

Master's programme in Computer, Communication, and Information Sciences

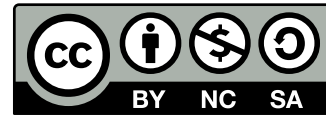
# Efficient estimation of selection-induced bias in Bayesian model selection

---

**Yann McLatchie**

© 2023

This work is licensed under a [Creative Commons](https://creativecommons.org/licenses/by-nc-sa/4.0/) “Attribution-NonCommercial-ShareAlike 4.0 International” license.



---

**Author** Yann McLatchie

---

**Title** Efficient estimation of selection-induced bias in Bayesian model selection

---

**Degree programme** Computer, Communication, and Information Sciences

---

**Major** Machine Learning, Data Science, and Artificial Intelligence

---

**Supervisor and advisor** Prof. Aki Vehtari

---

**Date** September 15, 2023

**Number of pages** 36+8

**Language** English

---

**Abstract**

Model selection aims to identify a sufficiently well performing model that is possibly simpler than the most complex model among a pool of candidates. However, the decision-making process itself can inadvertently introduce non-negligible bias when the cross-validation estimates of predictive performance are marred by excessive noise. In finite data regimes, cross-validated estimates can encourage the statistician to select one model over another when it is not actually better for future data. While this bias remains negligible in the case of few models, when the pool of candidates grows, and model selection decisions are compounded (as in forward search), the expected magnitude of selection-induced bias is likely to grow too. This paper introduces an efficient approach to estimate and correct selection-induced bias based on order statistics. Numerical experiments demonstrate the reliability of our approach in estimating both selection-induced bias and over-fitting along compounded model selection decisions, with specific application to forward search. This work represents a light-weight alternative to more computationally expensive approaches to correcting selection-induced bias, such as nested cross-validation and the bootstrap. Our approach rests on several theoretic assumptions, and we provide a diagnostic to help understand when these may not be valid and when to fall back on safer, albeit more computationally expensive approaches. The accompanying code facilitates its practical implementation and fosters further exploration in this area.

---

**Keywords** selection-induced bias, Bayesian model selection, cross-validation, bias correction

---

## **Preface**

I would like to take this opportunity to express my gratitude to the individuals whose support and contributions have been instrumental in the completion of this thesis.

First and foremost, I am immensely thankful to my supervisor, Prof. Aki Vehtari, for his guidance and support throughout the research process. His valuable insights, constructive feedback, and continuous encouragement have greatly enriched the quality of this work.

I would also like to acknowledge the constructive comments provided by Andrew R. Johnson, Frank Weber, Anna E. Riha, Noa Kallionen, and Leevi Lindgren.

To those mentioned above, as well as to all others who have supported me along this journey, I extend my thanks. Your help and encouragement have played an integral role in the completion of this endeavour.

Otaniemi, September 15, 2023

Yann McLatchie

# Contents

<b>Abstract</b>	<b>3</b>
<b>Preface</b>	<b>4</b>
<b>Contents</b>	<b>5</b>
<b>1 Introduction</b>	<b>6</b>
1.1 Relation to previous work . . . . .	6
1.2 Alternatives to model selection . . . . .	8
1.3 Structure of this paper . . . . .	9
<b>2 Predictive model assessment, selection, and comparison</b>	<b>9</b>
2.1 Predictive model assessment . . . . .	9
2.2 Selection-induced bias . . . . .	10
2.3 When is one model clearly better than another? . . . . .	11
<b>3 Bias grows with the number of candidate models</b>	<b>13</b>
3.1 Empirical demonstration . . . . .	13
3.2 Our proposal: modelling the elpd differences . . . . .	15
3.3 Diagnosing the theoretic assumptions . . . . .	17
<b>4 Compounding the bias through forward search</b>	<b>18</b>
4.1 Correcting for selection-induced bias in forward search . . . . .	19
4.2 A simulated example . . . . .	21
4.3 Reducing selection-induced bias with priors . . . . .	22
4.4 Projection predictive inference . . . . .	24
4.5 Real world data . . . . .	25
<b>5 Recommendations</b>	<b>28</b>
<b>6 Discussion</b>	<b>29</b>
6.1 Limitations . . . . .	29
6.2 Future directions . . . . .	29
<b>A Connection between elpd difference and model averaging</b>	<b>37</b>
<b>B Empirical demonstration of order statistics</b>	<b>38</b>
<b>C Bias estimation under different assumed asymptotic cross-validated correlation regimes</b>	<b>39</b>
<b>D Sparsity-inducing priors</b>	<b>41</b>
<b>E Alternative stopping rules</b>	<b>43</b>

