
Meta-Uncertainty in Bayesian Model Comparison

Marvin Schmitt

Cluster of Excellence SimTech
University of Stuttgart
Germany

Stefan T. Radev

Cluster of Excellence STRUCTURES
Heidelberg University
Germany

Paul-Christian Bürkner

Cluster of Excellence SimTech
University of Stuttgart
Germany

Abstract

Bayesian model comparison (BMC) offers a principled probabilistic approach to study and rank competing models. In standard BMC, we construct a discrete probability distribution over the set of possible models, conditional on the observed data of interest. These posterior model probabilities (PMPs) are measures of uncertainty, but—when derived from a finite number of observations—are also uncertain themselves. In this paper, we conceptualize distinct levels of uncertainty which arise in BMC. We explore a fully probabilistic framework for quantifying *meta-uncertainty*, resulting in an applied method to enhance any BMC workflow. Drawing on both Bayesian and frequentist techniques, we represent the uncertainty over the uncertain PMPs via *meta-models* which combine simulated and observed data into a predictive distribution for PMPs on new data. We demonstrate the utility of the proposed method in the context of conjugate Bayesian regression, likelihood-based inference with Markov chain Monte Carlo, and simulation-based inference with neural networks.

Yet, assigning probability distributions over parameters might not address every question of interest. In fact, substantial research questions frequently revolve around comparing several competing models for an empirical phenomenon (Etz & Vandekerckhove, 2018). Since any researcher can favor any model, assessing the quality of a given model is of utmost importance, and a vast number of metrics exist to evaluate a proposed model (e.g., Akaike, 1974; Konishi, 1996; Konishi & Kitagawa, 2008). These metrics become crucial to applied statistical modeling when several models are proposed and need to be compared against each other (Bürkner et al., 2022).

As a recent example, various researchers have put forward models of disease outbreaks in order to provide forecasts and estimate the utility of interventions during the COVID-19 pandemic (e.g., Dehning et al., 2020; Wagner et al., 2022). Consequently, policy makers need to compare the relative merits of plausible candidate models and decide which models to consider as a basis for implementing effective policies. Thus, model comparison (along with model selection) turns into a central topic with a potentially enormous societal impact.

The current work is located at the core of Bayesian model comparison research. We study how distinct sources of uncertainty arise in Bayesian model comparison settings and introduce *meta-uncertainty* as a novel concept to describe the model-implied variation of inference quantities (see Figure 1). Specifically, our main contributions are:

- (i) a fully probabilistic framework to describe qualitatively distinct *levels of uncertainty* arising in any Bayesian model comparison setting;
- (ii) a method to *augment* observed data through simulated data, leading to a principled embedding of observed posterior model probabilities into the probabilistic context of all competing models;
- (iii) a formal argument that our method is *consistent* and *mitigates overconfidence* under certain conditions;
- (iv) an actionable interpretation for applied researchers in terms of the *replicability* of observed results.

1 INTRODUCTION

Scientific reasoning rarely comes with absolute certainty. However, properly accounting for *uncertainty* at each step of a research project is challenging, let alone quantifying and communicating all sources of uncertainty adequately.

Bayesian statistics offer an end-to-end probabilistic framework for assigning probability distributions to model parameters and propagating the associated uncertainty throughout the entire workflow (Gelman et al., 2014).

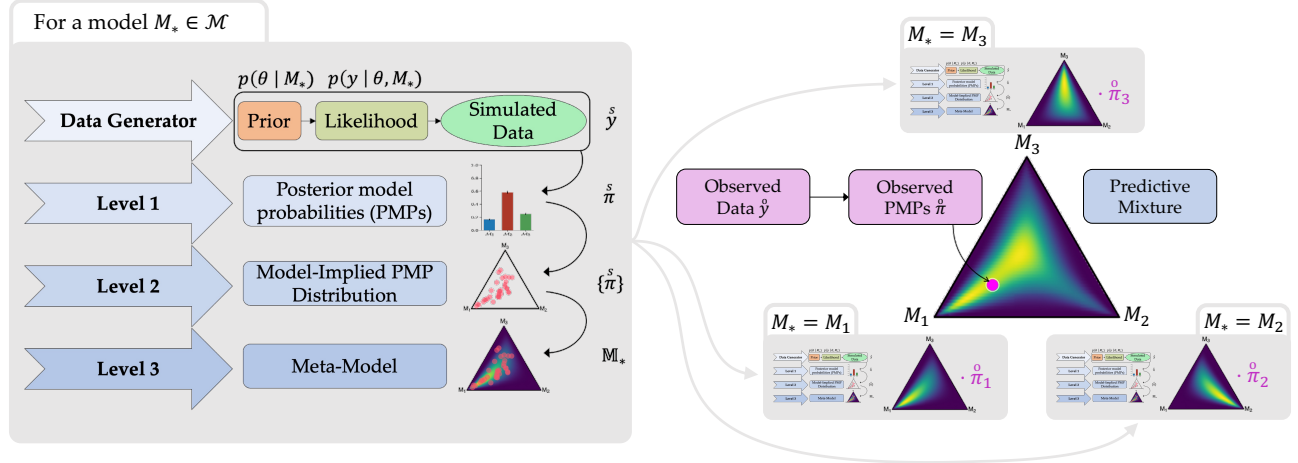


Figure 1: Meta-uncertainty acts as an overarching concept to capture the uncertainty in posterior model probabilities (PMPs), which are measures of uncertainty themselves. The result is a posterior predictive mixture distribution for PMPs on new observed data, for example from a replication study. This predictive mixture distribution is consistent, mitigates overconfidence, and gauges the reproducibility of observed model comparison results in future studies or experiments.

2 BACKGROUND

Bayesian Statistics Bayesian methods constitute a principled approach to address questions of statistical inference from a self-consistent probabilistic perspective (Gelman et al., 2014). Bayesian models formalize the assumed data generating process as

$$y \sim p(y | \theta, M) \quad \text{with} \quad \theta \sim p(\theta | M), \quad (1)$$

with observed quantities (data) $y \in \mathcal{Y}$, parameters $\theta \in \Theta$, and model $M \in \mathcal{M}$. The posterior distribution

$$p(\theta | y, M) = \frac{p(y | \theta, M) p(\theta | M)}{p(y | M)} \quad (2)$$

is the main endpoint of Bayesian inference as it captures our information gain regarding the parameters θ of model M , conditional on the data y .

Uncertainty Quantification In a Bayesian context, uncertainty is a paramount concept which often entertains distinct meanings and consequences (Hüllermeier & Waegeman, 2021). The posterior distribution in Equation 2 implies that our knowledge about the parameter values θ is not absolutely precise, but surrounded by *epistemic* uncertainty. Notably, epistemic uncertainty is reducible: Under certain regularity conditions, the posterior distribution $p(\theta | y, M)$ will shrink to a Dirac impulse, with no epistemic uncertainty left as the number of observations approaches infinity (van der Vaart, 2000).

However, even if the exact “true” values of the parameters θ were known and there were no remaining epistemic uncertainty, Equation 1 would still entail uncertainty due to the

likelihood function $p(y | \theta, M)$. This remaining uncertainty is irreducible and follows from the *aleatoric* “intrinsic randomness” of the data generating process (Hüllermeier & Waegeman, 2021).

Bayesian Model Comparison In light of observed data y , Bayesian model comparison (BMC) establishes a framework to compare J candidate models $M_1, \dots, M_J \in \mathcal{M}$ conditional on the observations y . A key quantity in BMC is the marginal likelihood

$$p(y | M_j) = \int_{\Theta_j} p(y | \theta, M_j) p(\theta | M_j) d\theta, \quad (3)$$

which quantifies the evidence for model M_j by averaging the likelihood function over all possible values of the parameters θ .

Once we have obtained (or approximated) the marginal likelihoods of individual models, we can compute their *posterior model probabilities* (PMPs) via Bayes’ rule

$$\pi_j := p(M_j | y) = \frac{p(y | M_j) p(M_j)}{\sum_{M' \in \mathcal{M}} p(y | M') p(M')}, \quad (4)$$

where $p(M)$ represents the prior model probability (e.g., uniform if models are deemed equally plausible before observing any data). To shorten notation, we collect the PMPs of all considered models into a vector $\pi = (\pi_1, \dots, \pi_J)$.

By definition, PMPs (as well as Bayes Factors; Kass & Raftery, 1995) are sensitive to the specification of the prior (Oelrich et al., 2020; Schad et al., 2021), and specifying sensible joint priors is extremely challenging in practical modeling scenarios (Aguilar & Bürkner, 2022; Dellaportas

et al., 2012; Lindley, 1957; Van Dongen, 2006; Zitzmann et al., 2021). Consequently, our method for augmenting PMPs can help practitioners evaluate their prior choices by unveiling undesired patterns in the resulting PMPs for both (i) samples from each model’s prior predictive distribution and (ii) actually observed real data.

Predictive Distribution Differently from marginal likelihood approaches, we can also perform BMC based on posterior predictive performance. This family of methods builds on the *posterior predictive distribution*, which expresses the probability of new data \tilde{y} by accounting for all uncertainty implied by the posterior (cf. Equation 2):

$$p(\tilde{y} | y, M_j) = \int_{\Theta_j} p(\tilde{y} | \theta, y, M_j) p(\theta | y, M_j) d\theta. \quad (5)$$

Predictive methods, typically paired with cross-validation, are less sensitive to prior specification, but do not enjoy the consistency guarantees of Bayes Factors and PMPs either (Bürkner et al., 2020; Vehtari & Lampinen, 2002; Vehtari & Ojanen, 2012; Vehtari et al., 2021).¹

Importantly, our method combines aspects from both marginal likelihood and posterior predictive approaches to arrive at a consistent framework for a more nuanced uncertainty quantification in model comparison settings.

