix C discusses the defensive

sampling approach and the proposal $m_j(\alpha_j|\theta, y_j)$. Hence,

$$\hat{p}_N(y|\theta) = \prod_{j=1}^S \hat{p}(y_j|\theta) \quad (6)$$

is an unbiased estimator of the likelihood $p(y|\theta)$. To choose an optimal number of particles N , discussed in Appendix B1, we use the delta method to estimate the variance

$$\mathbb{V}(\log \hat{p}_N(y|\theta)) = \sum_{j=1}^S \mathbb{V}(\log \hat{p}(y_j|\theta)) \quad (7)$$

of $\log \hat{p}_N(y|\theta)$. This gives the estimate

$$\hat{\mathbb{V}}(\log \hat{p}(y_j|\theta)) \approx \frac{\sum_{i=1}^N w(\alpha_j^{(i)}, \theta)^2}{\left(\sum_{i=1}^N w(\alpha_j^{(i)}, \theta)\right)^2} - \frac{1}{N}.$$

We now rewrite Eq. 4 to show how it can be estimated by IS. Let u consist of all random variables used to construct $\hat{p}_N(y|\theta)$ for a given value of θ , with $p_N(u|\theta)$ the density of u . In practice, as will be clear shortly, it is unnecessary to know $p_N(u|\theta)$. We also equivalently write $\hat{p}_N(y|\theta)$ as $\hat{p}_N(y|\theta, u)$. Then $\hat{p}_N(y|\theta)$ is unbiased if

$$p(y|\theta) = \int_U \hat{p}_N(y|\theta, u) p_N(u|\theta) du.$$

Let $g_{IS}(\theta)$ be the proposal for θ obtained by MCMC or otherwise. Then, the marginal likelihood is given by

$$\begin{aligned} p(y) &= \int_{\Theta} \int_U \hat{p}_N(y|\theta, u) p_N(u|\theta) p(\theta) du d\theta, \\ &= \int_{\Theta} \int_U \frac{\hat{p}_N(y|\theta, u) p_N(u|\theta) p(\theta)}{g_{IS}(\theta) p_N(u|\theta)} g_{IS}(\theta) p_N(u|\theta) du d\theta, \\ &= \int_{\Theta} \int_U w(\theta, u) g_{IS}(\theta) p_N(u|\theta) du d\theta, \quad \text{where} \\ w(\theta, u) &= \frac{\hat{p}_N(y|\theta, u) p(\theta)}{g_{IS}(\theta)}. \end{aligned} \quad (8)$$

This leads to the IS² estimator,

$$\hat{p}_{IS^2}(y) = \frac{1}{M} \sum_{i=1}^M \tilde{w}(\theta_i) \quad \text{where } \tilde{w}(\theta_i) = w(\theta_i, u_i), \quad i = 1, \dots, M. \quad (9)$$

Algorithm 1 outlines the IS² algorithm for obtaining an estimate of the marginal likelihood. On a practical note, the algorithm is well-suited to efficient parallelization, by processing particles (importance samples) independently in steps (1) and (2).

Algorithm 1 IS² algorithm for estimating the marginal likelihood

1. Generate $\theta_i \stackrel{iid}{\sim} g_{IS}(\theta)$ and compute the likelihood estimate $\hat{p}_N(y|\theta_i)$, $i = 1, \dots, M$ via Equations (5) and (6).
2. Compute the weights

$$\tilde{w}(\theta_i) = \frac{\hat{p}_N(y|\theta_i) p(\theta_i)}{g_{IS}(\theta_i)}, \quad i = 1, \dots, M \quad (10)$$

3. The IS² estimator of the marginal likelihood $p(y)$

$$\hat{p}_{IS^2}(y) = \frac{1}{M} \sum_{i=1}^M \tilde{w}(\theta_i). \quad (11)$$

The usual way to estimate the standard error of the marginal likelihood estimator is by replication, i.e., by estimating the marginal likelihood estimator a number of times and then computing the standard deviation of these estimates. However, reliable estimation by replication can take a long time. An important advantage of our method is that it is straightforward to estimate the standard error of the marginal likelihood estimator $\hat{p}_{IS^2}(y)$. This affords researchers greater confidence in application of the estimated marginal likelihood, and also permits a simpler investigation of the tradeoff between efficiency and bias. The variance estimator is

$$\hat{\mathbb{V}}(\hat{p}_{IS^2}(y)) = \frac{1}{M} \hat{\sigma}_{p_{IS^2}(y)}^2, \quad \text{where } \hat{\sigma}_{p_{IS^2}(y)}^2 = \frac{1}{M} \sum_{i=1}^M (\tilde{w}(\theta_i) - \hat{p}_{IS^2}(y))^2. \quad (12)$$

It is clear that under mild conditions the IS² estimator is unbiased, simulation consistent and tends to normality as it is an IS estimator. Theorem 1 formally states these results and Appendix E gives its proof.

Theorem 1 *Let M be the number of samples for θ and N the number of particles for estimating the likelihood. Under the assumptions in Theorem A1 in Section A.1 of the Appendix.*

- (i) $\mathbb{E}(\hat{p}_{IS^2}(y)) = p(y)$ and $\hat{p}_{IS^2}(y) \xrightarrow{a.s.} p(y)$ as $M \rightarrow \infty$ for any $N \geq 1$.
- (ii) $\sqrt{M}(\hat{p}_{IS^2}(y) - p(y)) \xrightarrow{d} N(0, \sigma_{p_{IS^2}(y)}^2)$ and $\hat{\sigma}_{p_{IS^2}(y)}^2 \xrightarrow{a.s.} \sigma_{p_{IS^2}(y)}^2 = \mathbb{V}(\tilde{w}(\theta))$ as $M \rightarrow \infty$ for any $N \geq 1$.

There is a trade-off between the statistical precision of the IS² estimator (in terms of the variance from Eq. 12) and the computational cost of obtaining the estimate. A large number of particles N gives a more accurate estimate of the likelihood with greater computational cost, while a small N results in a likelihood estimator $\hat{p}_N(y|\theta)$ that is cheap to evaluate but has larger variance. Appendix B

provides theoretical and practical guidelines of this tradeoff. Proposition A1 in Appendix B.2 shows that the optimal N in the tradeoff between accuracy and computational cost is such that the variance of the log-likelihood estimator $\mathbb{V}(\log \hat{p}(y|\theta))$ is approximately 1.

Applying the IS² method to data

This section describes two applications of the IS² method to real data. The applications are chosen from quite different domains – the first focuses on choice in a health context, the second focuses on an experimental manipulation in cognition. We chose these applications as they are challenging cases for estimating marginal likelihoods, and thus represent good test cases for illustrating the power of our method. They also emphasize that the IS² method is not restricted to a particular domain of application or model class. It is a general method for estimating the marginal likelihood and the standard error of that estimate and can be applied in a variety of other contexts where quantitative models are estimated from data in a hierarchical framework including the study of attention, decision making, categorization, and memory.

Individual differences in health-related decisions

The first application explores individual differences in preferences for health decision making, by examining choices for hypothetical appointments with a health practitioner. A common approach to understanding individual differences in applied contexts is to assume that different individuals have different *utility* parameters, where the utility represents the subjective value placed on the *attributes* that describe the products and services on offer. In this sense the utility is a quantitative estimate of what people like and what they dislike. A common way to model individual differences in choices is to then assume individual differences in *utility* parameters, commonly described as “taste heterogeneity”. The mixed logit model is the most common framework for modeling individual differences in utilities (Train, 2009), which assumes utility parameters are multivariate normal distributed in the population. Nevertheless, more recent research suggests that all individuals might hold the same (or have sufficiently similar) utilities except that those utilities can be scaled up or down for different individuals (Fiebig, Keane, Louviere, & Wasi, 2010). This approach assumes that the ratio of the utilities for pairs of items is constant across participants though the absolute difference between those utilities can differ. This “scale heterogeneity” means that some people make more random choices than others, though their latent preferences, in the limit of very many trials, are similar.

The key model comparison question in this context is whether the choices observed in a population of participants are more consistent with individual differences in *latent preferences* (taste) or individual differences in *choice consistency* (scale). We tested this idea for the data reported in Fiebig et al. (2010) in a study where $S = 79$ women made choices about $T = 32$ choice scenarios for a pap smear test, where each scenario was defined by a combination of the values of the $K = 5$ choice attributes listed in Table 1 (not all unique combinations were presented). For example, one choice scenario might be that the pap smear would be conducted by a female doctor who is unknown to the patient, the test is not due though the doctor recommends it, with a cost of \$20. The participant is then asked whether they would take this test or not.

Our goal is to test whether choices in the pap smear test data provide greater support for the presence of individual differences in latent preferences, or the presence of individual differences in latent preferences *and* choice consistency. For this goal, we estimated different models which instantiated the different hypotheses. We then used the IS² method to estimate the marginal likelihoods of those models, and used those estimates for model selection via Bayes factors.