# 1 Introduction

In model selection, we are usually interested in identifying the most predictive model from a set of candidates. Specifically, we are interested in identifying a model from a collection of models whose out-of-sample predictive performance is best, either in the hope of generalising to future unseen data (Vehtari and Ojanen, 2012), or as a proxy for parameter recoverability (Scholz and Bürkner, 2022). When the number of predictors to consider is large, we might also be interested in achieving a smaller model that is capable of replicating the predictive behaviour of a larger model and use this instead for its improved interpretability, or to decrease data collection cost (Piironen et al., 2020). In finite data regimes, however, there can be significant uncertainty in differentiating between models of similar performance. Over-optimism in the utility estimate of the selected model is known as *selection-induced bias* (Stone, 1974). Alongside this, when we select a model which performs non-negligibly worse than the oracle asymptotically, we say that we have *over-fit*. The kernel of this work is an efficient and actionable selection-induced bias estimation based on order statistics. Concretely, we:

1. discuss over-optimism when making model selection decisions based on noisy cross-validated predictive metrics;
2. propose a lightweight selection-induced bias estimation and correction tool based on order statistics;
3. provide a diagnostic for our estimate to understand when it is liable to be unsafe, and more computationally expensive approaches may be warranted;
4. show empirically that correcting selection-induced bias can help expose non-negligible over-fitting; and,
5. apply our bias correction to real and simulated data examples where we are interested in choosing between multiple candidate models, and in forward search.

## 1.1 Relation to previous work

Predictive approaches to Bayesian model comparison have been qualitatively and quantitatively reviewed by Vehtari and Ojanen (2012) and Piironen and Vehtari (2017a) respectively. One popular method is choosing the maximum *a posteriori* (MAP) model, which reduces to choosing the maximum marginal likelihood model in the case of a uniform prior over the model space, and can be performed with Bayes factors (Jeffreys, 1998; Kass and Raftery, 1995). The computational instability of Bayes factors, and the difficulty they have sampling from the combinatorially exploding model space, however, is well-documented (Han and Carlin, 2001; Oelrich et al., 2020). In the context of variable selection, one might otherwise consider investigating the probability of inclusion of an individual predictor (Barbieri and Berger, 2004; Brown et al., 1998; Narisetty and He, 2014). Most contemporary work in this field has centred on defining intelligent priors over model space and model parameters (O’Hara and Sillanpää, 2009). However, Piironen et al. (2020) show how such approaches provide limited understanding of the predictive improvement expected by adding a predictor, and can fail when predictors are collinear. It has further been argued

that priors and model selection exist as separate components to a Bayesian workflow (Bernardo and Smith, 1994; Key et al., 1999; Vehtari and Ojanen, 2012): a view we hold ourselves.

A fully decision theoretical Bayesian model selection procedure was first outlined by Lindley (1968), wherein the statistician is implored to build the largest, most complex model possible which encompasses all uncertainty relating to the modelling task, and then to identify smaller models capable of closely replicating its behaviour. Leamer (1979) motivated the use of the Kullback–Leibler (KL) divergence to quantify the distance between models’ predictive distributions. Goutis (1998) and Dupuis and Robert (2003) went further, and proposed that one should directly *project* the reference model onto parameter sub-spaces so as to minimise the distance between the induced projected posterior predictive distribution and the posterior predictive distribution of the reference model: essentially fitting sub-models to the fit of the reference model. McLatchie et al. (2023) present practical recommendations for projection predictive model selection, and a more complete overview of recent advances therein. We are interested in cases where either the number of models is too small to justify the added complexity of projection predictive inference, or when we don’t know how to perform the projection efficiently (e.g. in Student- $t$  observational family models, and some spatial models).