3 METHOD

Drawing on both Bayesian and frequentist concepts, we propose a mathematically sound description of uncertainty levels and use the resulting framework to augment observed PMPs with meta-uncertainty (see Figure 1 and Table 1).

3.1 Model-Implied Data and PMP Distributions

Each Bayesian model $M_j \in \mathcal{M}$ implies a recipe for simulating data \dot{y} by randomly sampling from its prior and likelihood (Equation 1). Thus, every Bayesian model with a proper prior is *generative*, and simulating data from each model $M_j \in \mathcal{M}$ is straightforward. The decorated symbol \dot{y} emphasizes that the data are simulated from the generative models and not actually observed. This distinction is essential in the context of our framework, as it integrates simulated and actually observed data. By applying Equation 3 and Equation 4, we can calculate a vector of PMPs

$$\dot{\pi} = \left(\dot{\pi}_1, \dots, \dot{\pi}_J \right) = \left(p(M_1 | \dot{y}), \dots, p(M_J | \dot{y}) \right) \quad (6)$$

for any fixed simulated data set \dot{y} . The resulting PMP vector $\dot{\pi}$ forms the basis of all subsequently discussed uncertainty levels.

¹See also Fong and Holmes (2020) for a connection between the marginal likelihood and exhaustive cross-validation.

Algorithm 1 Generation of a model-implied PMP distribution (level 2 uncertainty).

Input: \mathcal{M}, K, N

- 1: **for** $k = 1, \dots, K$ **do**
 - 2: $M_*^{(k)} \sim p(M)$ {sample true model}
 - 3: $\theta^{(k)} \sim p(\theta | M_*^{(k)})$ {sample parameter}
 - 4: $\dot{y}^{(k)} \sim p(y | \theta^{(k)}, M_*^{(k)})$ {sample N observations}
 - 5: $\dot{\pi}^{(k)} \leftarrow (p(M_1 | \dot{y}^{(k)}), \dots, p(M_J | \dot{y}^{(k)}))$ {PMPs}
 - 6: **end for**
 - 7: **return** $\{(\dot{\pi}^{(k)}, M_*^{(k)})\}_{k=1}^K$
-

Level 1 Uncertainty Each model $M_j \in \mathcal{M}$ implies a marginal likelihood of the data y , which integrates out the parameters θ from the joint distribution $p(\theta, y)$ and thus averages out the epistemic uncertainty encoded in the prior distribution (cf. Equation 3). The uncertainty in the PMP vector—i.e., deviations from a one-hot encoding (aka. dummy coding with J instead of $J-1$ variables) where one PMP equals one and all others equal zero—is essential to our framework and we refer to it as *level 1 uncertainty*. This uncertainty is epistemic, because we could, in principle, reduce it with more data: In the limit of an arbitrary number of observations, $N \rightarrow \infty$, the level 1 uncertainty vanishes, since PMPs are consistent (Barron et al., 1999).

Simulation Paradigm By applying Equation 1 repeatedly, we can generate a set of K simulated data sets $\{\dot{y}^{(k)}\}_{k=1}^K$, which we call “model-implied” data. This set of data sets corresponds to a sample from the prior predictive distribution. Computing the PMPs for each simulated data set in $\{\dot{y}^{(k)}\}_{k=1}^K$ leads to a set of “model-implied” PMPs $\{\dot{\pi}^{(k)}\}_{k=1}^K$. Algorithm 1 outlines the essential steps of these computations. Figure 3 shows the distribution of model-implied PMPs under each candidate model for a toy example in Experiment 1.

Level 2 Uncertainty The data generating process outlined in Algorithm 1 features aleatoric uncertainty, which translates into variation in the PMPs $\dot{\pi}^{(k)}$ across simulations from the same model. We call the variation inherent in the *model-implied* PMP distribution *level 2 uncertainty*. The level 2 uncertainty yields a notion of the aleatoric uncertainty in the model-implied distribution of PMPs due to the inherent randomness of the data generating process: The PMPs conditional on two finite data sets $\dot{y}^{(1)}, \dot{y}^{(2)}$ may differ even if the two data sets were generated by the same parameter configuration θ of the same true model M_* . The model-implied PMP distributions may unveil complex patterns and interactions between the models on an intuitive scale, that is, in terms of PMPs. In practical applications, this may guide practitioners’ endeavors for both model specification and design optimization.

Relation to Sampling Distributions Studying the statistical behavior of a sample of inference targets—PMPs in our case—under competing models is closely related to the concept of sampling distributions, which aim to capture sample variations of frequentist origin (Kulesa et al., 2015; Schad et al., 2021). Formulating the statement “model M_j is the true model M_* ” as a hypothesis and the PMPs π as an inference target in accordance with frequentist terminology, the simulated sample of PMPs can be interpreted as a random draw from the sampling distribution $p(\pi | M_j = M_*)$. Consequently, as the number of simulations grows arbitrarily large ($K \rightarrow \infty$), we reproduce the *exact* sampling distribution of the model-implied PMPs. In what follows, we will utilize the model-implied distribution of PMPs to contextualize observed PMPs.

Computational Efficiency Obtaining the model-implied PMPs (i.e., line 5 in Algorithm 1) can be quite demanding. On the one hand, computing the marginal likelihood for non-conjugate models is hard, since it requires integration. In these most typical cases, we can resort to approximations based on Markov chain Monte Carlo (MCMC), such as bridge sampling (Gronau et al., 2020; Meng & Schilling, 2002). On the other hand, the marginal likelihood becomes doubly intractable if the likelihood function itself is not available in a closed form (e.g., as in simulation-based inference, Cranmer et al., 2020). In this case, we can fall back on approximate Bayesian computation (Csilléry et al., 2010) or more scalable approximators relying on amortized neural estimation (Radev et al., 2021; Radev et al., 2023). Indeed, in our experiments, we demonstrate the integration of meta-uncertainty across all three settings (i.e., conjugate, likelihood-based, and simulation-based).

3.2 Meta-Models of Model-Implied PMP Distributions

In this section, we will analyze the model-implied distribution of PMPs using a Bayesian approach. First, we group the model-implied PMPs obtained via Algorithm 1 by the respective true model indices $M_* \in \mathcal{M}$, resulting in J groups. The indices $j = 1, \dots, J$ express which model generated the data in the respective group.

For each group, we fit a Bayesian model with parameters τ_j on the vectors of simulated PMPs $\tilde{\pi}$ for data from the respective true model $M_j = M_*$, namely the set $\{\tilde{\pi}^{(k)}\}_{M_*^{(k)}=M_j}$. We refer to these models as *meta-models* and denote them as \mathbb{M}_j . To shorten notation, we abbreviate the set by writing $\{\tilde{\pi}\}_j$ when the context is clear and no strictly formal notation is necessary. The likelihood $p_j(\pi | \tau_j)$ of each meta-model \mathbb{M}_j can be any density with support on the unit simplex, such as a Dirichlet density with $\tau = \alpha$, or a logistic Normal density with $\tau = (\mu, \Sigma)$.

Each meta-model \mathbb{M}_j yields a posterior distribution

$p_j(\tau_j | \{\tilde{\pi}\}_j)$ over its parameters $\tau_j \in \mathfrak{T}_j$. This posterior is informed by the sampling distribution of PMPs estimated from the corresponding model-implied data. In practice, we can readily fit the meta-models with the R package `brms` (Bürkner, 2017, 2018) as an interface to `Stan` (Stan Development Team, 2022) to obtain D draws $\{\tau_j^{(d)}\}_{d=1}^D$ from the posterior distribution $p_j(\tau_j | \{\tilde{\pi}\}_j)$ based on Hamiltonian Monte Carlo (HMC; Betancourt, 2017; Neal, 2011). To shorten notation, we will refer to the D posterior draws as $\{\tau_j^{(d)}\}$. Finally, each meta-model implies a posterior predictive distribution,

$$p_j(\tilde{\pi} | \{\tilde{\pi}\}_j) = \int_{\mathfrak{T}_j} p_j(\tilde{\pi} | \tau_j) p(\tau_j | \{\tilde{\pi}\}_j) d\tau_j, \quad (7)$$

which captures the distribution of PMPs for a new observed data set from the perspective of the meta-model \mathbb{M}_j .

Level 3 Uncertainty The J posterior distributions of the meta-models have epistemic uncertainty owing to the finite number of simulated data sets K . We refer to this uncertainty as *level 3 uncertainty*, which can be reduced with more simulations and vanishes for $K \rightarrow \infty$. What is more, the epistemic level 3 uncertainty is contingent on the aleatoric uncertainty from the meta-models’ likelihood functions. Thus, a substantial level 3 uncertainty might indicate that more simulations are required to improve the meta-model and learn about the patterns on level 2.

3.3 Predictive Mixture Distributions

We now describe how to use the information from model simulations and model-implied PMPs to inform the analysis of real data. For an actually observed data set \mathcal{y} and the set of candidate models $\mathcal{M} = \{M_1, \dots, M_J\}$, we first compute the observed PMPs $\hat{\pi} = (\hat{\pi}_1, \dots, \hat{\pi}_J)$ according to Equation 4. Then, we can embed the observed PMPs into the context of the model-implied PMPs distributions. We achieve this by constructing a mixture distribution, whose components are the posterior predictive distributions of the meta-models from Equation 7 and whose mixture weights are the observed PMPs $\hat{\pi}$:

$$f(\tilde{\pi}) = \sum_{j=1}^J \hat{\pi}_j p_j(\tilde{\pi} | \{\tilde{\pi}\}_j). \quad (8)$$

Accordingly, $f(\tilde{\pi})$ represents a data-informed mixture distribution of simulation-based meta-models. In other words, Equation 8 expresses a distribution over PMPs for a new (e.g., future or held-out) data set, based on the predictive distributions of the meta-models. In real observational studies or experiments, the predictive mixture distribution can thus provide a *notion of the reproducibility or replicability of a given model comparison result with regard to a subsequent study*.