Behavioural model

We analyse the choice data with the generalized multinomial logit (GMNL) model (Fiebig et al., 2010), a generalization of the mixed logit model, that accounts for individual differences in both latent preferences and choice consistency. We let the observed choice for participant j in scenario i be $y_{ij} = 1$ if the participant chooses to take the test and $y_{ij} = 0$ otherwise. The general form of the GMNL for the probability of individual j on trial i selecting the test is

$$p(y_{ij}=1|X_{ij}, \beta^{(j)}) = \frac{\exp(\beta_{0j} + \sum_{k=1}^K \beta_{kj}x_{ijk})}{1 + \exp(\beta_{0j} + \sum_{k=1}^K \beta_{kj}x_{ijk})}, \quad (13)$$

where $\beta^{(j)} = (\beta_{0j}, \beta_{1j}, \dots, \beta_{Kj})'$ and $X_{ij} = (x_{ij1}, \dots, x_{ijK})'$ are the vectors of utility weights and choice attributes respectively. As is standard in most consumer decision research, the GMNL model assumes that the attributes comprising an option (i.e., a hypothetical appointment) each have an associated utility, and the utility of an option as a whole is the sum of the component attribute utilities. The summed utility probabilistically generates a choice via the Luce choice rule, as shown in Eq. 13. The model also allows for an *a priori* bias toward accepting or rejecting the test at each trial independent of the feature values of the test, known as an alternative specific constant; $\beta_{0j} = \beta_0 + \eta_{0j}$, with $\eta_{0j} \sim \mathcal{N}(0, \sigma_0^2)$.

The utilities for individual participants are modeled as

$$\begin{aligned} \beta_{jk} &= \lambda_j \beta_k + \gamma \eta_{jk} + (1 - \gamma) \lambda_j \eta_{jk}, \\ \lambda_j &= \exp(-\delta^2/2 + \delta \zeta_j), \\ k &= 1, \dots, K, \end{aligned}$$

with $\eta_{jk} \sim \mathcal{N}(0, \sigma_k^2)$ and $\zeta_j \sim \mathcal{N}(0, 1)$. The expected value of the scaling coefficients λ_j is one, implying that $\mathbb{E}(\beta_{jk}) = \beta_k$. We set the across-participant variance parameter for the cost attribute to zero ($\sigma_5^2 = 0$) as we found no evidence of heterogeneity across individuals for this attribute, beyond the scaling effect. As the sum of the probabilities over the possible choice options in Eq. 13 is 1 (i.e., take the test, or not), to make the model identified we set the coefficients β with respect to $y_{ij} = 0$ (not taking the test) to zero.

The GMNL model accounts for individual differences in latent preferences (taste) and choice consistency (scale) through the parameters δ and γ , respectively. When $\delta = 0$ (so that $\lambda_j = 1$ for all individuals), the GMNL model reduces to the mixed logit model. The mixed logit model captures heterogeneity in consumer preferences by allowing individuals to weight the choice attributes differently (i.e., individual differences in attribute utilities). By introducing taste heterogeneity, the mixed logit model avoids the restrictive independence of irrelevant alternatives property of the standard multinomial logit model.

Table 1 Choice attributes for the pap smear data set and their associated parameters

Choice attributes	Values	Parameters
Alternative specific constant for “take the test”	1	β_0, σ_0
Whether patient knows doctor	0 (no), 1 (yes)	β_1, σ_1
Whether doctor is male	0 (no), 1 (yes)	β_2, σ_2
Whether test is due	0 (no), 1 (yes)	β_3, σ_3
Whether doctor recommends test	0 (no), 1 (yes)	β_4, σ_4
Test cost	{0, 10, 20, 30}	β_5

The GMNL model additionally allows for differing choice consistency across individuals, known as scale heterogeneity, through the random variable λ_j . This variable changes all attribute weights simultaneously, allowing the model to predict choices to be more random for some consumers than others, such that smaller values of δ (so that $\lambda_j \rightarrow 1$) indicate differences in latent preferences where larger values of δ (so that $\lambda_j \rightarrow 0$) indicate differences in choice consistency. The γ parameter weights the specification between two alternative ways of introducing heterogeneity into the model.

The parameter vector is $\theta = (\beta_0, \beta_1, \dots, \beta_K, \sigma_0^2, \sigma_1^2, \dots, \sigma_K^2, \delta^2, \gamma)'$, while the vector of random effects for each individual is $\alpha_j = (\eta_{0j}, \dots, \eta_{jK}, \lambda_j)$. The likelihood is therefore

$$p(y|\theta) = \prod_{j=1}^S p(y_j|\theta) = \prod_{j=1}^S \left[\int \left(\prod_{i=1}^T p(y_{ij}|\alpha_j, \theta) \right) p(\alpha_j|\theta) d\alpha_j \right], \quad (14)$$

where y_{ij} is the observed choice, $y = (y_{11}, \dots, y_{1T}, \dots, y_{j1}, \dots, y_{jT})'$ and $p(y_{ij}|\alpha_j, \theta)$ is given by the choice probability in Eq. 13. We use diffuse priors specified as $\beta_0 \sim \mathcal{N}(0, 100)$, $\sigma_0 \propto (1 + \sigma_0^2)^{-1}$, $\beta_k \sim \mathcal{N}(0, 100)$, $\sigma_k \propto (1 + \sigma_k^2)^{-1}$, for $k = 1, \dots, K$, $\delta \propto (1 + \delta/0.2)^{-1}$, and $\gamma \sim \text{U}(0, 1)$. The standard deviation parameters have half-Cauchy priors (Gelman, 2006). Appendix C provides complete details of the estimation procedure.

Results

We estimated the posterior distribution using $M = 50,000$ importance samples for the parameters, and estimated the Monte Carlo standard errors by bootstrapping the importance samples because we found that bootstrapping gave the most stable results of the estimator. Appendix C gives implementation details. The log marginal likelihood (with standard error) for the mixed logit and GMNL models were $-981.24 (.003)$ and $-978.27 (.012)$, respectively. The Monte Carlo standard errors are very small, indicating the IS² method is highly efficient. The estimates of the log marginal likelihood allow us to calculate a Bayes factor of approximately 20 for the GMNL model over the mixed logit model; $BF_{\text{GMNL:mixed}} = \exp(-978.27 - (-981.24)) \approx 19.5$.

The larger marginal likelihood for the GMNL model than the mixed logit model provides some evidence for the presence of scale heterogeneity for this data set. Corroborating evidence comes from the parameter estimates. The posterior mean for γ was .17 (90% credible interval [.014–.452]), where the weight for scale heterogeneity in the GMNL model was $1 - \gamma$. This means that assuming individuals only differed in their latent

preferences (taste heterogeneity) did not sufficiently capture the trends in data. Some participants also made more variable choices than others (scale heterogeneity).

The speed-accuracy tradeoff in perceptual decisions

Our second application explores the cognitive processes involved in perceptual decisions. We examine decision making in the context of the well-studied speed-accuracy tradeoff: the finding that electing to make faster decisions decreases decision accuracy and, conversely, that electing to increase decision accuracy causes decisions to become slower (see Heitz, 2014 for a review). The speed-accuracy tradeoff is typically attributed to changes in caution – the quantity of evidence required to trigger a response (i.e., the response threshold). Here, we apply the IS² method to a class of decision making models in the Linear Ballistic Accumulator (LBA; Brown & Heathcote, 2008) framework to confirm that model choice via the marginal likelihood is consistent with existing methods of model choice in the literature.

We draw upon data reported by Forstmann, Dutilh, Brown, Neumann, and von Cramon (2008) coming from an experiment in which $S = 19$ participants were required to make difficult motion discrimination decisions under varying degrees of time pressure. The participants made repeated decisions about whether a cloud of semi-randomly moving dots appeared to move to the left or to the right of a display. Participants made these decisions under three conditions that differed in their task instructions: they were asked to respond as accurately as possible (condition 1: “accuracy emphasis”), at their own pace (condition 2: “neutral emphasis”), or as quickly as possible (condition 3: “speed emphasis”). The three conditions were randomized across trials. A visual cue described the decision emphasis on a trial-by-trial basis prior to stimulus onset. Each subject made 280 decisions in each condition for a total of $T = 840$ trials. See Forstmann et al. (2008) for all remaining details.

Freely available code applying the IS² method to one of the hierarchical LBA models described here is available at osf.io/xv59c. Appendix D also provides two additional applications of the IS² method to the hierarchical LBA: one application to the speed-accuracy tradeoff in lexical decisions, and a second application to response bias in lexical decisions. Code for these two examples is also available at the same repository.

Behavioural model

We analyse the choice and response time data with a hierarchical LBA model. In a typical perceptual decision making experiment, the i^{th} observation for the j^{th} participant contains two pieces of information. The first is

the response choice, denoted by $RE_{ij} \in \{1, \dots, C\}$, where C is the number of response alternatives. The second is the response time (RT), denoted by $RT_{ij} \in (0, \infty)$. Brown and Heathcote (2008) derived the joint density and cumulative distribution function of the finishing time of an LBA accumulator over response choice and response time. The finishing time distribution for each of the LBA accumulators is determined by five parameters: the mean and standard deviation of the drift rate distribution (v, s), the threshold parameter b , the width of the start point distribution A , and the non-decision time τ . It is common to set the standard deviation s of the drift rate distribution to one, i.e. $s = 1$ (see also Donkin, Brown, & Heathcote, 2009). The probability density of the c^{th} accumulator reaching the threshold at time t , and the other accumulators not reaching the threshold prior to time t , is

$$\text{LBA}(RE_{ij} = c, RT_{ij} = t | b, A, v, s, \tau) = f_c(t) \prod_{k \neq c} (1 - F_k(t)),$$

where f_c and F_c are the density and distribution functions for the finishing times of a single LBA accumulator (for details see Brown & Heathcote, 2008; Terry et al., 2015). In a two-choice experiment where response bias is not expected, it is convenient to define the means of the drift rate distributions as $v = (v^{(e)}, v^{(c)})$, where $v^{(e)}$ is the mean of the drift rate for the accumulator corresponding to the incorrect response and $v^{(c)}$ is the mean of the drift rate for the accumulator corresponding to the correct response. Together, these assumptions imply that each subject j , $j = 1, \dots, S$, has the vector of random effects

$$\{b_j, A_j, v_j^{(e)}, v_j^{(c)}, \tau_j\}.$$

With the usual assumption of independence, the conditional density of all the observations is

$$p(RE, RT | b, A, \tau, v) = \prod_{j=1}^S \prod_{i=1}^T \text{LBA}(RE_{ij}, RT_{ij} | b_j, A_j, v_j^{(e)}, v_j^{(c)}, \tau_j).$$