As a computationally efficient, intuitive, and non-parametric alternative, and one of the most common in non-Bayesian literature, statisticians compare models based on their out-of-sample performance with either cross-validation (CV; Geisser and Eddy, 1979; Shao, 1993; Vehtari and Lampinen, 2002), or various information criteria (Burnham and Anderson, 2002; Gelfand and Ghosh, 1998; Laud and Ibrahim, 1995; Marriott et al., 2001; Spiegelhalter et al., 2002; Watanabe, 2013). Such procedures are able to produce an almost unbiased estimate of a model’s predictive ability (Watanabe, 2010), and have been reviewed extensively by Arlot and Celisse (2010), and by Gelman et al. (2014) in a Bayesian context. When using noisy predictive performance estimates such as cross-validation to select a best model, we may be over-optimistic in the utility estimate of the eventually selected model. This selection-induced bias has been discussed in depth by Ambroise and McLachlan (2002), Reunanen (2003), and Cawley and Talbot (2010). Cawley and Talbot (2010), in addition to demonstrating the potential risks of selection-induced bias, emphasised that this bias does not always lead to non-negligible over-fitting. In many cases, when one model is much better than others, or if there is sufficient regularisation (via priors in Bayesian context), the selection-induced bias is negligible even if the risk of selecting the “wrong” model is not small. We do not look to eliminate bias, but identify when it is non-negligible. In particular, we agree with past authors that selection-induced bias is not always a problem, and rather than a zero-one assessment of its existence, we endeavour to talk about negligible and non-negligible bias.

Very little work has dealt with bias correction in model selection, and those procedures that have been previously presented require intense computation. For instance, Stone (1974) proposed a procedure of nested CV (which they called the “double cross”) which is insensitive to sampling and partitioning schemes. This procedure operates with complexity  $O(K^2)$  in terms of the number of folds  $K$ : often

fiendishly expensive.<sup>1</sup> An alternative was provided by [Tibshirani and Tibshirani \(2009\)](#) who made use of the bootstrap ([Efron and Tibshirani, 1993](#)) to correct CV estimates. While less expensive, the complexity nonetheless scales in  $O(KB)$  with  $B$  denoting the number of bootstrap samples. Our contribution is a lightweight alternative to these procedures, requiring only the sufficient statistics from the standard CV process and *no* additional computation. In other words, it has linear complexity  $O(K)$ . While our proposal is simple, and thus cannot be expected to produce perfect results, it can serve as a useful demonstration of the dynamics of selection-induced bias: when one model is sufficiently better than other candidates then we can proceed more comfortably, and when models are close we must take care and perform a correction.

We provide the statistician with a diagnostic for our procedure, and should one estimate it to not be safe for their task, then one can fall-back on safer albeit more computationally expensive approaches. While we don't claim that CV estimates are the best way to do the model selection, it can be sometimes computationally convenient and sufficiently good; our approach helps to understand when it is sufficiently good.

## 1.2 Alternatives to model selection

While predictive performance-based model selection can operate as a convenient non-parametric method of determining whether there is a significant difference between models, simply investigating parameter posteriors can sometimes provide sharper resolution to the problem ([Gelman et al., 2014, 1996](#); [Wang and Gelman, 2015](#)). When the model parameters are interpretable, causally consistent, nested, and there is no collinearity between predictors, then we can often investigate their posteriors directly with more success. Since the predictive distribution integrates over the uncertainty over all parameters, it will often induce more variance in the estimator than is required for our use. For instance, we could investigate only the posterior of the extra parameter in the larger model and decide whether or not to include this parameter in our analysis based on whether its posterior differs greatly from the prior.

When we are not in the nested case, we might follow the suggestions of [Yates et al. \(2021\)](#) and either: (1) perform Bayesian model averaging (BMA; [Hoeting et al., 1999](#)) or model stacking ([Yao et al., 2023, 2021, 2018](#)); (2) use the largest model thoroughly-considered priors; or, (3) consider alternative model structures ([Garthwaite and Mubwandarikwa, 2010](#)).

BMA is an attractive initial first consideration, and has been shown by [Le and Clarke \(2022\)](#) to even be better in terms of prediction than model selection under some conditions. The good performance of BMA in the predictive sense has been covered previously by [Hoeting et al. \(1999\)](#) and [Raftery and Zheng \(2003\)](#). Bayesian model averaging weights, however, can be expensive to compute, and [Yao et al. \(2018\)](#) introduce pseudo-BMA weights as a lightweight alternative (we discuss the connection between pseudo-BMA weights and model selection based on elpd in [Appendix A](#)).