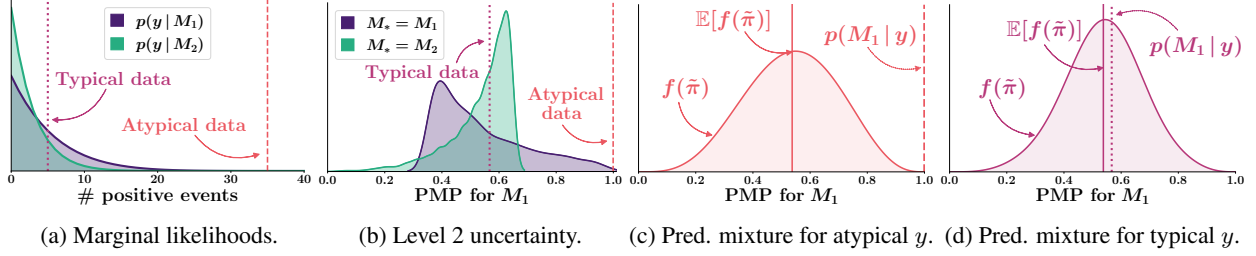


Figure 2: Illustration of the predictive mixture method for the minimal motivating example comparing two conjugate Beta-Binomial models. Level 2 uncertainty in (b) is visualized by a kernel density estimate (KDE) on $K = 10\,000$ draws from the model-implied PMP distribution. The KDE is not part of our method but only used for visualization.

In most analyses, the predictive distributions from Equation 7 and Equation 8 will be available to modelers as a set of (potentially thousands of) posterior draws. Each posterior draw thus leads to a separate *pushforward likelihood* and a separate density plot. To ease visualization and analysis, we propose to compute embeddings (summary statistics) of the posterior draws $\{\tau_j^{(d)}\}$ with a set of C suitable embeddings $\{\mathbb{T}_1 \dots \mathbb{T}_C\}$, $\mathbb{T}_c : \mathfrak{T}^D \rightarrow \mathfrak{T}$, leading to the computationally more feasible construction

$$f_{\mathbb{T}_c}(\tilde{\pi}) = \sum_{j=1}^J \tilde{\pi}_j p_j(\tilde{\pi} \mid \mathbb{T}_c(\{\tau_j^{(d)}\})). \quad (9)$$

While the expected value would seem like a natural choice for \mathbb{T} with $C=1$, it might not behave as expected for multivariate data.² Instead, we can choose functions $\mathbb{T}_c(\{\tau_j^{(d)}\})$ which extract C clusters from the posterior draws and yield the center within each cluster (cf. Piironen et al., 2020, for a similar clustering approach in projective inference). Thus, for a set of D posterior draws $\{\tau_j^{(d)}\}$, the clustering approach reduces the number of resulting plots by a factor of D/C . While our experiments in the main paper feature a mean embedding ($\mathbb{T} = \mathbb{E}$) due to space constraints, the **Supplementary Material** contains results from the clustering technique for all experiments.

3.4 Implications and Guarantees

Proposition 1 (Consistency). *For $K > 0$ and $M_* \in \mathcal{M}$ we have $f(\tilde{\pi}) \xrightarrow{\text{a.s.}} \mathbb{I}_{M_*}$ as $N \rightarrow \infty$.*

For an arbitrarily large number of observations $N \rightarrow \infty$, the PMPs (simulated and observed) converge almost surely to one-hot encoded vectors, $\pi \xrightarrow{\text{a.s.}} \mathbb{I}_{M_*}$ with $p(M | y) = 1$ for $M = M_*$ and 0 for all other models (i.e., level 1 uncertainty approaches zero). Consequently, level 2 uncertainty vanishes since all model-implied PMPs under each model are equal. Accordingly, the meta-models will collapse to Dirac distributions, which in turn imply no level 3 uncer-

tainty. Thus, both the observed PMP and the mixture components in Equation 8 approach \mathbb{I}_{M_*} . Finally, the predictive mixture distribution for a data set \hat{y} from model M_* equals \mathbb{I}_{M_*} , which is consistent by definition. We prove Proposition 1 in the **Supplementary Material** and illustrate it in **Experiment 1** (see Figure 3, $N = 100$, bottom row).

Proposition 2 (Mitigating Overconfidence). *For $f(\tilde{\pi}) \neq \mathbb{I}_{M_j}$, the joint distribution $p(\tilde{\pi}, M_j) = \tilde{\pi}_j p_j(\tilde{\pi} \mid \{\hat{\pi}\}_j)$ has non-vanishing variance even if $\hat{\pi}$ has no level 1 uncertainty.*

In the case of spuriously high PMPs in favor of a model M_j (i.e., $\hat{\pi}_j \approx 1$), the predictive mixture distribution selects the sampling distribution of the respective model M_j . Given that the corresponding mixture component $p_j(\tilde{\pi} \mid \{\hat{\pi}\}_j)$ has a non-vanishing variance, we mitigate overconfidence: Even if $\hat{\pi} = \mathbb{I}_{M_j}$, the predictive mixture still contains all variance from level 2. We prove Proposition 2 in the **Supplementary Material** and illustrate it in **Experiment 1** (cf. Figure 4c).

3.5 Minimal Example

In order to convey some intuition about the applications and benefits of our framework, we consider a minimal example with two competing toy models M_1 and M_2 . The models follow the Beta-Binomial conjugacy

$$\begin{aligned} \theta &\sim \text{Beta}(\alpha, \beta), \\ y_n &\sim \text{Bernoulli}(\theta), \end{aligned} \quad (10)$$

with two different prior configurations: $M_1: \alpha=1, \beta=10$; $M_2: \alpha=1, \beta=20$. These priors assume that positive events (i.e., $y_n = 1$) are rare and differ solely with regard to how much mass they assign to this assumption. The marginal likelihoods and associated PMPs are analytic (Salvatier et al., 2016). Clearly, the resulting PMPs will contain level 1 uncertainty for a finite number N of observations.

To obtain a notion of level 2 uncertainty, we use Algorithm 1 to generate $K = 10\,000$ draws from the two model-implied PMP distributions ($N = 50$ for all simulated data sets) and display these draws in Figure 2b. We then fit two Beta meta-models to these PMP draws, which we will use

²Pun intended. The expected value might not even lie in the “typical set” of a distribution (Betancourt, 2017).

for constructing the predictive mixture distribution according to Equation 8. To simplify exposition, we will focus on interpreting the mean of the predictive distribution,

$$\mathbb{E}[f(\tilde{\pi})] = \overset{\circ}{\pi}_1 \mathbb{E}[p_1(\tilde{\pi} | \{\overset{\circ}{\pi}\}_1)] + \overset{\circ}{\pi}_2 \mathbb{E}[p_2(\tilde{\pi} | \{\overset{\circ}{\pi}\}_2)], \quad (11)$$

with $\overset{\circ}{\pi}_2 = (1 - \overset{\circ}{\pi}_1)$ since we only compare two models. As the epistemic uncertainty over meta-model parameters (level 3) is negligible due to the large number of simulations, we can further simplify the above expectation by representing it as a linear transformation of the observed PMPs

$$\mathbb{E}[f(\tilde{\pi})] \approx \begin{pmatrix} \bar{\pi}_1 & 1 - \bar{\pi}_2 \\ 1 - \bar{\pi}_1 & \bar{\pi}_2 \end{pmatrix} \begin{pmatrix} \overset{\circ}{\pi}_1 \\ \overset{\circ}{\pi}_2 \end{pmatrix}, \quad (12)$$

where $\bar{\pi}_1$ and $\bar{\pi}_2$ denote the corresponding means of the meta-models. These means correspond to the *average magnitude of model-implied PMPs which we expect under the assumption that each of the two models is true*.

Consider now what happens if, for instance, the data generating process undergoes a shift and positive events become much more likely (see Figure 2a). In this case, both models will be severely misspecified, but M_1 will completely “dominate” M_2 , since it has a slightly fatter right tail. As observed by Yang and Zhu (2018), researchers would ideally want an indication that the models are misspecified, but instead, PMPs will converge to 1 for the “least misspecified” model (Barron et al., 1999) and model selection based on PMPs will be markedly overconfident.

Indeed, assuming an atypical observed data set $\overset{\circ}{y}$ with 35 positive and 15 negative events (cf. Figure 2a), the PMP for model M_1 is $\overset{\circ}{\pi}_1 \approx 0.999$. However, Equation 12 suggests a more modest preference for M_1 and reduces the observed PMP to the average model-implied PMP $\bar{\pi}_1$, amounting to no more than $\bar{\pi}_1 \approx 0.55$ (cf. Figure 2c). In contrast, the PMPs obtained on a typical data set are close to the mean of the corresponding predictive distribution (cf. Figure 2d). Thus, in this example, the predictive mixture acts as a “simulation-based correction” for overconfident model selection in the presence of atypical data. In other settings, it can be used to predict the expected PMPs in an upcoming experiment or a replication study (see Experiment 1).