We follow the notation and assumptions in Gunawan, Hawkins, Tran, Kohn, and Brown (2019). For each subject $j = 1, \dots, S$, let the vector of random effects

$$\alpha_j = (\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j}),$$

where $\alpha_{b_j} = \log(b_j)$, $\alpha_{A_j} = \log(A_j)$, $\alpha_{v_j^{(e)}} = \log(v_j^{(e)})$, $\alpha_{v_j^{(c)}} = \log(v_j^{(c)})$ and $\alpha_{\tau_j} = \log(\tau_j)$. We take the following prior densities (as in Gunawan et al., 2019):

α_j is $N(\alpha_j; \mu_\alpha, \Sigma_\alpha)$, $\mu_\alpha \sim N(0, I_{d_\alpha})$ where d_α is the dimension of α_j , and the hierarchical prior for Σ_α is

$$\Sigma_\alpha | a_1, \dots, a_{d_\alpha} \sim IW(v_\alpha + d_\alpha - 1, 2v_\alpha \text{diag}(1/a_1, \dots, 1/a_{d_\alpha})), \quad (15)$$

$$a_1, \dots, a_{d_\alpha} \sim IG\left(\frac{1}{2}, \frac{1}{\mathcal{A}_d^2}\right), \quad d = 1, \dots, d_\alpha, \quad (16)$$

where $v_\alpha, \mathcal{A}_1, \dots, \mathcal{A}_{d_\alpha}$ are positive scalars and $\text{diag}(1/a_1, \dots, 1/a_{d_\alpha})$ is the diagonal matrix with diagonal elements $1/a_1, \dots, 1/a_{d_\alpha}$.¹ This is the marginally non-informative prior of Huang and Wand (2013). Following Gunawan et al. (2019), we set $v_\alpha = 2$ and $\mathcal{A}_d = 1$ for all d . These settings lead to half- t marginal prior distributions for the standard deviations of Σ_α and uniform marginal prior distributions for the off-diagonal correlations implied by Σ_α .

We tested whether the changes in task instructions across the three experimental conditions caused changes in observed decision behaviour that is consistent with changes in decision caution. We compared four LBA models. For consistency with the notation above, we refer to instructions that emphasize accuracy, neutral or speedy decisions as a, n and s , respectively, and the correct and error accumulators as c and e , respectively.

Model I assumes no differences in parameters across the three conditions, i.e., a null model. This model corresponds to the psychological assumption that performance is not influenced by the instructions given to participants. The vector of random effects for subject j is

$$\alpha_j = (\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j}).$$

Model II assumes there are two response threshold parameters. One threshold parameter defines decision-making in the accuracy- and neutral-emphasis conditions, and the second threshold parameter in the speed-emphasis condition. This makes the simplifying assumption that performance in two of the three conditions is so similar as to be effectively identically distributed,

$$\alpha_j = (\alpha_{b_j^{(a/n)}}, \alpha_{b_j^{(s)}}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j}).$$

Model III assumes there are three response threshold parameters, one for each condition. This model assumes that the participant responds to each condition with

¹The notation $IW(a, A)$ means an inverse Wishart distribution with degrees of freedom a and scale matrix A and the notation $IG(a, b)$ means an inverse Gamma distribution with scale parameter a and shape parameter b .

a different level of caution, though otherwise with identically distributed parameters,

$$\alpha_j = \left(\alpha_{b_j^{(a)}}, \alpha_{b_j^{(n)}}, \alpha_{b_j^{(s)}}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j} \right).$$

Model IV assumes there are separate response threshold and non-decision time parameters for each condition. This model assumes that instructions to emphasise speedy, neutral or accurate decisions influence not only cautiousness but also the time required to perceptually encode the stimulus and produce a motor response (non-decision time),

$$\alpha_j = \left(\alpha_{b_j^{(a)}}, \alpha_{b_j^{(n)}}, \alpha_{b_j^{(s)}}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j^{(a)}}, \alpha_{\tau_j^{(n)}}, \alpha_{\tau_j^{(s)}} \right).$$

As in the first application, Appendix C provides complete details of the estimation procedure.

Results

Once we obtained samples from the posterior, we ran the IS² algorithm with $M = 10,000$ draws to estimate the log of the marginal likelihood, $\log \hat{p}(y)$, and adaptively set the number of particles N that were used to estimate the likelihood such that the variance of the log-likelihood estimates did not exceed 1; recall that $\sigma_{opt}^2 = 1$ is the optimal variance of the log-likelihood (cf. “[Estimating the marginal likelihood by IS²](#)”). We initiated this adaptive procedure by starting with $N = 250$ particles and computed the variance of the log-likelihood given in Eq. 7. If the variance was greater than 1, we increased the number of particles. We again estimated the Monte Carlo standard error of the log of the marginal likelihood estimates by bootstrapping the importance samples because we found that bootstrapping gave the most stable results.

The log marginal likelihoods (with standard errors) for Models I, II, III and IV are, respectively, 5204.17 (0.11), 7352.75 (0.06), 7453.73 (0.10) and 7521.44 (0.17), with a total computation time of around 80 minutes. As in the first application, the Monte Carlo standard errors are very small, suggesting that IS² is highly efficient. The marginal likelihoods favor Model IV – allowing a separate response threshold and a separate non-decision time for each instruction condition. This outcome is partially consistent with Forstmann et al. (2008), but also suggests that the non-decision time – the combined time required to perceptually encode a stimulus and produce a motor response – changes with instruction condition.

General discussion and future work

Model comparison is the means by which competing theories are rigorously evaluated and compared, which is fundamental to advancing psychological science. Many quantitative approaches to model comparison in psychology have struggled to appropriately account for model flexibility – the range of data patterns that a model can predict. For example, commonly used methods such as AIC and BIC measure the flexibility of a model simply by counting the number of freely estimated model parameters. This is problematic because not all parameters are equal, in terms of complexity, and models with the same numbers of parameters can be quite different in complexity. In some cases, adding extra parameters to a model can even *decrease* the complexity, for example when a hierarchical distribution structure is added to constrain otherwise-independent models for individual participants. Model selection via the marginal likelihood is one of the few methods that naturally accounts for model flexibility.

Despite its theoretical advantages, the marginal likelihood has seen limited uptake in practice for all but the simplest of psychological models, owing to its prohibitive computational cost. Although many psychologically interesting models have analytic, numerical or rapidly-simulated solutions for the responses of individual participants (e.g., the density of observed choices and response times of individual participants), the likelihood for hierarchical versions of those same models can be analytically intractable because they contain an integral over the individual random effects. This has proved to be the stumbling block in previous attempts to compute the marginal likelihood for hierarchical implementations of psychological models.

Our article develops a method for Bayesian model comparison when the likelihood (of a hierarchical model) is intractable but can be estimated unbiasedly. We show that when the density of individual observations conditioned on group- and individual-level parameters is available, an importance sampling algorithm can provide an unbiased estimate of the marginal likelihood of hierarchical models. We term this approach importance sampling squared (IS²). The IS² method can be applied to samples from the posterior distribution over a model’s parameters obtained following any sampling method. Appendix B studies the convergence properties of the IS² estimator and provides practical guidelines for optimally selecting the number of particles required to estimate the likelihood.

We show how the IS² method can be used to estimate the marginal likelihood in hierarchical models of health decision making and choice response time models. In both applications, the method estimates the marginal likelihood in a principled way, providing conclusions that are consistent with the current literature. In all cases, the

marginal likelihood is estimated with very small Monte-Carlo standard errors. To aid researchers in using the IS^2 method in their own research, we provide scripts implementing the key elements of the method, as applied to data from Forstmann et al. (2008) and Wagenmakers, Ratcliff, Gomez, and McKoon (2008) – see osf.io/xv59c for more details.

The data applications highlight two important properties of the IS^2 method: it is an efficient and unbiased estimator of the marginal likelihood, and it provides a standard error of the estimator. The marginal likelihood is a key quantity in appropriately accounting for model flexibility in quantitative model comparison. The standard error of an estimator is essential for interpreting any model comparison metric, not just marginal likelihoods, as it expresses the variability of the estimated metric. This is equally important in quantitatively comparing psychological models as it is in the interpretation of conventional statistics, despite being routinely overlooked. No researcher would conclude the population means of two groups differ solely on the basis of a difference in their sampled group means – we would demand a measure that expresses the magnitude of the mean difference as a function of its variability, like a t -test. The IS^2 method provides the researcher with the tools to do precisely this with potentially complex and non-linear psychological models: it estimates the magnitude of differences (marginal likelihoods for different models) as a function of the variability of those estimates (the standard errors).

We emphasize that the IS^2 method is general and can be used to estimate the marginal likelihood and its standard error from models where the density of individual observations conditioned on group-level parameters and a vector of individual-level parameters is available. This is quite a general scenario – reaching beyond the models studied here, and decision-making models more generally – that applies to a very large category of psychological models for which hierarchical Bayesian parameter estimation has been used to date. This generality holds great promise for IS^2 as a vehicle for performing model comparison via the marginal likelihood in psychological research.

We also note that the IS^2 method is robust, by which we mean that the MCMC draws used to form the proposals do not have to converge to draws from the exact posterior as they are used in IS^2 to form importance sampling proposals. Hence, as long as they are roughly in the same region as the posterior, the marginal likelihood estimates will be simulation consistent. The same remarks hold if the MCMC targets a slightly perturbed posterior.