---

<sup>1</sup>For the rest of the paper, we will use  $K$  to denote the number of models, but here it is used to denote the number of CV folds. We will focus on the case of LOO-CV where we have the same number of folds as data observations:  $K = n$ .



The authors discuss that when none of the models considered is the true model, this method can be flawed, and motivate stacking of the models’ posterior predictive distributions as more robust alternative. This consists of identifying model weights such that the combination of densities minimises the LOO-CV mean-squared error. More recently, Yao et al. (2023) have extended this paradigm to include more exotic model combinations.

### 1.3 Structure of this paper

We begin by discussing how model selection decisions themselves are liable to induce over-confidence in the expected performance of a selected model in Section 2. We demonstrate how the magnitude of selection-induced bias grows with the number of candidate models in Section 3, wherein we present our first contribution: a novel order statistic-based tool to determine when a model selection decision is liable to induce non-negligible bias. In Section 4 we demonstrate how, while a single decision might introduce negligible bias, when we make multiple sequential decisions, and with specific reference to forward search, the bias is likely to compound. We then extend our first contribution to correct for selection-induced bias in forward search, and use it as a stopping rule. We manifest the reliability, computational tractability, and flexibility of our correction on both simulated and real-world data exercises, wherein we compare our stopping rule with traditional alternatives.<sup>2</sup>

## 2 Predictive model assessment, selection, and comparison

The focus of our paper is on developing a framework for modelling the predictive difference estimates between models, and specifically identifying when models are sufficiently different to safely make a model selection decision based on sufficient statistics.

### 2.1 Predictive model assessment

Vehtari et al. (2017) define the expected log point-wise predictive density (elpd) as a measure of predictive performance for a new dataset of  $n$  observations:

$$\text{elpd}(M_k | y) = \sum_{i=1}^n \int p_t(\tilde{y}_i) \log p_k(\tilde{y}_i | y) d\tilde{y}_i, \quad (1)$$

where  $p_t(\tilde{y}_i)$  denotes the density of the true data-generating process for  $\tilde{y}_i$ , and  $p_k(\tilde{y}_i | y) = \int p_k(\tilde{y}_i | \theta) p_k(\theta | y) d\theta$  is the posterior predictive distribution of model

<sup>2</sup>Code to replicate the results in this paper is freely available at <https://github.com/yanmclatchie/model-selection-uncertainty>.

$M_k$ . In practice, we do not have access to  $p_t(\cdot)$  and approximate the elpd with its leave-one-out cross-validated (LOO-CV) estimate,

$$\text{elpd}_{\text{LOO}}(M_k | y) := \sum_{i=1}^n \log p_k(y_i | y_{-i}) = \sum_{i=1}^n \log \int p_k(y_i | \theta) p_k(\theta | y_{-i}) d\theta. \quad (2)$$

If the data  $y$  are conditionally independent then [Gelfand et al. \(1992\)](#), and later [Gelfand \(1996\)](#), demonstrate how this quantity can be efficiently estimated with importance sampling by re-weighting posterior samples  $\theta^s$ , for  $s = 1, \dots, S$  indexing samples from the posterior, with weights

$$r_i^s = \frac{1}{p_k(y_i | \theta^s)} \propto \frac{p_k(\theta^s | y_{-i})}{p_k(\theta^s | y)} \quad (3)$$

leading to the importance sampling estimate

$$\widehat{\text{elpd}}_{\text{LOO}}(M_k | y) = \sum_{i=1}^n \log \left( \frac{\sum_{s=1}^S r_i^s p_k(y_i | \theta^s)}{\sum_{s=1}^S r_i^s} \right). \quad (4)$$

[Vehtari et al. \(2017\)](#) further demonstrate how this estimator can be stabilised and self-diagnosed with Pareto-smoothed importance sampling ([Vehtari et al., 2022](#)). Here we focused on LOO-CV, but our approach can also be used in the context of other cross-validation approaches, for example leave-future-out ([Bürkner et al., 2020](#); [Cooper et al., 2023](#)) and leave-one-group-out ([Merkle et al., 2019](#)).