4 RELATED WORK

Overconfidence and Dance Researchers have observed that PMPs tend to concentrate on one candidate model, even if other evaluation metrics (such as predictive performance or information criteria) do not indicate a strong preference for that model (Oelrich et al., 2020) or the candidate models are misspecified (Yang & Zhu, 2018). This phenomenon is called *overconfidence* and amounts to spuriously selecting one of the compared models as the only probable model in the set \mathcal{M} (Oelrich et al., 2020; Yang & Zhu, 2018). What is more, a moderate change in the data often suffices for the PMPs to shift and concentrate on another model. This behavior—large jumps of PMPs when the data is slightly altered—motivates the term “dance of Bayes factors” (Oelrich et al., 2020). Accordingly, our predictive mixture method could reduce overconfidence through its “simulation-based correction” of observed PMPs and its stronger emphasis on uncertainty.

Higher-Order Uncertainty Cheeseman (1985) argues that single probability assignments are (often) sufficient for reasoning and deciding under uncertainty. Nevertheless, various generalizations of probability have been suggested over the years, such as the theory of random sets (Matheron, 1975; Nguyen, 2008) or (generalized) evidence theory (Deng, 2015; Shafer, 1976). Roughly speaking, these ideas conceptualize higher-order uncertainty as *sets of distributions* or *distributions over distributions* (see Hüllermeier & Waegeman, 2021). Our method is inspired by these ideas and follows the premise that a single probability assignment is not sufficient for all practical purposes, as witnessed, for example, by the dangers of overconfident model choice (Oelrich et al., 2020; Yang & Zhu, 2018).

Closed vs. Open World Evidence theory (Deng, 2015; Shafer, 1976) differentiates between a *closed world* and an *open world* assumption. This terminology is commonly acknowledged in statistical inference (Bernardo & Smith, 2009; Yao et al., 2018), where it describes fundamentally different settings for model comparison. In a closed world setting, the “true” data generating model M_* is assumed to be one of the considered models, $M_* \in \mathcal{M}$. Differently, in the open world setting, the “true” model may not be a member of \mathcal{M} or may not even exist.

Table 1: Overview of the meta-uncertainty framework for J models M_1, \dots, M_J . We simulate K data sets $\{\overset{\circ}{y}^{(k)}\}$ with N observations from each of the models’ prior predictive distributions.

Level	Core concepts	Core implications
Level 1	Posterior model probabilities $\overset{\circ}{\pi} = p(M \overset{\circ}{y})$	Consistency of PMPs: $N \rightarrow \infty \xrightarrow{\text{a.s.}} \overset{\circ}{\pi} \rightarrow \mathbb{I}_{M_*}$ (Prop. 1)
Level 2	Model-implied PMP distributions $\{\overset{\circ}{\pi}^{(k)}\}$	Approach PMP sampling distribution for $K \rightarrow \infty$
Level 3	Meta models on PMP distributions $p_j(\tau_j \{\overset{\circ}{\pi}\}_j)$ Predictive mixture $f(\tilde{\pi}) = \sum \overset{\circ}{\pi}_j p_j(\tilde{\pi} \{\overset{\circ}{\pi}\}_j)$	Epistemic uncertainty of meta models fades for $K \rightarrow \infty$ Retain level 2 uncertainty \rightarrow ease overconfidence (Prop. 2)





Figure 3: **Experiment 1.** For each $N \in \{5, 10, 100\}$, we simulate $K = 100$ data sets and compute the PMPs under each model (red scatter). The pushforward likelihood distribution of the expected value for each meta-distribution is illustrated as a density plot in the *viridis* palette. The density in the row $N = 100$ is heavily concentrated at the respective vertex M_* and thus hardly visible.

In the context of deep learning, the relaxation of the closed world assumption is known as “out-of-distribution detection” (DeVries & Taylor, 2018; Hendrycks & Gimpel, 2018; Liang et al., 2020). In the Bayesian literature, the problem is framed as “model misspecification” and has been studied both in the context of standard and generalized Bayesian inference (Giummolè et al., 2019; Grünwald, Van Ommen, et al., 2017; Knoblauch et al., 2019; Masegosa, 2020; Thomas & Corander, 2019) as well as in the context of simulation-based inference (Alquier & Ridgway, 2019; Frazier & Drovandi, 2021; Frazier et al., 2020; Hermans et al., 2021; Schmitt et al., 2022; Zhang & Gao, 2020).

Our method can be viewed as a contribution to this growing literature, as it explores a way to integrate information from a closed world (i.e., simulations from a finite set of candidate models) and an open world (i.e., the real data obtained from a potentially unknown source).

5 EXPERIMENTS

In the following, we will illustrate the utility of our uncertainty framework in three BMC scenarios. For all experiments, we assume uniform model and meta-model priors: $p(M_j) = p(\mathbb{M}_j) = 1/J$. Further, we use the default priors of the `brms` R package for the meta-model param-



(a) $\hat{\pi} = (0.38, 0.31, 0.32)$ (b) $\hat{\pi} = (0.10, 0.55, 0.36)$ (c) $\hat{\pi} = (0.05, 0.91, 0.03)$

Figure 4: **Experiment 1.** Predictive mixtures for three *observed* data sets with $N = 10$ observations each.

eters and draw $D = 2000$ posterior samples from each meta-model. All meta-models are represented by logistic normal distributions. The code is publicly available at <https://github.com/marvinschmitt/MetaUncertaintyPaper>.

5.1 Experiment 1: Toy Conjugate Model

Setup Building on a conjugate Bayesian linear regression model, we aim to characterize the uncertainty at levels 1–3 in a controlled setting. Thus, we assume three competing regression models $\mathcal{M} = \{M_1, M_2, M_3\}$, which follow the Normal-Inverse-Gamma conjugacy (Murphy, 2007),

$$\begin{aligned} \sigma^2 &\sim \Gamma^{-1}(a_0, b_0), \\ \beta &\sim \mathcal{N}(\mu_0, \sigma^2 \Lambda_0^{-1}), \\ y &\sim \mathcal{N}(\mathbf{X}\beta, \sigma^2 \mathbb{I}), \end{aligned} \quad (13)$$

where $\Gamma^{-1}(a, b)$ is the inverse gamma distribution with shape a and rate b , and $\mathcal{N}(\mu, \Sigma)$ is the multivariate normal distribution with location μ and covariance matrix Σ . All models share the prior parameters $a_0 = 1, b_0 = 1, \mu_0 = \mathbf{0}, \Lambda_0 = \text{diag}(5)$. While all models feature the predictor variables x_1, x_2 , each model has a *distinct* third predictor.

Results As expected, Figure 3 shows that there is substantial epistemic level 1 uncertainty in the resulting PMPs when the number of observations per data set is small (i.e., $N = 5$). As the number of observations increases, the epistemic uncertainty decreases: The PMPs gravitate towards the true data generating model, that is, the corresponding vertex M_* . The aleatoric variation in the data generating process of y through the nuisance parameter Λ propagates into level 2 uncertainty of the corresponding PMPs. This level 2 uncertainty manifests itself as a scatter across the probability simplex, which corresponds to the *Dance of Bayes Factors* (Oelrich et al., 2020). In practical applications, this simulation setup can help researchers understand the variation (i.e., *dance*; Oelrich et al., 2020) of PMPs due to aleatoric uncertainty in the data generating process.

In order to study meta-uncertainty, we fit meta-models on $K = 200$ simulated data sets from the prior predictive distribution for each $N \in \{5, 10, 100\}$. We aggregate the un-

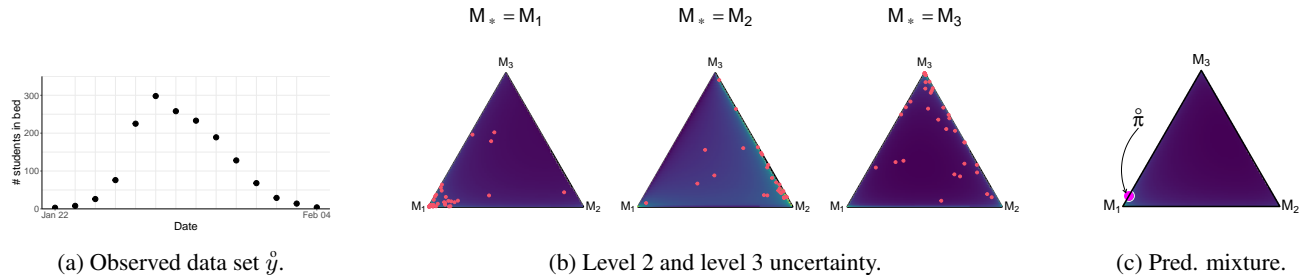


Figure 5: **Experiment 2.** Observed time series \hat{y} , level 2 and level 3 uncertainty, as well as the predictive mixture for the observed time series \hat{y} of interest. All density plots use a 4th root color mapping to the viridis palette to enhance visibility.

certainty over posterior draws by computing the pushforward likelihood density for the posterior expectation (see Figure 3). As expected, the uncertainty of the pushforward likelihood distribution depends on the uncertainties at level 1 and level 2: For $N = 100$, the vanishing level 1 and level 2 uncertainties imply that the posterior predictive distributions concentrate at the corresponding true model M_* .

For three *observed* data sets with $N = 10$ observations each, the predictive mixture distribution does in fact deviate from the observed PMPs (see Figure 4): In (a) the observed PMPs are uncertain which model is true, so the predictive mixture weighs all three components roughly equally. In (b) the observed PMPs rule out M_1 while being balanced between M_2 and M_3 . This pattern, in turn, manifests itself in the predictive mixture, which is strongly influenced by the components 2 and 3 and shows low density towards M_1 . In (c) the observed PMPs provide strong evidence for M_2 . However, even if M_2 was actually true, another researcher gathering new data for a replication study would be likely to find PMPs that provide *much less* evidence for M_2 , as indicated by the density of the predictive mixture.