For faster computation, we conjecture that future work could speed up the estimation of the marginal likelihood as follows: a) First, run the MCMC procedure in parallel on

K chains without worrying about the precise convergence of each chain. In practice this means we can run each chain for far fewer iterates than would be required for carrying out Bayesian inference if the inference was based only on the output of each chain. b) Second, form a proposal density based on the output of all the chains. c) Third, use IS to get K robust and unbiased estimates of the marginal likelihood from each of K processors. d) Finally, average the K estimates to get a final unbiased estimate of the marginal likelihood whose variance is $1/K$ th the variance of each individual estimator. The robustness and unbiasedness properties of the IS^2 estimator are crucial in ensuring that bias does not dominate the variance as K becomes large.

We note in Appendix A that posterior expectations of functions of the parameters can similarly be estimated with IS^2 , albeit with some very minor bias. We conjecture that such posterior moment estimators will be far more efficient than those obtained from standard MCMC output, but leave to future work such speeding up of the computations of the marginal likelihood and posterior moments.

Finally, we note that the marginal likelihood is sometimes sensitive to the prior, and this might lead to some controversy in using marginal likelihood for model comparison. Addressing this issue is beyond the scope of this paper, and we focus only on how to estimate the marginal likelihood efficiently when the likelihood function is intractable.

Acknowledgments The research of Tran, Gunawan, Kohn and Brown was partially supported by Australian Research Council grant DP180102195. Hawkins was partially supported by Australian Research Council grant DE170100177. Tran and Kohn thank Michael Pitt for some useful discussions that led to the development of Proposition A1.

Open Practices Statement The two applications cover previously published data sets (Fiebig et al., 2010; Forstmann, Dutilh, Brown, Neumann, & von Cramon, 2008). Data and code are available at osf.io/xv59c.

Appendix A: Using the IS^2 method to obtain posterior expectations

The motivation for using IS^2 so far in this article was to develop an efficient and robust estimator for the marginal likelihood and the standard error of the estimator. This section shows that it is also straightforward to use IS^2 to robustly estimate expectations with respect to the posterior distribution, as well as the standard errors of these estimators. We also discuss the convergence properties of these estimators.

Using the same notation as in “[Marginal likelihood estimation for hierarchical models](#)”, the posterior expectation of the function φ of θ is

$$\mathbb{E}_\pi(\varphi) = \int_{\Theta} \varphi(\theta) \pi(\theta) d\theta = \frac{\int_{\Theta} \varphi(\theta) p(y|\theta) p(\theta) d\theta}{\int_{\Theta} p(y|\theta) p(\theta) d\theta}. \quad (17)$$

When the likelihood is intractable, but can be estimated unbiasedly, we can use IS^2 to estimate unbiasedly and robustly both the numerator and denominator on the right side of Eq. 17, similarly to “[Estimating the marginal likelihood by \$\text{IS}^2\$](#) ”, to obtain

$$\hat{\mathbb{E}}_\pi(\varphi) = \frac{\frac{1}{M} \sum_{i=1}^M \varphi(\theta_i) \tilde{w}(\theta_i)}{\frac{1}{M} \sum_{i=1}^M \tilde{w}(\theta_i)}, \text{ with weights } \tilde{w}(\theta_i) = \frac{p(\theta_i) \hat{p}_N(y|\theta_i)}{g_{\text{IS}}(\theta_i)}, \quad \theta_i \stackrel{iid}{\sim} g_{\text{IS}}(\theta). \quad (18)$$

In Eq. 18, $\hat{p}_N(y|\theta)$ is the unbiased estimate of $p(y|\theta)$, with N the number of samples or particles used to estimate the likelihood. It is clear that under mild conditions both the numerator and denominator of Eq. 18 will converge to the numerator and denominator respectively of Eq. 17 as $M \rightarrow \infty$, and hence $\hat{\mathbb{E}}_\pi(\varphi)$ converges to $\mathbb{E}_\pi(\varphi)$. It is also clear that under very mild conditions the numerator of Eq. 18 becomes normally distributed as $M \rightarrow \infty$, which then means that $\hat{\mathbb{E}}_\pi(\varphi)$ also converges to normality. These issues are discussed more rigorously below.

We note that while both the numerator and denominator in Eq. 17 are estimated unbiasedly, $\hat{\mathbb{E}}_\pi(\varphi)$ is a biased estimator of $\mathbb{E}_\pi(\varphi)$, although the bias goes to 0 as $M \rightarrow \infty$.

In practice, the variances of both the numerator and denominator in Eq. 18, as well as their covariance, can be estimated by the bootstrap and hence both the variance and bias of $\hat{\mathbb{E}}_\pi(\varphi)$ can be estimated.

A.1 Some technical results

We now give some large sample (in M) convergence results for $\hat{\mathbb{E}}_\pi(\varphi)$ assuming that,

Assumption A1 $\mathbb{E}[\hat{p}_N(y|\theta)] = p(y|\theta)$ for every $\theta \in \Theta$, where the expectation is with respect to the random variables generated in the process of estimating the likelihood.

It is useful in the rest of this section and Appendix E to follow (Pitt, Silva, Giordani, & Kohn, 2012) and write $\hat{p}_N(y|\theta)$ as $p(y|\theta)e^z$, where $z := \log \hat{p}_N(y|\theta) - \log p(y|\theta)$ is a scalar random variable whose distribution conditional on θ is governed by the randomness due to estimating the likelihood $p(y|\theta)$. Thus, the scalar z replaces the multivariate u in “[Estimating the marginal likelihood by \$\text{IS}^2\$](#) ”. Let $g_N(z|\theta)$ be the density of z conditional on θ . Assumption A1 implies that $\mathbb{E}(e^z|\theta) = \int_{\mathbb{R}} e^z g_N(z|\theta) dz = 1$.

Theorem A1 Suppose that Assumption A1 holds, $\mathbb{E}_\pi(|\varphi|)$ exists and is finite, and $\text{Supp}(\pi) \subseteq \text{Supp}(g_{\text{IS}})$, where $\text{Supp}(\pi)$ denotes the support of the distribution π .

- (i) For any $N \geq 1$, $\hat{\mathbb{E}}_\pi(\varphi) \xrightarrow{a.s.} \mathbb{E}_\pi(\varphi)$ as $M \rightarrow \infty$.
- (ii) If

$$\int h(\theta)^2 \left(\frac{\pi(\theta)}{g_{\text{IS}}(\theta)} \right)^2 \left(\int \exp(2z) g_N(z|\theta) dz \right) g_{\text{IS}}(\theta) d\theta \quad (19)$$

is finite for $h(\theta) = \varphi(\theta)$ and $h(\theta) = 1$ for all N , then

$$\sqrt{M} \left(\hat{\mathbb{E}}_\pi(\varphi) - \mathbb{E}_\pi(\varphi) \right) \xrightarrow{d} \mathcal{N} \left(0, \sigma_{\text{IS}^2}^2(\varphi) \right), \text{ as } M \rightarrow \infty, \quad (20)$$

where the asymptotic variance in M for fixed N is given by

$$\sigma_{\text{IS}^2}^2(\varphi) = \mathbb{E}_\pi \left\{ \left(\varphi(\theta) - \mathbb{E}_\pi(\varphi) \right)^2 \frac{\pi(\theta)}{g_{\text{IS}}(\theta)} \mathbb{E}_{g_N}[\exp(2z)] \right\}. \quad (21)$$

- (iii) Define

$$\hat{\sigma}_{\text{IS}^2}^2(\varphi) := \frac{M \sum_{i=1}^M \left(\varphi(\theta_i) - \hat{\mathbb{E}}_\pi(\varphi) \right)^2 \tilde{w}(\theta_i)^2}{\left(\sum_{i=1}^M \tilde{w}(\theta_i) \right)^2}. \quad (22)$$

If the conditions in (ii) hold, then $\hat{\sigma}_{\text{IS}^2}^2(\varphi) \xrightarrow{a.s.} \sigma_{\text{IS}^2}^2(\varphi)$ as $M \rightarrow \infty$, for given N . The proof is in Appendix E.

Although both $\sigma_{\text{IS}^2}^2(\varphi)$ and $\hat{\sigma}_{\text{IS}^2}^2(\varphi)$ depend on N , we do not show this dependence explicitly to simplify the notation. Here, all the probabilistic statements, such as the almost sure convergence, must be understood on the extended probability space that takes into account the extra randomness occurring when estimating the likelihood. When the likelihood can be computed, the analogous results are well known in the literature (see, e.g., p.1318, Geweke, 1989).

Appendix B: The effect on importance sampling of estimating the likelihood

The results in the previous section show that it is straightforward to use importance sampling even when the likelihood is intractable but unbiasedly estimated. This section addresses the question of how much asymptotic efficiency is lost when working with an estimated likelihood. We follow Pitt et al. (2012) and Doucet, Pitt, Deligiannidis, and Kohn (2015) and make the following idealized assumption

to make it possible to develop some intuition and practical guidelines for selecting N . Appendix E contains all the proofs unless stated otherwise.