## 2.2 Selection-induced bias

We can then use these predictive estimates to select one model from a pool of candidates. Even if each predictive estimate is almost unbiased for each model, however, when we use these for model selection the estimate for the selected model may incur some selection-induced bias.

Selection-induced bias emerges from selecting a model from a collection of models based on a noisy CV predictive estimate which could have been shown by the true elpd to be sub-optimal with respect to the models considered.<sup>3</sup> For the rest of this paper, we suppose that we choose the model with the highest LOO-CV elpd point estimate from a collection of  $K$  candidate models,

$$j^* = \arg \max_{k=1, \dots, K} \widehat{\text{elpd}}_{\text{LOO}}(M_k | y), \quad (5)$$

and define selection-induced bias concretely as

$$\text{bias}(M_1, \dots, M_K | y) = \widehat{\text{elpd}}_{\text{LOO}}(M_{j^*} | y) - \text{elpd}(M_{j^*} | y). \quad (6)$$

We then wish to understand under which situations choosing the model with the best LOO-CV elpd estimate may lead to non-negligible over-optimism in expected

---

<sup>3</sup>Most of the noisiness of these estimates is borne from finite data, and not the importance sampling scheme employed.

LOO-CV elpd (selection-induced bias), and how bad this bias is as a function of model similarity and the number of candidate models.

Rather than measure the performance of any model selection decision in terms of whether or not we choose the correct model (since the notion of a true model is difficult to justify in practice), we instead discuss how far we are from selecting the oracle model (in terms of minimal excess loss in expectation, as defined by [Arlot and Celisse, 2010](#)). In a word, we interest ourselves in the expected selection-induced bias injected by a given model selection decision, aim to minimise this as much as possible, and correct it when it is non-negligible. When we expect a model selection decision to induce only negligible amounts of selection-induced bias, we consider it a “safe” decision.

### 2.3 When is one model clearly better than another?

We begin by illustrating the  $K = 2$  regime, in which we have only two candidate models. A typical question arises: is one of these models better than the other? These two models may be achieved with different priors, be nested (as in one iteration of forward search, say), or may be completely non-nested. Not only are we interested in discerning when it is safe to select one of the two models over the other, we would also like to be able to say something about the potential danger associated with simply selecting the better model in terms of LOO-CV elpd. In this section, we firstly present the probability of selecting the wrong model under some oracle knowledge, before discussing the selection-induced bias arising from this decision when the models are indistinguishable. We interest ourselves later on in this paper primarily in discerning when models are indistinguishable in terms of their predictive performance.

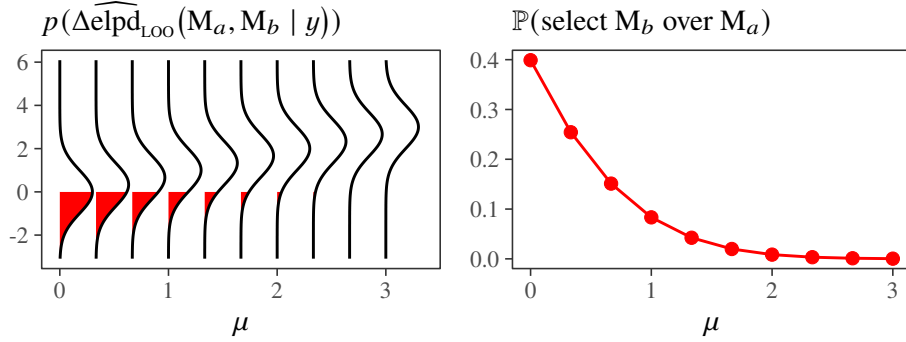
Consider two models,  $M_a$  and  $M_b$ , and suppose we know that  $M_a$  is the oracle model. We follow the definition of [Arlot and Celisse \(Section 2.2; 2010\)](#) in that the oracle model is the model having the minimum excess loss in expectation over multiple datasets from a collection of models. In particular, it is possible for the oracle model to differ from the model containing the truly relevant predictor set (for simplicity we call that the “true model”) in cases where that model can over-fit. In general, we don’t always need to choose the oracle model; the excess loss can just be negligible. It is often more useful to discuss this difference in terms of a continuous scale, rather than focus on the (often unrepresentative) case of either selecting the oracle model or not.