5.2 Experiment 2: Likelihood-Based Inference

Setup This experiment studies meta-uncertainty in likelihood-based BMC for epidemiological compartmental models of influenza outbreaks using Stan (Stan Development Team, 2022) and bridge sampling (Gronau et al., 2020). The experiment and software code are adapted from Grinsztajn et al. (2021). The three assumed models M_1, M_2, M_3 differ with respect to their latent disease model (SIR vs. SEIR) and their likelihood (Poisson vs. Negative-Binomial). In this experiment, M_1 is the least complex model (SIR, Poisson), followed by M_2 (SEIR, Poisson), and the most complex model M_3 (SEIR, Negative-Binomial). The **Supplementary Material** contains implementation details and model configurations.

The observed data set \hat{y} stems from an influenza (H1N1) outbreak in 1978 at a British boarding school with $N = 763$ students, available in the `outbreaks` R pack-

age (Jombart et al., 2020). The data set contains the number of students in bed y at each day from January 22 to February 4, see Figure 5a.

Results The predictive mixture method on $K = 100$ simulated data sets illustrates that the model-implied PMPs in Figure 5b are aligned with the strong observed PMPs $\hat{\pi} = (0.92, <0.0001, 0.08)$, which supports the conclusion that the latter are *not* overconfident. Accordingly, the predictive mixture distribution is concentrated at the observed PMPs, too (see Figure 5c), and we would expect a replication study to show equally strong evidence for M_1 .

5.3 Experiment 3: Simulation-Based Inference

Setup This experiment will demonstrate that the predictive mixture method by itself is essentially independent of the complexity of the data generating process: Since it acts directly on PMPs and these PMPs are agnostic to the employed approximation method, the difficulty of calculating PMPs does not affect our predictive mixture method. Accordingly, we will compare three Hodgkin-Huxley simulator models of neural activity $\mathcal{M} = \{M_1, M_2, M_3\}$ with different parameter configurations $\Theta_{\mathcal{M}} = \{\theta_1, \theta_2, \theta_3\}$ (Hodgkin & Huxley, 1952; Pospischil et al., 2008) using the BayesFlow methodology (Radev et al., 2021).

Similarly to Radev et al. (2021), we train an unregularized evidential network for 25 epochs with 500 iterations (i.e., batch simulations) per epoch. Training details and loss history are described in the **Supplementary Material**. In the inference phase, the network estimates the parameters of an approximate Dirichlet distribution $f_{\phi}(M | y)$ over the PMPs for each simulated neural time series y . Here, ϕ denotes all parameters of the trained network. We derive a point estimate $\hat{p}(M | y)$ from the expected value of the Dirichlet density as $\mathbb{E}_{\pi \sim \text{Dir}(\alpha)}[\pi_j] = \alpha_j / \sum_j \alpha_j$ (Balakrishnan & Nevzorov, 2003).

Within our meta-uncertainty framework, we use the trained network to perform fast amortized model comparison on $K = 200$ simulated neural time series, effectively leverag-




Figure 6: **Experiment 3.** Model-implied distributions of PMPs on $K = 200$ simulated time series (level 2) are depicted as scatter across the probability simplices. The pushforward likelihood distributions for the meta-models’ expected posterior values are illustrated in the viridis palette.

ing the network’s ability to calculate the PMPs in line 5 of Algorithm 1 in a fraction of a second. The fast amortized inference of PMPs enables us to use the predictive mixture method on remarkably complex models, since it bypasses the bottleneck of fitting a large number of models (J models for each of the K data sets) to obtain the required set of model-implied PMPs. This renders the second phase of our method (i.e., fitting the meta-models) essentially independent of the complexity of the data generating process.

Results The model-implied PMP distributions (level 2) as well as the pushforward likelihoods for the expectations of the meta-models (level 3) are depicted in Figure 6. There is substantial level 2 uncertainty across all data generating models. However, it is worth noting that data from M_2 rarely ever features high probability for M_2 itself. In contrast, a large fraction of data sets simulated by M_1 attains high posterior model probability for M_1 . The estimated meta-models are able to reveal these patterns.

Finally, approximating the PMPs for an observed neural time series \hat{y} (see Figure 7a) yields $\hat{\pi} = (0.25, 0.30, 0.45)$. Our predictive mixture method equips the observed PMPs with information from realizations of the complex data generating process. Correspondingly, Figure 7 shows that the predictive mixture distribution deviates from the observed PMPs estimated using BayesFlow. While the observed PMPs suggest a slight tendency towards M_3 , the simulation-informed predictive mixture distribution unveils the strongly multimodal structure of the predicted PMPs for a new data set: Areas of high density in Figure 7c anticipate either a large PMP for M_1 , or else balanced PMPs for M_2 and M_3 , with low posterior probability for M_1 .

6 CONCLUSION

We have proposed a method to extend the methodological repertoire of probabilistic modeling workflows. Our method does not purport to provide a “better” point estimate than standard posterior model probabilities or Bayes factors. Using the generative property of Bayesian mod-

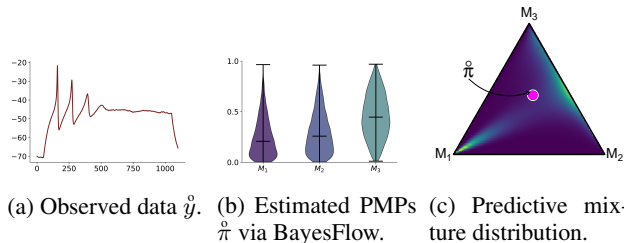


Figure 7: **Experiment 3.** For the observed data set \hat{y} , BayesFlow yields an estimated PMP distribution. Our predictive mixture augments this estimate with model-implied uncertainty, revealing the multi-model distribution of predicted PMPs for new data.

els, our *predictive mixture* method merely *augments* posterior model probabilities with simulation-based information. This information yields a more nuanced picture of uncertainty in finite-data settings and can prevent overconfident model selection in practice.

Meta-uncertainty can readily be derived for any (standard or adjusted) PMPs by replacing the marginal likelihood in line 5 of Algorithm 1, for example by conditional (aka. partial; Lotfi et al., 2022; O’Hagan, 1995), fractional (O’Hagan, 1995), or outlier-adjusted marginal likelihoods.

Future work could systematically investigate the connections between meta-models and related concepts from evidence theory (Deng, 2015; Shafer, 1976). In addition, the consequences of using different meta-models (i.e., probability distributions over probabilities) for the predictive mixture distribution can be further explored.

We hope that our framework can reveal intriguing empirical properties of Bayesian model probabilities and guide applied model building and model comparison workflows in virtually any research area. A selection of frequently asked questions we encountered prior to publication is answered in the FAQ section of the **Supplementary Material**.

Acknowledgments

We thank Lasse Elsemüller for his valuable feedback. We thank all reviewers for their constructive and thought-provoking feedback, which considerably improved the quality and clarity of the paper. MS was supported by the Cyber Valley Research Fund (grant number: CyVy-RF-2021-16). MS and PCB were supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2075 - 390740016 (the Stuttgart Cluster of Excellence SimTech). STR was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy — EXC-2181 - 390900948 (the Heidelberg Cluster of Excellence STRUCTURES).

References

- Aguilar, J. E., & Bürkner, P.-C. (2022). Intuitive Joint Priors for Bayesian Linear Multilevel Models: The R2-D2-M2 prior.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6), 716–723.
- Alquier, P., & Ridgway, J. (2019). Concentration of tempered posteriors and of their variational approximations. *arXiv:1706.09293*.
- Balakrishnan, N., & Nevzorov, V. B. (2003). *A primer on statistical distributions*. Wiley.
- Barron, A., Schervish, M. J., & Wasserman, L. (1999). The consistency of posterior distributions in nonparametric problems. *The Annals of Statistics*, 27(2).
- Bernardo, J. M., & Smith, A. F. M. (2009). *Bayesian Theory*.
- Betancourt, M. (2017). A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint*.
- Bürkner, P.-C. (2017). Brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28.
- Bürkner, P.-C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. *The R Journal*, 10(1), 395–411.
- Bürkner, P.-C., Gabry, J., & Vehtari, A. (2020). Approximate leave-future-out cross-validation for Bayesian time series models. *Journal of Statistical Computation and Simulation*, 90(14), 2499–2523.
- Bürkner, P.-C., Scholz, M., & Radev, S. T. (2022). Some models are useful, but how do we know which ones? Towards a unified Bayesian model taxonomy.
- Cheeseman, P. C. (1985). In defense of probability. *IJCAI*, 85, 1002–1009.
- Cranmer, K., Brehmer, J., & Louppe, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48), 30055–30062.
- Csilléry, K., Blum, M. G., Gaggiotti, O. E., & François, O. (2010). Approximate bayesian computation (ABC) in practice. *Trends in ecology & evolution*, 25(7), 410–418.
- Dehning, J., Zierenberg, J., Spitzner, F. P., Wibral, M., Neto, J. P., Wilczek, M., & Priesemann, V. (2020). Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions. *Science*, 369(6500).
- Dellaportas, P., Forster, J. J., & Ntzoufras, I. (2012). Joint Specification of Model Space and Parameter Space Prior Distributions. *Statistical Science*, 27(2).
- Deng, Y. (2015). Generalized evidence theory. *Applied Intelligence*, 43(3), 530–543.
- DeVries, T., & Taylor, G. W. (2018). Learning Confidence for Out-of-Distribution Detection in Neural Networks.
- Etz, A., & Vandekerckhove, J. (2018). Introduction to Bayesian Inference for Psychology. *Psychonomic Bulletin & Review*, 25(1), 5–34.
- Fong, E., & Holmes, C. C. (2020). On the marginal likelihood and cross-validation. *Biometrika*, 107(2), 489–496.
- Frazier, D. T., & Drovandi, C. (2021). Robust Approximate Bayesian Inference With Synthetic Likelihood. *Journal of Computational and Graphical Statistics*, 30(4), 958–976.
- Frazier, D. T., Robert, C. P., & Rousseau, J. (2020). Model misspecification in approximate Bayesian computation: Consequences and diagnostics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(2), 421–444.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (Third edition). CRC Press.
- Giummolè, F., Mameli, V., Ruli, E., & Ventura, L. (2019). Objective bayesian inference with proper scoring rules. *Test*, 28(3), 728–755.
- Grinsztajn, L., Semenova, E., Margossian, C. C., & Riou, J. (2021). Bayesian workflow for disease transmission modeling in Stan. *Statistics in Medicine*, 40(27), 6209–6234.
- Gronau, Q. F., Singmann, H., & Wagenmakers, E.-J. (2020). Bridgesampling: An R Package for Estimating Normalizing Constants. *Journal of Statistical Software*, 92(10).
- Grünwald, P., Van Ommen, T., et al. (2017). Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis*, 12(4), 1069–1103.
- Hendrycks, D., & Gimpel, K. (2018). A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks.
- Hermans, J., Delaunoy, A., Rozet, F., Wehenkel, A., & Louppe, G. (2021). Averting a crisis in simulation-based inference. *arXiv preprint arXiv:2110.06581*.
- Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), 500–544.
- Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods. *Machine Learning*, 110(3), 457–506.
- Jombart, T., Frost, S., Nouvellet, P., Campbell, F., & Sudre, B. (2020). *Outbreaks: A Collection of Disease Outbreak Data*.

- Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Knoblauch, J., Jewson, J., & Damoulas, T. (2019). Generalized variational inference: Three arguments for deriving new posteriors. *arXiv preprint arXiv:1904.02063*.
- Konishi, S. (1996). Generalised information criteria in model selection. *Biometrika*, 83(4), 875–890.
- Konishi, S., & Kitagawa, G. (2008). *Information criteria and statistical modeling*. Springer.
- Kulesa, A., Krzywinski, M., Blainey, P., & Altman, N. (2015). Sampling distributions and the bootstrap. *Nature Methods*, 12(6), 477–478.
- Liang, S., Li, Y., & Srikant, R. (2020). Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks.
- Lindley, D. V. (1957). A Statistical Paradox. *Biometrika*, 44(1-2), 187–192.
- Lotfi, S., Izmailov, P., Benton, G., Goldblum, M., & Wilson, A. G. (2022). Bayesian Model Selection, the Marginal Likelihood, and Generalization. *arXiv:2202.11678 [cs, stat]*.
- Masegosa, A. (2020). Learning under model misspecification: Applications to variational and ensemble methods. *Advances in Neural Information Processing Systems*, 33, 5479–5491.
- Matheron, G. (1975). *Random sets and integral geometry*. Hoboken: John Wiley & Sons.
- Meng, X.-L., & Schilling, S. (2002). Warp Bridge Sampling. *Journal of Computational and Graphical Statistics*, 11(3), 552–586.
- Murphy, K. (2007). Conjugate Bayesian analysis of the Gaussian distribution.
- Neal, R. M. (2011). *MCMC using Hamiltonian dynamics*.
- Nguyen, H. T. (2008). On random sets and belief functions. In *Classic works of the dempster-shafer theory of belief functions* (pp. 105–116). Springer.
- Oelrich, O., Ding, S., Magnusson, M., Vehtari, A., & Villani, M. (2020). When are Bayesian model probabilities overconfident? *arXiv:2003.04026 [math, stat]*.
- O’Hagan, A. (1995). Fractional Bayes Factors for Model Comparison. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 99–118.
- O’Hagan, A., Kendall, M. G., & Stuart, A. (1994). *Kendall’s Advanced theory of statistics: Volume 2B, Bayesian inference*. Edward Arnold.
- Piironen, J., Paasiniemi, M., & Vehtari, A. (2020). Projective inference in high-dimensional problems: Prediction and feature selection. *Electronic Journal of Statistics*, 14(1).
- Pospischil, M., Toledo-Rodriguez, M., Monier, C., Pwowska, Z., Bal, T., Frégnac, Y., Markram, H., & Destexhe, A. (2008). Minimal Hodgkin–Huxley type models for different classes of cortical and thalamic neurons. *Biological Cybernetics*, 99(4-5), 427–441.
- Radev, S. T., D’Alessandro, M., Mertens, U. K., Voss, A., Köthe, U., & Bürkner, P.-C. (2021). Amortized bayesian model comparison with evidential deep learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Radev, S. T., Mertens, U. K., Voss, A., Ardizzone, L., & Köthe, U. (2020). BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*.
- Radev, S. T., Schmitt, M., Pratz, V., Picchini, U., Köthe, U., & Bürkner, P.-C. (2023). JANA: Jointly Amortized Neural Approximation of Complex Bayesian Models. *arXiv preprint arXiv:2302.09125*.
- Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. (2016). Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*, 2, e55.
- Schad, D. J., Nicenboim, B., Bürkner, P.-C., Betancourt, M., & Vasisht, S. (2021). Workflow Techniques for the Robust Use of Bayes Factors.
- Schmitt, M., Bürkner, P.-C., Köthe, U., & Radev, S. T. (2022). Detecting Model Misspecification in Amortized Bayesian Inference with Neural Networks.
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton Univ. Press.
- Stan Development Team. (2022). Stan Modeling Language Users Guide and Reference Manual [https://mc-stan.org].
- Thomas, O., & Corander, J. (2019). Diagnosing model misspecification and performing generalized Bayes’ updates via probabilistic classifiers. *arXiv preprint arXiv:1912.05810*.
- van der Vaart, A. W. (2000). *Asymptotic statistics* (Vol. 3). Cambridge university press.
- Van Dongen, S. (2006). Prior specification in Bayesian statistics: Three cautionary tales. *Journal of Theoretical Biology*, 242(1), 90–100.
- Vehtari, A., & Lampinen, J. (2002). Bayesian Model Assessment and Comparison Using Cross-Validation Predictive Densities. *Neural Computation*, 14(10), 2439–2468.
- Vehtari, A., & Ojanen, J. (2012). A survey of Bayesian predictive methods for model assessment, selection and comparison. *Statistics Surveys*, 6(none).
- Vehtari, A., Simpson, D., Gelman, A., Yao, Y., & Gabry, J. (2021). Pareto Smoothed Importance Sampling.
- Wagner, C. E., Saad-Roy, C. M., & Grenfell, B. T. (2022). Modelling vaccination strategies for covid-19. *Nature Reviews Immunology*, 22(3), 139–141.

- Yang, Z., & Zhu, T. (2018). Bayesian selection of misspecified models is overconfident and may cause spurious posterior probabilities for phylogenetic trees. *Proceedings of the National Academy of Sciences*, *115*(8), 1854–1859.
- Yao, Y., Vehtari, A., Simpson, D., & Gelman, A. (2018). Using Stacking to Average Bayesian Predictive Distributions (with Discussion). *Bayesian Analysis*, *13*(3).
- Zhang, F., & Gao, C. (2020). Convergence rates of variational posterior distributions. *The Annals of Statistics*, *48*(4).
- Zitzmann, S., Helm, C., & Hecht, M. (2021). Prior Specification for More Stable Bayesian Estimation of Multilevel Latent Variable Models in Small Samples: A Comparative Investigation of Two Different Approaches. *Frontiers in Psychology*, *11*, 611267.

A Frequently Asked Questions (FAQ)

How can I reproduce the results?

The full software code (R and Python) to reproduce the minimal example in the Methods section, as well as all experiments, is available in the public repository at <https://github.com/marvinschmitt/MetaUncertaintyPaper>. The code structure is detailed in Section B in the Appendix.

Is meta-uncertainty better than standard posterior model probabilities?

We do not intend to provide a "better" point estimate than standard posterior model probabilities or Bayes Factors. Instead, we *enrich* observed PMPs by embedding them into the probabilistic context of the compared models to enable more realistic interpretations in practice. The epistemic uncertainty in PMPs vanishes for the infinite data case $N \rightarrow \infty$, and this is a desired property known as consistency. However, in the *finite data* case, we might encounter *misleadingly* extreme preference for one model, i.e., overconfidence. The paramount importance of overconfidence is well-known and acknowledged in the scientific community (cf. Oelrich et al., 2020; Yang & Zhu, 2018). A replication study might come to an utterly different result, but both observed results are compatible with the models. Yet, we cannot learn this from any single study, but it is clearly visible through the level 2 uncertainty of PMPs across many simulated data sets. Thus, our method mitigates overconfidence by answering the question "which PMPs do we expect in a replication study?" with a principled Bayesian approach. In all three experiments, we show how our method yields additional insights beyond single observed PMPs.

How is meta-uncertainty related to modified marginal likelihoods?

The idea of conditional marginal likelihoods to use parts of the data to pre-condition marginal likelihoods reaches back to O'Hagan (1995), and it is not an alternative to our method. In fact, our pipeline readily works for modified computations of PMPs by adjusting line 5 in Algorithm 1, for example through conditional (partial) marginal likelihoods, fractional marginal likelihoods, or outlier-adjusted marginal likelihoods.

How is meta-uncertainty related to Bayesian model averaging?

Bayesian model averaging (BMA) is concerned with parameter estimation under several competing models, and thus not an alternative to our method which is exclusively concerned with model comparison. However, our predictive mixture distribution might serve as a basis for averaged parameter estimates (replacing PMPs in Bayesian model averaging), and this might be a promising future research endeavor.