Assumption A2 (i) *There exists a function $\gamma^2(\theta)$ such that the density $g_N(z|\theta)$ of z is $\mathcal{N}(-\frac{\gamma^2(\theta)}{2N}, \frac{\gamma^2(\theta)}{N})$, where $\mathcal{N}(a, b^2)$ is a univariate normal density with mean a and variance b^2 .*
(ii) *For a given $\sigma^2 > 0$, define $N_{\sigma^2}(\theta) := \gamma^2(\theta)/\sigma^2$. Then, $\mathbb{V}(z|\theta, N = N_{\sigma^2}(\theta)) \equiv \sigma^2$ for all $\theta \in \Theta$.*

If $g_N(z|\theta)$ is Gaussian, then, by Assumption A1, its mean is $-\frac{1}{2}$ times its variance because $\mathbb{E}_{g_N}(\exp(z)) = 1$. Assumption A2 (ii) keeps the variance $\mathbb{V}(z|\theta, N)$ constant across different values of θ , thus making it easy to associate the IS^2 asymptotic variances with σ . Under Assumption A2, the density $g_N(z|\theta)$ depends only on σ and we write it as $g(z|\sigma)$.

Lemma A1 *If Assumption A2 holds for a fixed σ^2 , then Eq. 19 becomes*

$$\int h(\theta)^2 \left(\frac{\pi(\theta)}{g_{\text{IS}}(\theta)} \right)^2 g_{\text{IS}}(\theta) d\theta < \infty \quad (23)$$

for both $h = \varphi$ and $h = 1$.

These are the standard conditions for IS (Geweke, 1989). The proof of this lemma is straightforward and omitted.

Recall that $\sigma_{\text{IS}}^2(\varphi)$ and $\sigma_{\text{IS}^2}^2(\varphi)$ are respectively the asymptotic variances of the IS estimators we would obtain when the likelihood is available and when it is estimated. We refer to the ratio $\sigma_{\text{IS}^2}^2(\varphi)/\sigma_{\text{IS}}^2(\varphi)$ as the inflation factor, which measures how much the asymptotic variance is inflated when working with an estimated likelihood. Theorem A2 obtains an expression for the inflation factor, shows that it is independent of φ , greater than or equal to 1 and increases exponentially with σ^2 , which is the variance of z .

Theorem A2 *Under Assumption A2 and the conditions in Theorem A1,*

$$\frac{\sigma_{\text{IS}^2}^2(\varphi)}{\sigma_{\text{IS}}^2(\varphi)} = \exp(\sigma^2). \quad (24)$$

B.1 Optimal N for estimating posterior expectations

This section aims to explicitly take into account both the statistical precision of the IS^2 estimators and their computational cost. It is apparent that there is a trade-off between these two considerations. A large value of N results in a precise estimator so that σ^2 is small and the relative variance $\mathbb{V}(\hat{\varphi}_{\text{IS}^2})/\mathbb{V}(\hat{\varphi}_{\text{IS}}) = \sigma_{\text{IS}^2}^2(\varphi)/\sigma_{\text{IS}}^2(\varphi) = \exp(\sigma^2)$

given by Eq. 24 is close to 1. However, the cost of such an estimator will be large due to the large number of particles, N , required. Conversely, a small value of N results in an estimator that is cheap to evaluate but has a large value of σ^2 and hence the variance of the IS^2 estimator relative to the IS estimator will be large. To explicitly trade-off these considerations, we introduce the computational time $\text{CT}(\sigma^2)$ which is a product of the relative variance and the computational effort. Minimising $\text{CT}(\sigma^2)$ results in an optimal value for σ^2 and hence N .

Under Assumption A2, $N = N_{\sigma^2}(\theta) = \gamma^2(\theta)/\sigma^2$, so that the expected number of particles, over the draws of θ , is $\bar{N} = \bar{\gamma}^2/\sigma^2$, where $\bar{\gamma}^2 := \mathbb{E}_{g_{\text{IS}}}[\gamma^2(\theta)]$. This motivates the assumption that the computational cost is proportional to $1/\sigma^2$. From Theorems A1 and A2, the variance of the estimator $\hat{\varphi}_{\text{IS}^2}$ based on M importance samples from $g_{\text{IS}}(\theta)$ is approximated by

$$\mathbb{V}(\hat{\varphi}_{\text{IS}^2}) \approx \frac{\sigma_{\text{IS}^2}^2(\varphi)}{M} = \frac{\sigma_{\text{IS}}^2(\varphi)}{M} \exp(\sigma^2). \quad (25)$$

If κ^* is the target precision of the estimator $\hat{\varphi}_{\text{IS}^2}$, then the required number of samples is $M = \sigma_{\text{IS}}^2(\varphi)e^{\sigma^2}/\kappa^*$, and the required computational cost is proportional to

$$\frac{\sigma_{\text{IS}}^2(\varphi)e^{\sigma^2}}{\kappa^*} \times \frac{1}{\sigma^2} = \frac{\sigma_{\text{IS}}^2(\varphi)}{\kappa^*} \times \frac{e^{\sigma^2}}{\sigma^2}.$$

Therefore, it makes sense to define the measure of the computing time of the IS^2 method (in order to obtain a given precision for the IS^2 estimators) as

$$\text{CT}(\sigma^2) := \frac{\exp(\sigma^2)}{\sigma^2}. \quad (26)$$

It is straightforward to check that $\text{CT}(\sigma^2)$ is minimized at

$$\sigma_{\text{opt}}^2 = 1, \quad (27)$$

and the optimal number of particles N is such that $\mathbb{V}(z|\theta, N) = \mathbb{V}(\log \hat{p}_N(y|\theta)) = \sigma_{\text{opt}}^2$.

B.2 Optimal N for estimating the marginal likelihood

Section B.1 derived the optimal N for estimating integrals of the form in Eq. 17. We now show that the optimal N is the same when the main goal is to estimate the marginal likelihood.

The IS^2 marginal likelihood estimator with an estimated likelihood is $\hat{p}_{\text{IS}^2}(y) = \sum_{i=1}^M \tilde{w}(\theta_i)/M$, with the weights $\tilde{w}(\theta_i)$ from Eq. 10. We noted that $\mathbb{E}_{g_{\text{IS}}}[\omega(\theta)] = \mathbb{E}_{\tilde{g}_{\text{IS}}}[\tilde{\omega}(\theta)] = p(y)$, where $\tilde{\omega}(\theta) = e^z \omega(\theta)$. We again wish to compare the variance of the IS^2 estimator to that of the

IS estimator. Under Assumption A2 that z is Gaussian and independent of θ ,

$$\begin{aligned}\mathbb{V}_{\text{GIS}}[\tilde{\omega}(\theta)/p(y)] &= \mathbb{E}_{\text{GIS}}[e^{2z}\omega^2(\theta)/p(y)^2] - 1 = e^{\sigma^2}\mathbb{E}_{\text{GIS}}[\omega^2(\theta)/p(y)^2] - 1 \\ &= e^{\sigma^2}(\mathbb{V}_{\text{GIS}}[\omega(\theta)/p(y)] + 1) - 1.\end{aligned}$$

Consequently, the relative variance of the two schemes is given by

$$\frac{\mathbb{V}_{\text{GIS}}[\hat{p}_{\text{IS}}^2(y)]}{\mathbb{V}_{\text{GIS}}[\hat{p}_{\text{IS}}(y)]} = \frac{\mathbb{V}_{\text{GIS}}[\tilde{\omega}(\theta)/p(y)]}{\mathbb{V}_{\text{GIS}}[\omega(\theta)/p(y)]} = \frac{e^{\sigma^2}(v+1)-1}{v}, \quad (28)$$

where $v = \mathbb{V}_{\text{GIS}}[\omega(\theta)/p(y)] = \mathbb{V}_{\text{GIS}}[\pi(\theta)/g_{\text{IS}}(\theta)]$. Following the argument of Appendix B.1, the corresponding computing time can be defined as

$$\text{CT}_{\text{ML}}(\sigma^2) := \frac{e^{\sigma^2}(v+1)-1}{\sigma^2},$$

where the subscript ML indicates this is for the marginal likelihood estimator.

Let $\sigma_{\min}^2(v)$ minimize $\text{CT}_{\text{ML}}(\sigma^2)$ for a given v . The following proposition, whose proof is obvious and omitted, summarizes some properties of $\sigma_{\min}^2(v)$, where σ_{opt}^2 below is given by Eq. 27.

Proposition A1 (i) For any value of v , $\text{CT}_{\text{ML}}(\sigma^2)$ is a convex function of σ^2 ; therefore $\sigma_{\min}^2(v)$ is unique.
(ii) $\sigma_{\min}^2(v)$ increases as v increases and $\sigma_{\min}^2(v) \rightarrow \sigma_{\text{opt}}^2 = 1$ in Eq. 27 as $v \rightarrow \infty$.

Table 2 illustrates some of the results of Proposition A1 and shows $\sigma_{\min}^2(v)$ and the ratio $\text{CT}_{\text{ML}}(\sigma_{\min}^2)/\text{CT}_{\text{ML}}(\sigma_{\text{opt}}^2)$ for some common values of v used in practice. The table shows that the values of $\sigma_{\min}^2$ and the ratio are insensitive to v , and $\sigma_{\min}^2(v) \rightarrow \sigma_{\text{opt}}^2 = 1$ as v increases. From these observations, we advocate using $\sigma_{\text{opt}}^2 = 1$ as the optimal value of the variance of the log likelihood estimates in estimating the marginal likelihood.