The structure of these two models does not influence model selection since we only consider their elpds. Denote the difference between the models’ theoretical elpds as

$$\Delta \text{elpd}(M_a, M_b \mid y) = \text{elpd}(M_a \mid y) - \text{elpd}(M_b \mid y). \quad (7)$$

As previously mentioned, we are most interested in the situation where no model is truly better than any other, that is that  $\Delta \text{elpd}(M_a, M_b \mid y) = 0$ . When this is the case, the estimated  $\widehat{\Delta \text{elpd}}_{\text{LOO}}(M_a, M_b \mid y)$  represents pure selection-induced bias, regardless of the model chosen.

Consider the following (purely illustrative) setting: assume the distribution of elpd *point estimates*,  $\widehat{\Delta \text{elpd}}_{\text{LOO}}(M_a, M_b \mid y)$ , is a Gaussian centred on the mean  $\mu$



**Figure 1:** The left plot shows illustrative LOO-CV elpd difference distributions between the true model  $M_a$  and alternative model  $M_b$ , as described in Equation 8. If we observe a LOO-CV estimated elpd difference the region coloured in red, then we are at risk of selecting the wrong model. The right plot shows this probability across different hypothetical mean values as computed in Equation 9.

and having variance  $\sigma^2 = 1$  (similar to the setup employed by [Sivula et al., 2022](#)). From this distribution, we “sample” one observation of the LOO-CV elpd difference, denoted<sup>4</sup>

$$\widehat{\Delta\text{elpd}}_{\text{LOO}}(M_a, M_b | y) := \widehat{\Delta\text{elpd}}_{a,b} \sim p(\widehat{\Delta\text{elpd}}_{a,b}) \equiv \text{normal}(\mu, 1^2). \quad (8)$$

When this observed elpd difference  $\widehat{\Delta\text{elpd}}_{a,b}$  is negative and we subsequently select  $M_b$  over  $M_a$  based on the LOO-CV estimate then we are at risk of inducing non-negligible selection-induced bias. Naturally, since  $M_a$  is the oracle model, the true elpd difference is non-negative.

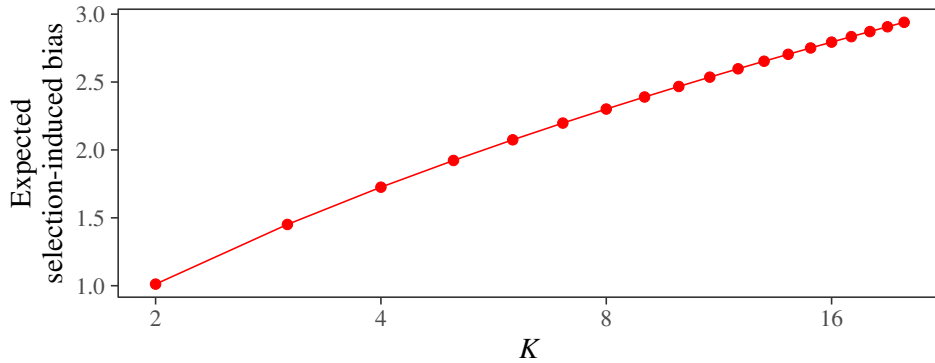
In this toy example, we can compute the theoretical probability of selecting the sub-optimal model, with the knowledge that  $M_a$  is the oracle model, as

$$\mathbb{P}(\text{select } M_b \text{ over } M_a) = \int_{-\infty}^0 \widehat{\Delta\text{elpd}}_{a,b} \cdot p(\widehat{\Delta\text{elpd}}_{a,b}) d(\widehat{\Delta\text{elpd}}_{a,b}), \quad (9)$$

where  $p(\widehat{\Delta\text{elpd}}_{a,b})$  is the oracle distribution of the elpd difference between the two models (omitting the dependence on the data  $y$  for brevity). We visualise this theoretical illustration in Figure 1, where we let the mean slowly grow away from zero, keeping the standard deviation the same. The probability of selecting the wrong model then is greatest for a given  $\sigma$  when the two models are most similar ( $\mu$  closest to zero), and decays more slowly with an increasing mean of the elpd difference distribution.