How is meta-uncertainty related to the Bayesian information criterion?

The Bayesian information criterion (BIC) is neither fully Bayesian, nor an information criterion in a strict sense. Instead, it is a (possibly crude) approximation of the marginal likelihood, which lies at the core of our method

How does the uncertainty change with more samples?

As the number of observations per data set grows arbitrarily large ($N \rightarrow \infty$), the epistemic uncertainty in the PMP vectors vanishes. For an increasing number of simulated data sets ($K \rightarrow \infty$), we argue that (i) the level 2 distribution of PMPs across data sets approaches the sampling distribution (cf. Section 3.1, "Relation to Sampling Distributions"); and (ii) the epistemic uncertainty of the meta models vanishes (cf. Section 3.2, "Level 3 Uncertainty"). Furthermore, the levels of the framework are briefly outlined in Table 1.

How do outliers in the data influence the results?

We acknowledge that outliers in the data y can lead to overconfident PMPs for single data sets. Yet, if outliers only occur occasionally in few data sets, the level 2 distribution should mirror this behavior, and the meta model should reflect it too. Thus, our method naturally puts data sets with outliers into an appropriate probabilistic context of the models without altering the data y by removing said outliers.

Does the method work for models with high-dimensional parameter spaces?

We agree that settings with high parameter dimensions constitute a paramount application for (Bayesian) model selection, and established Bayesian methods (ABC, MCMC) are infeasible in high dimensions with large data sets. As discussed in Section 5.3, our method is essentially independent of the data dimensionality, and the method to obtain PMPs (line 5 in Algorithm 1) is the actual bottleneck. Recent advances in simulation-based inference (SBI) allow model comparison in high dimensions through special neural network architectures, such as BayesFlow (Radev et al., 2020). As demonstrated in Experiment 3, our method is fully compatible with these state-of-the-art neural approaches and thus capable of augmenting model comparison workflows in high dimensions. Furthermore, Experiment 3 shows that our method has a negligible computational cost when an amortized PMP approximator (such as BayesFlow) is available, which further facilitates the integration of our methods to simulation-based inference workflows.

Does the method work for large-scale Bayesian variable selection?

The current state-of-the-art method for large-scale Bayesian variable selection is projection predictive inference, while using PMPs for this purpose is not advisable (Piiroinen et al., 2020). The most important application of PMPs, and thus also of our method, are comparisons of qualitatively different models arising from competing scientific theories. While the parameter spaces may be high-dimensional (and simulation-based inference might be employed), the model space is usually low-dimensional (i.e., only few models are being compared).

B CODE

Code for all methods and experiments is available in the public repository at <https://github.com/marvinschmitt/MetaUncertaintyPaper>. The code is structured as outlined in Figure 8. To facilitate reproducibility, we use seeds for the random number generators throughout the analyses. Note that the R code is structured as a package, thus requiring a local installation with subsequent loading via `library(MetaUncertaintyPaper)`.




Figure 8: Outline of the experimental code.




Figure 9: **Experiment 1.** Pushforward densities of cluster means from the posterior draws of the meta models on simulated data.

C EXPERIMENTS: IMPLEMENTATION DETAILS AND ADDITIONAL RESULTS

C.1 Experiment 1: Conjugate Model

The conjugate construction yields an analytic posterior distribution for the inference parameters (Murphy, 2007), which again follow a Normal-Inverse-Gamma distribution $\mathcal{N}\text{-}\Gamma^{-1}(a_n, b_n, \mu_n, \Lambda_n)$ with

$$\begin{aligned}
 a_n &= a_0 + \frac{n}{2}, \\
 b_n &= b_0 + \frac{1}{2}(y^\top y + \mu_0^\top \Lambda_0 \mu_0 - \mu_n^\top \Lambda_n \mu_n), \\
 \mu_n &= (\mathbf{X}^\top \mathbf{X} + \Lambda_0)^{-1}(\mathbf{X}^\top \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top y + \Lambda_0 \mu_0) \\
 \Lambda_n &= \mathbf{X}^\top \mathbf{X} + \Lambda_0.
 \end{aligned} \tag{14}$$

Consequently, the marginal likelihood can be directly calculated (Murphy, 2007; O’Hagan et al., 1994) to obtain the analytic expression

$$p(y | M) = \frac{1}{(2\pi)^{n/2}} \sqrt{\frac{\det(\Lambda_0)}{\det(\Lambda_n)}} \frac{b_0^{a_0}}{b_n^{a_n}} \frac{\Gamma(a_n)}{\Gamma(a_0)}. \tag{15}$$

C.2 Experiment 2: Likelihood-based Inference

Following the implementation of Grinsztajn et al. (2021), the data generating process for the influenza model is formulated either as an SIR model or as an SEIR model

$$\begin{aligned}
 \text{SIR:} \quad & \frac{dS}{dt} = -\beta S \frac{I}{N}, & \frac{dI}{dt} &= \beta S \frac{I}{N} - \gamma I, & \frac{dR}{dt} &= \gamma I, \\
 \text{SEIR:} \quad & \frac{dS}{dt} = -\beta S \frac{I}{N}, & \frac{dE}{dt} &= \beta S \frac{I}{N} - \eta E, & \frac{dI}{dt} &= \eta E - \gamma I, & \frac{dR}{dt} &= \gamma I,
 \end{aligned} \tag{16}$$

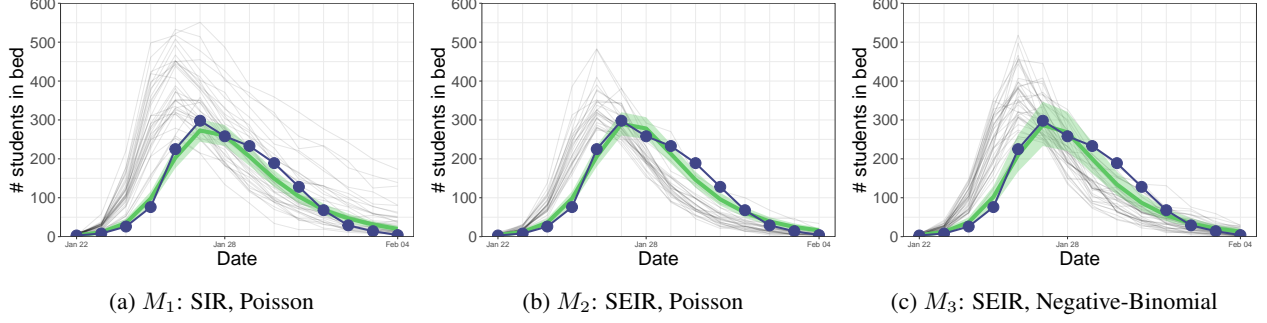


Figure 10: **Experiment 2.** Simulated data sets $\{\tilde{y}^s\}$ from the prior predictive distribution (gray, thin). Real observed data set \tilde{y}^o of interest (blue, thick). Posterior predictive distributions $p(\tilde{y} | \tilde{y}^o, M_j)$ for each model M_j , conditional on the observed data: median (green line) and symmetric 89% credible interval (green ribbon).

with susceptible (S), exposed (E), infected (I), and recovered (R) population compartments, respectively. The likelihood function represents the stochastic relationship between the latent counts in the I compartment and the actually reported infections (i.e., students in bed) y as either a Negative-Binomial distribution

$$p(y | \theta) = \text{Negative-Binomial}(y | I(t), \phi) \quad (17)$$

or as a Poisson distribution

$$p(y | \theta) = \text{Poisson}(y | I(t)) \quad (18)$$

where θ expresses all parameters in the respective model, e.g., $\theta = \{\beta, \gamma, \eta, \phi\}$ for a latent SEIR model with Negative-Binomial likelihood. The prior distributions are specified as

$$\beta \sim \mathcal{N}^+(\mu_\beta, \sigma_\beta), \quad \gamma \sim \mathcal{N}^+(\mu_\gamma, \sigma_\gamma), \quad \eta \sim \mathcal{N}^+(\mu_\eta, \sigma_\eta), \quad \phi \sim \mathcal{N}^+(\mu_\phi, \sigma_\phi), \quad (19)$$

where \mathcal{N}^+ denotes the (positive) half-Normal distribution truncated at 0. The model configurations are outlined in Table 2.

The observed data \tilde{y}^o from the actual outbreak in 1978 yield the posterior model probabilities $\hat{\pi} = (0.92, <0.0001, 0.08)$. Figure 10 shows the posterior predictive distribution of the Bayesian model on the data y (not the meta models, which act on PMPs π). While all posterior predictive distributions heuristically suggest a somewhat good fit, the relatively low number of time steps (i.e., $T = 14$) might explain the strong evidence for the simplest model M_1 (SIR+Poisson), effectively illustrating Occam's razor. We hypothesize that longer time series (i.e., more observations) might support more signal, and might thus have the potential to render the more complex models M_2 and M_3 more likely.

The model-implied PMPs (level 2) show that (i) simulated data from M_1 typically shows high PMPs for M_1 ; (ii) data emerging from M_2 mostly features low probability for M_1 , moderate to large PMPs for M_2 and moderate probability for M_3 ; and (iii) PMPs for data from M_3 either concentrate at extremely large PMPs for M_3 , or scatter unsystematically across most parts of the probability simplex. As expected, the meta models (level 3) capture this pattern.