Table 2 Sensitivity of the computing time to $v = \mathbb{V}_{\text{GIS}}[\pi(\theta)/g_{\text{IS}}(\theta)]$ in estimating the marginal likelihood

v	$\sigma_{\min}^2(v)$	$\text{CT}_{\text{ML}}(\sigma_{\min}^2)/\text{CT}_{\text{ML}}(\sigma_{\text{opt}}^2)$
1	0.77	0.97
5	0.93	0.99
10	0.97	1.00
100	1.00	1.00
∞	1.00	1.00

Appendix C: Estimation details for the two applications

C.1 Individual differences in health-related decisions

We estimate the likelihood (Eq. 14 of the main text) by integrating out the random effects vector for each individual separately using different approaches for the mixed logit and GMNL models. For the mixed logit model, we combine the efficient importance sampling (EIS) method of (Richard & Zhang, 2007) with the defensive sampling approach of Hesterberg (1995). The importance density is the two component defensive mixture

$$h(\alpha_j|y_{j1}, \dots, y_{jT}) = \pi h^{\text{EIS}}(\alpha_j|y_{j1}, \dots, y_{jT}) + (1 - \pi)p(\alpha_j),$$

where $h^{\text{EIS}}(\alpha_j|y_{j1}, \dots, y_{jT})$ is a multivariate Gaussian importance density obtained using EIS. Following (Hesterberg, 1995), including the natural sampler $p(\alpha_j)$ in the mixture ensures that the importance weights are bounded. We set the mixture weight as $\pi = 0.5$. For the GMNL model, we follow Fiebig et al. (2010) and use the model density $p(\alpha_j)$ as an importance sampler. We implement this simpler approach for the GMNL model because the occurrence of large values of λ_j causes the defensive mixture estimates of the log-likelihood to be right skewed in this case.

The repeated questioning of each individual (i.e., $T = 32$ choice scenarios) implies that the log-likelihood estimates are sums of independent estimates $\log \hat{p}(y_j|\theta)$ for each individual. To target a certain variance σ^2 for the log-likelihood estimator, we chose the number of particles for each individual (N_j) and parameter combination θ such that $\mathbb{V}(\log \hat{p}(y_j|\theta)) \approx \sigma^2/S$. We implemented this scheme by using a fixed number of initial importance samples and then used the jackknife method to estimate $\gamma_j^2(\theta)$, the asymptotic variance of $\log \hat{p}(y_j|\theta)$, and selected $N_j(\theta) = \hat{\gamma}_j^2(\theta)S/\sigma^2$. The preliminary number of particles is $N = 20$ for the mixed logit model and $N = 2,500$ for the GMNL model.

To construct efficient and reliable proposal densities for the parameters $g_{\text{IS}}(\theta)$, we used the “Mixture of t by Importance Sampling Weighted Expectation Maximization” approach (MitISEM; Hoogerheide, Opschoor, and Van Dijk (2012)). MitISEM constructs a mixture of t densities for approximating the target distribution by minimizing the Kullback–Leibler divergence between the target density and the t mixture, and it can handle target distributions that have non-standard shapes such as multimodality and skewness; MitISEM effectively approximates the posterior of the two models as two component mixtures of multivariate Student’s t distributions. We write $\hat{p}_N(y|\theta) = \hat{p}_N(y|\theta, u)$, with u a fixed random number stream for all θ . The target distribution in the MitISEM is $p(\theta)\hat{p}_N(y|\theta, u)/p(y)$, which

can be considered as the posterior $p(\theta|y, u)$ conditional on y and the common random numbers u . Our procedure is analogous to using common random numbers u to obtain simulated maximum likelihood estimates of θ (see, e.g., Gouriéroux & Monfort 1995), except that we obtain a histogram estimate of the “posterior” $p(\theta|u, y) \propto p(y|\theta, u)p(u)$. This “posterior” is biased but is sufficient to obtain a good proposal density.

We implemented two standard variance reduction methods at each IS stage: stratified mixture sampling and antithetic sampling. The first consists of sampling from each component at the exact proportion of the mixture weights. For example, when estimating the likelihood for the mixed logit model we generate exactly πN draws from the efficient importance density $h^{\text{EIS}}(\alpha_j|y_{j1}, \dots, y_{jT})$ and $(1 - \pi)N$ draws from $p(\alpha_j)$. The antithetic sampling method consists of generating pairs of perfectly negatively correlated draws from each mixture component (see, e.g., Ripley, 1987).

C.2 The speed-accuracy tradeoff in perceptual decisions

We used the Particle Metropolis within Gibbs (PMwG) sampler of Gunawan et al. (2019) to obtain 10,000 posterior draws of θ and $\alpha_{1:S}$ for each of the four models described in the main text. We fitted a mixture of normal distributions to the samples from the posterior of θ to obtain the proposal density for the group-level parameters

$$g_{\text{IS}}(\theta) = \sum_{k=1}^K w_k^{\text{MIX}} \phi(\theta|\mu_k, \Sigma_k), \quad (29)$$

where $\phi(\mu, \Sigma)$ denotes the multivariate normal density function with mean μ and variance-covariance matrix Σ , and the w_k^{MIX} are the component weights. For practical purposes, we select the number of components K using the Bayesian Information Criterion (BIC), and estimate the mixture of normals via Matlab’s built-in function `fitgmdist`.

Proposal densities for the random effects ($m_j(\alpha_j|\theta, y_j)$) were constructed participant-by-participant. For the j^{th} participant, we first fitted a normal distribution to the samples of (α_j, θ) , then derived the conditional distribution $g(\alpha_j|\theta) \sim N(\alpha_j; \mu_{j, \text{prop}}, \Sigma_{j, \text{prop}})$ for $j = 1, \dots, S$. The proposal density for subject j is the two-component defensive mixture

$$m_j(\alpha_j|\theta, y_j) = w_{\alpha}^{\text{MIX}} N(\alpha_j; \mu_{j, \text{prop}}, \Sigma_{j, \text{prop}}) + (1 - w_{\alpha}^{\text{MIX}}) p(\alpha_j|\theta) \quad (30)$$

with $p(\alpha_j|\theta)$ the prior density of α_j . The inclusion of the prior $p(\alpha_j|\theta)$ ensures that the importance weights in Eq. 5

are bounded (Hesterberg, 1995). We set the mixture weight in Eq. 30 to $w_{\alpha}^{\text{MIX}} = 0.95$.

Appendix D: Additional Applications of the Hierarchical LBA to Data

This Appendix applies the hierarchical LBA model discussed in “The speed-accuracy tradeoff in perceptual decisions” to two additional data sets.

D.1 The speed-accuracy tradeoff in lexical decisions

This section extends the analysis of the speed-accuracy tradeoff covered in the main text to judgments about the identity of strings of letters. Experiment 1 of Wagenmakers et al. (2008) had 17 participants perform a basic lexical decision task that involved repeated decisions regarding whether letter strings were valid words or not (‘non-words’). Prior to each block of trials, participants were given instructions to respond as quickly as possible (condition 1: speed emphasis) or as accurately as possible (condition 2: accuracy emphasis), where the instruction emphasis alternated between blocks. Concurrent to the speed-accuracy manipulation, word frequency was manipulated across four levels: words of very low frequency (vlf), low frequency (lf) and high frequency (hf), and non-words (nw). Participants completed 20 blocks with 96 trials per block for a total of 1920 lexical decisions per participant. See Wagenmakers et al. (2008) for all remaining details.

Rae, Heathcote, Donkin, Averell, and Brown (2014) re-analysed Wagenmakers et al. (2008)’s data to test whether instructions manipulating the speed-accuracy tradeoff only cause changes in decision caution (response threshold parameters) or decision caution *and* the speed of information processing (drift rate parameters). Through a large-scale maximum likelihood-based parameter estimation exercise, Rae et al. (2014) found evidence that emphasising the speed of decisions caused participants to lower their response threshold *and* increase their drift rates for correct and incorrect responses. The latter result suggests that time pressure increased the overall drive to respond (an overall increase in both correct and incorrect drift rates) while decreasing accuracy. We reassessed these findings, through the lens of the marginal likelihood, comparing five LBA models. For clarity in the following, we refer to instructions that emphasise accuracy and speed as a and s , respectively, and the correct and error accumulators as c and e , respectively.

Model I assumes there are no behavioural differences across all conditions (i.e., a null model), corresponding to a single set of LBA parameters over conditions. The

vector of random effects for subject j is

$$\alpha_j = \left(\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j} \right).$$

Model II assumes that decision caution differs as a function of the task instructions (speed vs accuracy), thus allowing two response threshold parameters,

$$\alpha_j = \left(\alpha_{b_j^{(s)}}, \alpha_{b_j^{(a)}}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j} \right).$$

Model III assumes the speed of information processing for the correct response, but not the incorrect response, differs as a function of task instructions,

$$\alpha_j = \left(\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c,s)}}, \alpha_{v_j^{(c,a)}}, \alpha_{\tau_j} \right),$$

where $v_j^{(c,s)}$ and $v_j^{(c,a)}$ are the mean drift rates for the accumulators corresponding to the correct response in the speed- and accuracy-emphasis conditions, respectively.

Model IV extends Model III to also allow variation in the speed of information processing for the incorrect response across task instructions,

$$\alpha_j = \left(\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e,s)}}, \alpha_{v_j^{(e,a)}}, \alpha_{v_j^{(c,s)}}, \alpha_{v_j^{(c,a)}}, \alpha_{\tau_j} \right),$$

where $v_j^{(e,s)}$ and $v_j^{(e,a)}$ are the mean drift rates for the incorrect accumulators in the speed- and accuracy-emphasis conditions, respectively.