The selection-induced bias is likely to be non-negligible when it is much larger than the variation in the LOO-CV point estimates. In terms of Figure 1, if the selection-induced bias is less than the illustrative  $\sigma^2 = 1$ , then it is less critical to

<sup>4</sup>In practice, the data induce a LOO-CV point estimate. For the purposes of this demonstration, we assume that we can sample from the distribution of LOO-CV point estimates directly. Throughout this paper it is the distribution of these point estimates that we interest ourselves in primarily.



**Figure 2:** Theoretical illustration. The hypothetical expected selection-induced bias in the presence of  $K$  candidate models. While the bias is negligible for  $K = 2$ , it can grow rapidly with the number of candidate models. We discuss this trend later.

model selection than if it were much larger. In a word, selection-induced bias is said to be negligible if it is dominated by the variation in the LOO-CV point estimates.

Nonetheless, in our experience (outwith this paper) we are able in most cases to simply choose the more complex model fitted with reasonable priors since the bias is negligible in the  $K = 2$  case. However, as the number of effectively equivalent models grows (e.g., if the number of candidate models grows exponentially with the number of observations  $n$ , or if the number of predictors  $p$  is much larger than  $\log(n)$  as described by [Gelman et al., 2013](#)) the potential bias grows too. While selection-induced bias remains negligible in the two model case, it can balloon as  $K$  grows. The specifics of the rate shown in Figure 2 is described later in this paper. This many model regime is the more interesting case.

### 3 Bias grows with the number of candidate models

In the regime of many candidate models, when none is truly more predictive than any other, the probability of one model *appearing* better than the rest grows with the number of models being compared. This is true even if independent test data would show that there is in fact no difference from it to the true model (or that it is in fact worse). As this probability grows, so too does the magnitude of selection-induced bias. In this section, we first show how these high risk regions behave as the number of models grows, and propose a selection-induced bias estimation scheme based on order statistics to identify when the bias is no longer negligible.

#### 3.1 Empirical demonstration

We fit a series of nested linear models to  $n = 100$  data points sampled from

$$y = X\beta + \epsilon \tag{10}$$

$$\epsilon \sim \text{normal}(0, \sigma^2 I), \tag{11}$$

where  $I$  is the  $n \times n$  identity matrix, and  $0 \in \mathbb{R}^n$  is the  $n$ -dimensional zero vector. The data  $X_i = (1, X_{i,1}, \dots, X_{i,K-1})$  are  $K$  dimensional, with  $X_{i,1}, \dots, X_{i,K-1} \stackrel{\text{iid}}{\sim} \text{normal}(0, 1)$ , and the true regression coefficients take the form  $\beta = (1, \beta_\Delta, 0, \dots, 0) \in \mathbb{R}^K$ . The variance of the residuals,  $\sigma$ , is chosen to vary to the signal-to-noise with  $\beta_\Delta$  while keeping the marginal variance constant. In particular, we only consider  $\beta_\Delta < 1$ , and set  $\sigma^2 = 1 - \beta_\Delta^2$ . We will then fit a baseline model comprising only the intercept term and denoted  $M_{\text{base}} : y_i | \beta_1, \tau \sim \text{normal}(\beta_1, \tau^2)$ . We compare this to the remaining  $K - 1$  one-predictor models of the form

$$M_k : y_i | \beta_1, \beta_k, \tau \sim \text{normal}(\beta_1 + X_{i,k}\beta_k, \tau^2), \quad (12)$$

for  $k = 2, \dots, K$  so that model  $M_2$  alone contains the two predictor set, and the scale factor  $\tau$  is treated as an unknown parameter across all models. Naturally, letting  $\beta_\Delta = 0$  collapses all models to be equivalent.<sup>5</sup>