Employing our proposed method in the meta-uncertainty framework, the resulting *predictive mixture* distribution shows high density towards M_1 . This is consistent with both the observed PMP ($\hat{\pi}_1 = 0.92$) as well as the model-implied PMPs under M_1 , which are usually strongly favoring M_1 as well. The alignment of high observed PMP for M_1 on the one hand with the model-implied PMP distribution under M_1 on the other hand gives rise to the belief that the observed PMPs are *not* overconfident. In other words, we expect PMPs for new data sets to feature equally large support for M_1 , and we believe that the observed PMPs are not a mere artifact of the dance of Bayes Factors (Oelrich et al., 2020).

Table 2: Competing models in **Experiment 2**. The models differ with respect to their latent epidemiological compartmental model (CM) and the functional form of the likelihood.

	Latent CM	Likelihood	μ_β	σ_β	μ_γ	σ_γ	μ_η	σ_η	μ_ϕ	σ_ϕ
M_1	SIR	Poisson	2	0.1	0.4	0.1	—	—	—	—
M_2	SEIR	Poisson	3	0.1	0.5	0.1	3	0.1	—	—
M_3	SEIR	Negative-Binomial	3	0.1	0.5	0.1	3	0.1	100	1

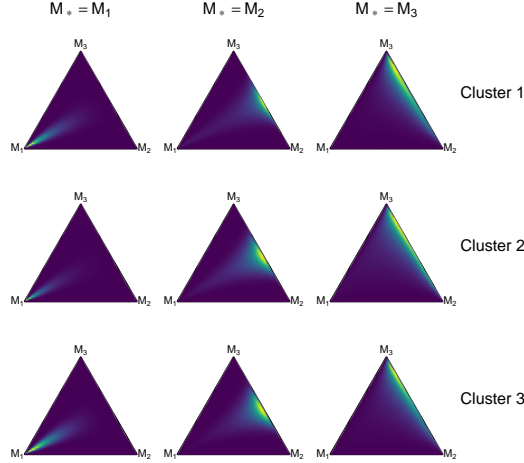


Figure 11: **Experiment 3.** Pushforward densities of cluster means from the posterior draws of the meta models on simulated data.

C.3 Experiment 3: Simulation-based Inference

The forward model is adapted from Radev et al. (2021) under MIT license. It is formulated as a set of five ordinary differential equations (ODEs) describing how the neuron membrane potential $V(t)$ unfolds in time as a function of an injected current $I_{inj}(t)$, and ion channels properties. The change in membrane potential is defined by the membrane stochastic differential equation (SDE)

$$C \frac{dV}{dt} = -I_L - I_{Na} - I_K - I_M + I_{inj} + \sigma\eta(t) \quad (20)$$

where C is the specific membrane capacitance, $\sigma\eta(t)$ the intrinsic neural noise, and the I_j s denote the ionic currents flowing through channels, such that:

$$I_L = g_L(V - E_L) \quad (21)$$

$$I_{Na} = \bar{g}_{Na} m^3 h (V - E_{Na}) \quad (22)$$

$$I_K = \bar{g}_K n^4 (V - E_K) \quad (23)$$

$$I_M = \bar{g}_M p (V - E_M). \quad (24)$$

Here, g_L is the leak conductance, while \bar{g}_{Na} , \bar{g}_K , \bar{g}_M are the sodium, potassium and M-type channel maximum conductances, respectively. E_L , E_{Na} and E_K denote the leak equilibrium potential, the sodium and potassium reversal potentials, respectively. In particular, g_L is assumed constant through time, whilst the other conductances vary over time. Consistently, (m, h, n, p) indicates the vector of the state variables accounting for ion channel gating kinetics evolving according to the following set of ODEs:

$$\frac{di}{dt} = \alpha_i(V)(1 - i) - \beta_i(V)i \quad (25)$$

$$\frac{dp}{dt} = \frac{p_\infty(V) - p}{\tau_p(V)} \quad (26)$$

where $i \in \{m, h, n\}$, and $\alpha_i(V)$, $\beta_i(V)$, $p_\infty(V)$ and $\tau_p(V)$ are nonlinear functions of the membrane potential (see Pospischil et al., 2008, for details).

In our simulated experiment, we treat conductances g_L , \bar{g}_{Na} , \bar{g}_K and \bar{g}_M as free parameters, and consider different neuronal models based on different parameter configurations. It is also assumed that such configurations allow to affect the span of the possible firing patterns attainable by each model. In particular, we consider 3 models $\mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}$ defined

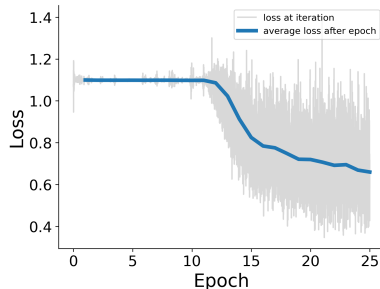


Figure 12: **Experiment 3.** Loss history of the training phase with BayesFlow.

by the parameter sets $\theta_1 = (\bar{g}_{Na}, \bar{g}_K)$, $\theta_2 = (\bar{g}_{Na}, \bar{g}_K, \bar{g}_M)$, $\theta_3 = (\bar{g}_{Na}, \bar{g}_K, \bar{g}_M, g_L)$. The parameter priors are listed in Table 3.

For training, we use the Adam optimizer with an initial learning rate of 10^{-4} . We also apply an exponential learning rate decay schedule with a decay parameter $\gamma = 0.99$ and a staircase reduction every 1000 backpropagation steps. The architecture follows the evidential neural network scheme in Radev et al. (2021) with 3 dense layers with ReLU activation and 128 units each, and a softplus activation in the output layer. The training took 5h on a consumer-grade CPU, with a drastically reduced expected training time on a GPU. The loss history is depicted in Figure 12.

Table 3: Prior distributions of the parameters in **Experiment 3**. If a cell contains a scalar value instead of a distribution, the respective parameter is fixed to that value in the model.

	\bar{g}_{Na}	\bar{g}_K	\bar{g}_M	g_L
\mathcal{M}_1	Uniform(1.5, 15)	Uniform(1, 3)	0.07	0.1
\mathcal{M}_2	Uniform(1.5, 15)	Uniform(1, 3)	Uniform(0.02, 0.15)	0.1
\mathcal{M}_3	Uniform(1.5, 15)	Uniform(1, 3)	Uniform(0.02, 0.15)	Uniform(0.05, 0.15)

D PROOFS

D.1 Proof of Proposition 1.

Proof. Assume a closed world setting ($M_* \in \mathcal{M}$) and at least one simulated data set ($K > 0$), as well as the regularity conditions (Barron et al., 1999) that (i) the prior does not put high mass near distributions with very rough densities; and (ii) the prior puts positive mass in Kullback-Leibler neighborhoods of the true distribution.

For an arbitrarily large number of observations $N \rightarrow \infty$, the posterior model probabilities π (level 1) converge almost surely to one-hot encoded vectors (Barron et al., 1999):

$$\pi \xrightarrow{\text{a.s.}} \mathbb{I}_{M_*} \iff \pi_j = \begin{cases} 1 & \text{if } M_j = M_* \\ 0 & \text{else} \end{cases} \quad (27)$$

For a set of K simulated data sets $\{y^{(k)}\}_{k=1}^K$ with $N \rightarrow \infty$ observations each, it follows from Equation 27 that the resulting set of model-implied PMPs $\{\hat{\pi}^{(k)}\}_{k=1}^K$ is a set of one-hot encoded vectors $\{\mathbb{I}_{M_j}^{(k)}\}_{k=1}^K$ with a 1 at the respective true model index and 0 otherwise. After grouping $\{\hat{\pi}^{(k)}\}_{k=1}^K$ by the respective true data generating model, we are left with J groups G_1, \dots, G_J . Recall that all simulated data sets underlying the PMPs in group G_j stem from the corresponding true model M_j . Thus, according to Equation 27, all PMPs in group G_j equal \mathbb{I}_{M_j} . Specifically, the variance of PMPs within each group G_j equals zero, and there is no level 2 uncertainty.

Consequently, fitting a meta-model \mathbb{M}_j on PMPs from group G_j yields a Dirac distribution, which in turn implies no level 3 uncertainty. Finally, the observed data set y° features $N \rightarrow \infty$ observations as well, implying $\hat{\pi}^\circ = \mathbb{I}_{M_*}$ (recall that we assume a closed world). Thus, both the observed PMPs as well as the mixture components (i.e., meta model predictive distributions) equal \mathbb{I}_{M_*} , resulting in a one-hot encoded predictive mixture $f(\hat{\pi}^\circ)$ in the limit, which is consistent by definition. \square

D.2 Proof of Proposition 2.

Proof. Assume that the predictive mixture does not collapse to a one-hot encoded vector: $f(\hat{\pi}) \neq \mathbb{I}_{M_j}$. In our framework, this translates to the statement that there is remaining uncertainty at level 2 and level 3. Given (spuriously) extreme observed PMPs for an arbitrary but fixed model M_q , $\hat{\pi}^\circ \approx \mathbb{I}_{M_q}$, the predictive mixture reduces approximately to the q^{th} component: $f(\hat{\pi}) \approx \hat{\pi}_q p_q(\hat{\pi} | \{\hat{\pi}\}_j)$. Since this component has non-vanishing variance by our initial assumption, the predictive mixture $f(\hat{\pi})$ has non-vanishing variance as well. Even if the observed PMPs were one-hot encoded—i.e., $\hat{\pi}^\circ = \mathbb{I}_{M_q}$ —, the predictive mixture consists exclusively of the q^{th} component, $f(\hat{\pi}) = 1 \cdot \underbrace{p_q(\hat{\pi} | \{\hat{\pi}\}_j)}_{\text{level 2 variance}}$, which still contains variance from level 2 as per the assumption. \square