Model V combines Models II and IV, thus allowing the response threshold and the drift rates for correct and incorrect responses to vary as a function of task instructions,

$$\alpha_j = \left(\alpha_{b_j^{(s)}}, \alpha_{b_j^{(a)}}, \alpha_{A_j}, \alpha_{v_j^{(e,s)}}, \alpha_{v_j^{(e,a)}}, \alpha_{v_j^{(c,s)}}, \alpha_{v_j^{(c,a)}}, \alpha_{\tau_j} \right).$$

Table 3 shows the log of the marginal likelihood estimates (with standard errors in parentheses) for the five models. The Monte Carlo standard errors for the estimates are again small for all models, suggesting that the IS² method is very efficient. Unsurprisingly, Model I, the most rigid model that allows no parameter variation across conditions, has the smallest marginal likelihood estimate. Model II, which allows response thresholds to vary with task instruction, provides a better representation of the data than Model III, which only allows the mean correct drift rates to vary with task instruction. Interestingly, Model IV, which allows the correct and incorrect drift rates to vary across speed and accuracy instructions, performs better than a model that only allows the response thresholds to differ over those conditions (i.e., Model II). Overall, the results favour Model V, which allows the response thresholds and the correct and

incorrect drift rates to vary with task instructions. This is consistent with the general conclusions of Rae et al. (2014). Relative to instructions emphasising decision accuracy, speed-emphasis instructions cause people to lower their response thresholds and increase their rates of accumulation for both the correct and incorrect response options.

D.2 Biasing lexical decisions

Here we analyse decisions in the presence of response bias. A common method for inducing biased decisions is to manipulate the proportion of stimuli from different response categories; for example, presenting more stimuli from category 1 than category 2 will result in a greater proportion of responses in favour of category 1 over 2. The effect of this form of bias manipulation is most commonly explained as a decision process that is primed to give a response for the more common stimulus on the basis of less evidence than would be required to give a response for the less common stimulus (Voss, Rothermund, and Voss, 2004). The LBA model captures this effect of response-relevant information through differences in the threshold across response options: the more common stimulus has a *lower* response threshold than the less common stimulus (Brown & Heathcote, 2008), thus requiring less evidence to trigger a response. This parameter change has been confirmed experimentally (Forstmann, Brown, Dutilh, Neumann, & Wagenmakers, 2010).

Here, we investigate whether changing the proportion of words and non-words induces response biases in lexical decisions, as reflected in the parameter estimates of the LBA model. Experiment 2 of Wagenmakers et al. (2008) had 19 participants perform a lexical decision task where the proportion of word vs non-word stimuli alternated across blocks: in a ‘word’ block, 75% of the stimuli were words and 25% were non-words; in a ‘non-word’ block the proportions switched such that 75% of the stimuli were non-words and 25% were words. Each block had an approximately equal amount of high, low, and very low frequency words. Participants completed 20 blocks with 96 trials per block for a total of 1920 lexical decisions per participant. See Wagenmakers et al. (2008) for all other details.

We test whether the word vs non-word response proportion manipulation affected the response threshold parameters of the LBA model in the expected direction. We used the marginal likelihood to compare five LBA models. In addition to the notation introduced above, we refer to the stimulus manipulation of the proportion of words and non-words as w and nw , where w refers to the condition with 75% words and 25% non-words and vice versa for nw , and W and NW to refer to a response of word and non-word, respectively.

Table 3 Log of the marginal likelihoods for Experiment 1 of Wagenmakers et al. (2008). Bootstrapped standard errors are in parentheses. Computation time is in minutes

Model	I	II	III	IV	V
$\log \hat{p}(y)$	−200.53 (0.20)	4959.32 (0.06)	2364.53 (0.15)	5413.50 (0.10)	5719.70 (0.16)
PMwG (in minutes)	54.70	62.76	55.88	57.05	62.03
IS ² (in minutes)	31.88	34.27	34.44	33.54	35.49
Total CPU time (in minutes)	86.58	97.03	90.32	90.59	97.52

Model I again assumes constancy in LBA parameters across conditions,

$$\alpha_j = (\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j}).$$

Model II assumes that the response threshold varies as a function of the proportion manipulation (w, nw) but does not differ for responses of words or non-words (W, NW),

$$\alpha_j = (\alpha_{b_j^{(w)}}, \alpha_{b_j^{(nw)}}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j}).$$

Model III assumes that the two responses (“word”, W ; “non-word”, NW) are governed by accumulators that had different response thresholds, and these thresholds are also free to vary across the proportion manipulation,

$$\alpha_j = (\alpha_{b_j^{(w,W)}}, \alpha_{b_j^{(w,NW)}}, \alpha_{b_j^{(nw,W)}}, \alpha_{b_j^{(nw,NW)}}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j}),$$

where $b_j^{(w,W)}$ and $b_j^{(w,NW)}$ refer to the threshold parameters for word and non-word responses in the 75% word condition, respectively, and similarly for $b_j^{(nw,W)}$ and $b_j^{(nw,NW)}$ in the 75% non-word condition.

Model IV has a parallel structure to Model III except that the proportion manipulation (w, nw) and response (W, NW) are assumed to selectively influence non-decision time rather than the response threshold,

$$\alpha_j = (\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e)}}, \alpha_{v_j^{(c)}}, \alpha_{\tau_j^{(w,W)}}, \alpha_{\tau_j^{(w,NW)}}, \alpha_{\tau_j^{(nw,W)}}, \alpha_{\tau_j^{(nw,NW)}}),$$

where the random effects $\tau_j^{(w,W)}$, $\tau_j^{(w,NW)}$, $\tau_j^{(nw,W)}$, and $\tau_j^{(nw,NW)}$ were similarly defined as in the threshold parameters of Model III.

Model V assumes that the proportion manipulation influences the rate of evidence accumulation for correct and incorrect “word” and “non-word” responses, separately for each proportion condition,

$$\alpha_j = (\alpha_{b_j}, \alpha_{A_j}, \alpha_{v_j^{(e,w)}}, \alpha_{v_j^{(e,nw)}}, \alpha_{v_j^{(c,w)}}, \alpha_{v_j^{(c,nw)}}, \alpha_{\tau_j^{(w,W)}}, \alpha_{\tau_j^{(w,NW)}}, \alpha_{\tau_j^{(nw,W)}}, \alpha_{\tau_j^{(nw,NW)}}),$$

where $v_j^{(e, \cdot)}$ and $v_j^{(c, \cdot)}$ are the vectors of mean drift rates for incorrect and correct responses, respectively. The relevant pair of correct-incorrect drift rates for a given datum are contingent on the proportion condition of the trial (75% words or non-words; w, nw) and the observed response (word or non-word; W, NW). Phrased differently, each datum (RT and response choice pair from a single trial) contains information to update 2 of the 8 drift rate parameters.

Table 4 shows the log of the marginal likelihoods for the five models. As in the previous applications, the Monte Carlo standard errors for the estimates are small for all models, although the more complex model V has slightly higher standard errors; nevertheless, even this larger standard error is still much smaller than the differences in the log of the marginal likelihoods between models.

Models that do not allow any parameter variation (I) or a simple threshold change across proportion condition but not response type (II) provide a poor representation of the data; these models do not account for response bias, which suggests there were considerable bias effects present in the data. Assuming independent drift rates for each combination of correct/incorrect, response type and proportion condition (V) provided a much improved fit, though it was still inferior to models that assume a selective influence of the response threshold (III) or non-decision time (IV) on the response type $\times$ proportion condition interaction. The best account of the data is one in which the biased responses are assumed to arise from a change in the response thresholds (III). The parameter estimates of the preferred model suggest that the biased responses arose from lowered response thresholds for more probable responses.

Appendix E: Proofs

Using the same notation as in Section A.1, we follow Pitt et al. (2012) and define the joint prior density of θ and z as $p(\theta) \exp(z) g_N(z|\theta)$, so that the posterior density of θ and z is

$$\tilde{\pi}(\theta, z) = \frac{p(y|\theta) p(\theta) \exp(z) g_N(z|\theta)}{p(y)} = \frac{\hat{p}_N(y|\theta) p(\theta) g_N(z|\theta)}{p(y)}, \quad (31)$$

Table 4 Log of the marginal likelihoods for Experiment 2 of Wagenmakers et al. (2008). Bootstrapped standard errors are in parentheses. Computation time is in minutes

Model	I	II	III	IV	V
log $\hat{p}(y)$	5769.33 (0.09)	5876.11 (0.05)	8449.08 (0.11)	8259.95 (0.13)	7501.00 (0.39)
PMwG (in minutes)	100.52	103.01	107.60	105.40	116.25
IS ² (in minutes)	35.82	39.58	42.96	45.24	57.88
Total CPU time (in minutes)	136.34	142.59	150.56	150.64	174.13

and has $p(\theta|y)$ as a marginal. Hence,

$$p(y) = \int_{\Theta} \int_{\mathcal{Z}} \hat{p}_N(y|\theta) p(\theta) g_N(z|\theta) dz d\theta = \int_{\Theta} \int_{\mathcal{Z}} \tilde{w}(\theta, z) \tilde{g}_{IS}(\theta, z) dz d\theta \quad (32)$$

where

$$\tilde{w}(\theta, z) = \frac{\hat{p}_N(y|\theta) p(\theta) g_N(z|\theta)}{\tilde{g}_{IS}(\theta, z)} \quad \text{and} \quad \tilde{g}_{IS}(\theta, z) = p(\theta) g_N(z|\theta) \quad (33)$$

Thus, we can consider $\tilde{g}_{IS}$ in Eq. 32 as an importance density in z and θ , which will lead to the same IS² estimate for the marginal likelihood as in “[Estimating the marginal likelihood by IS²](#)”.

Proof of Theorem 1 Unbiasedness holds because we are dealing with IS. The assumptions in Theorem A1 ensure that $\tilde{w}(\theta_i)$ ’s are i.i.d with a finite second moment. The results of (i) and (ii) then follow. $\square$

Proof of Theorem A1 If $\text{Supp}(\pi) \subseteq \text{Supp}(g_{IS})$ then $\text{Supp}(\pi_N) \subseteq \text{Supp}(\tilde{g}_{IS})$. This, together with the existence and finiteness of $\mathbb{E}_{\pi}(\varphi)$ ensures that $\mathbb{E}_{\tilde{g}_{IS}}[\varphi(\theta)\tilde{w}(\theta, z)] = p(y)\mathbb{E}_{\pi}(\varphi)$ and $\mathbb{E}_{\tilde{g}_{IS}}[\hat{w}(\theta, z)] = p(y)$ exist and are finite. Result (i) follows from the strong law of large numbers. To prove (ii), write

$$\hat{E}_{\pi}(\varphi) - \mathbb{E}_{\pi}(\varphi) = \frac{\frac{1}{M} \sum_{i=1}^M (\varphi(\theta_i) - \mathbb{E}_{\pi}(\varphi)) \tilde{w}(\theta_i, z_i)}{\frac{1}{M} \sum_{i=1}^M \tilde{w}(\theta_i, z_i)}. \quad (34)$$

Let $X_i = (\varphi(\theta_i) - \mathbb{E}_{\pi}(\varphi)) \tilde{w}(\theta_i, z_i)$, $i = 1, \dots, M$, $S_M = \frac{1}{M} \sum_{i=1}^M X_i$ and $Y_M = \frac{1}{M} \sum_{i=1}^M \tilde{w}(\theta_i, z_i)$. The X_i are independently and identically distributed with $\mathbb{E}_{\tilde{g}_{IS}}(X_i) = 0$ and

$$\begin{aligned} v_N^2 &= \mathbb{V}_{\tilde{g}_{IS}}(X_i) = \mathbb{E}_{\tilde{g}_{IS}}(X_i^2) = \int_{\Theta} \int_{\mathcal{Z}} (\varphi(\theta) \\ &\quad - \mathbb{E}_{\pi}(\varphi))^2 \tilde{w}(\theta, z) p(y) \pi_N(\theta, z) d\theta dz \\ &= p(y) \mathbb{E}_{\pi} \left\{ (\varphi(\theta) - \mathbb{E}_{\pi}(\varphi))^2 \tilde{w}(\theta, z) \right\} < \infty. \end{aligned}$$

By the central limit theorem for a sum of independently and identically distributed random variables with a finite second moment, $\sqrt{M} S_M \xrightarrow{d} \mathcal{N}(0, v_N^2)$. By the strong law of large numbers, $Y_M \xrightarrow{P} \mathbb{E}_{\tilde{g}_{IS}}(\tilde{w}(\theta, z)) = p(y)$. By Slutsky’s theorem,

$$\sqrt{M} \left(\hat{E}_{\pi}(\varphi) - \mathbb{E}_{\pi}(\varphi) \right) = \frac{\sqrt{M} S_M}{Y_M} \xrightarrow{d} \mathcal{N}(0, v_N^2 / p(y)^2).$$

The asymptotic variance is given by

$$\begin{aligned} \hat{E}_{\pi}(\varphi) &= \frac{v_N^2}{p(y)^2} = \frac{1}{p(y)} \mathbb{E}_{\pi_N} \left\{ (\varphi(\theta) - \mathbb{E}_{\pi}(\varphi))^2 \tilde{w}(\theta, z) \right\} \\ &= \mathbb{E}_{\pi} \left\{ (\varphi(\theta) - \mathbb{E}_{\pi}(\varphi))^2 \frac{\pi(\theta)}{g_{IS}(\theta)} \mathbb{E}_{g_N(z|\theta)}[\exp(2z)] \right\}. \end{aligned}$$

To prove (iii), write

$$\begin{aligned} \hat{\sigma}_{IS^2}^2(\varphi) &= \frac{\frac{1}{M} \sum_{i=1}^M \left(\varphi(\theta_i) - \hat{E}_{\pi}(\varphi) \right)^2 \tilde{w}(\theta_i, z_i)^2}{\left(\frac{1}{M} \sum_{i=1}^M \tilde{w}(\theta_i, z_i) \right)^2} \\ &\xrightarrow{a.s.} \frac{\mathbb{E}_{\tilde{g}_{IS}} \left\{ (\varphi(\theta) - \mathbb{E}_{\pi}(\varphi))^2 \tilde{w}(\theta, z)^2 \right\}}{\left(\mathbb{E}_{\tilde{g}_{IS}}(\tilde{w}(\theta, z)) \right)^2} \\ &= \sigma_{IS^2}^2(\varphi). \end{aligned}$$

$\square$

Proof of Theorem A2 Under Assumption 2, $g_N(z|\theta) = \mathcal{N}(-\sigma^2/2, \sigma^2)$ and $\mathbb{E}_{g_N(z|\theta)}[\exp(2z)] = \exp(\sigma^2)$. From Eqs. 21 and 22,

$$\begin{aligned} \sigma_{IS^2}^2(\varphi) &= \mathbb{E}_{\pi} \left\{ (\varphi(\theta) - \mathbb{E}_{\pi}(\varphi))^2 \frac{\pi(\theta)}{g_{IS}(\theta)} \exp(\sigma^2) \right\} \\ &= \exp(\sigma^2) \sigma_{IS}^2(\varphi). \end{aligned} \quad (35)$$

$\square$

References

Andrieu, C., Doucet, A., & Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society Series B*, 72, 1–33.

- Andrieu, C., & Roberts, G. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37, 697–725.
- Brown, S., & Heathcote, A. (2008). The simple complete model of choice reaction time: Linear Ballistic accumulation. *Cognitive Psychology*, 57, 153–178.
- Chopin, N., Jacob, P. E., & Papaspiliopoulos, O. (2013). SMC²: An efficient algorithm for sequential analysis of state space models. *Journal of Royal Statistical Society Series B*, 75(3), 397–426.
- Donkin, C., Brown, S. D., & Heathcote, A. J. (2009). The over-constraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, 16, 1129–1135.
- Doucet, A., Pitt, M. K., Deligiannidis, G., & Kohn, R. (2015). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2), 295–313.
- Evans, N. J., & Annis, J. (2019). Thermodynamic integration via differential evolution: A method for estimating marginal likelihoods. *Behavior Research Methods*.
- Evans, N. J., & Brown, S. D. (2018). Bayes factors for the linear ballistic accumulator model of decision-making. *Behavior Research Methods*, 50(2), 589–603.
- Fiebig, D. G., Keane, M. P., Louviere, J., & Wasi, N. (2010). The generalized multinomial logit model: accounting for scale and coefficient heterogeneity. *Marketing Science*, 29(3), 393–421.
- Fitzmaurice, G. M., Laird, N. M., & Ware, J. H. (2011). *Applied longitudinal analysis*. (2nd ed.). New Jersey: Wiley.
- Forstmann, B. U., Brown, S., Dutilh, G., Neumann, J., & Wagenmakers, E.-J. (2010). The neural substrate of prior information in perceptual decision making: a model-based analysis. *Frontiers in Human Neuroscience*, 4, 40.
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., & von Cramon, D. Y. (2008). Striatum and pre-sma facilitate decision making under time pressure. *Proceedings of the National Academy of Sciences*, 105, 17538–17542.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, 1(3), 515–534.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 57(6), 1317–1339.
- Gourieroux, C., & Monfort, A. (1995). *Statistics and econometric models* Vol. 1. Cambridge: Cambridge University Press.
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., . . . , Steingroever, H. (2017). A tutorial on bridge sampling. *Journal of Mathematical Psychology*, 81, 80–97.
- Gronau, Q. F., Wagenmakers, E.-J., Heck, D. W., & Matzke, D. (2019). A simple method for comparing complex models: Bayesian model comparison for hierarchical multinomial processing tree models using warp-iii bridge sampling. *Psychometrika*, 84(1), 261–284.
- Gunawan, D., Hawkins, G. E., Tran, M.-N., Kohn, R., & Brown, S. D. (2019). New estimation approaches for the linear ballistic accumulator model. Working paper. arXiv:1806.10089.
- Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in neuroscience*, 8, 150.
- Hesterberg, T. (1995). Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37, 185–194.
- Hoogerheide, L., Opschoor, A., & Van Dijk, H. K. (2012). A class of adaptive importance sampling weighted EM algorithms for efficient and robust posterior and predictive simulation. *Journal of Econometrics*, 171(2), 101–120.
- Huang, A., & Wand, M. P. (2013). Simple marginally noninformative prior distributions for covariance matrices. *Bayesian Analysis*, 8(2), 439–452.
- Pitt, M. K., Silva, R. S., Giordani, P., & Kohn, R. (2012). On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171(2), 134–151.
- Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the tortoise: emphasizing speed can change the evidence used to make decisions. *Journal of experimental psychology: learning, memory, and cognition*, 40(5), 1226–1243.
- Richard, J.-F., & Zhang, W. (2007). Efficient high-dimensional importance sampling. *Journal of Econometrics*, 141(2), 1385–1411.
- Ripley, B. (1987). *Stochastic simulation*. New York: Wiley.
- Terry, A., Marley, A., Barnwal, A., Wagenmakers, E.-J., Heathcote, A., & Brown, S. D. (2015). Generalising the drift rate distribution for linear ballistic accumulators. *Journal of Mathematical Psychology*, 68, 49–58.
- Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge: Cambridge university press.
- Tran, M. N., Kohn, R., Quiroz, M., & Villani, M. (2016). Block-wise pseudo marginal Metropolis-Hastings. arXiv:1603.02485v2.
- Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, 18(3), 368–384.
- Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: an empirical validation. *Memory & Cognition*, 32(7), 1206–1220.
- Wagenmakers, E. J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58, 140–159.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.