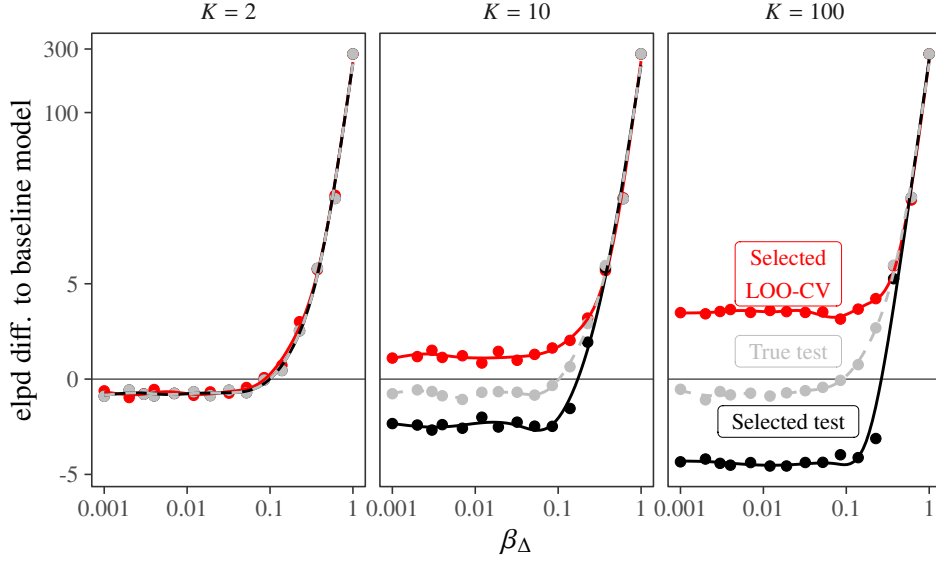
We vary the true  $\beta_\Delta$  an  $\epsilon$  distance from zero. For each value of  $\beta_\Delta$ , we investigate the LOO-CV elpd difference point estimates of the  $K - 1$  models from the baseline model, and identify the model with the highest (best) cross-validated score. This is then compared to its test elpd, and to that of the model with the known true predictor set (which we call the “true model”, even if we do not use the true values of the predictors). We show this for different numbers of candidate models  $K$  and along  $\beta_\Delta$  in Figure 3. As previously mentioned, the selection-induced bias (difference between the LOO-CV elpd and the test elpd of the selected model) is negligible when  $K = 2$ , but grows with  $K$ . For  $K = 100$ , it would be unsafe to make a decision based on LOO-CV elpd point estimates up until the theoretical  $\beta_\Delta$  grows above  $\approx 0.3$ , as the selection-induced bias is non-negligible. There is no formal threshold for non-negligible bias, but in Appendix A we show how the difference between elpd estimates can be connected to the notion of non-zero weights in model averaging. The magnitude of the bias also grows with  $K$ , seen in the gap between the LOO-CV elpd and test elpd of the selected model.

When the true value of  $\beta_\Delta$  is low, the model containing the known true predictors does not out-perform the baseline model in terms of test elpd. It is in these situations where the difference between the oracle model and so-called true model becomes apparent. In this case, the baseline model has the lowest excess loss in expectation over data realisations. When the true model performs worse than the oracle model, but we proceed to select the former, we have over-fitting. This excess loss from the oracle model is the difference from the baseline model to the selected model’s test elpd. As  $\beta_\Delta$  grows sufficiently large to consistently select the true model, this tends to zero, as too does selection-induced bias. The region between the LOO-CV and test elpd of the selected model is the selection-induced bias, while the region between the oracle elpd and the test elpd of the selected model is over-fitting. In this example, the magnitude of over-fitting is negligible across all values of  $\beta_\Delta$ .

Knowing then that the risk of selection-induced bias exists, and not knowing the true value of  $\beta_\Delta$ , we presently concern ourselves with producing a selection-induced

---

<sup>5</sup>We fit the model using independent  $\text{normal}(0, 10^2)$  priors over the regression coefficients, and a half-normal prior over  $\tau^2$  with mean zero and unit variance.



**Figure 3:** Simulated nested linear regression example. The red lines indicate the best LOO-CV elpd difference of  $K$  candidates from a baseline model (the horizontal zero line), and the black lines are the test elpd difference. In grey is the test elpd of the model fit with the true predictors. The region between the red and black curves is selection-induced bias, whereas the region between the oracle model (this is the zero line up to  $\beta_\Delta \approx 0.3$ , after which it is the grey line) and black curves is what we refer to as over-fitting. The magnitude of the selection-induced bias grows with  $K$ .

bias estimation procedure geared towards identifying high risk regimes in the presence of many candidate models by modelling the distribution of LOO-CV elpd difference estimates.

### 3.2 Our proposal: modelling the elpd differences

Suppose now that we have  $K$  candidate models, none of which is better than any other. The best LOO-CV elpd difference point estimate from a baseline model may appear, purely by chance, to be greater than zero. By modelling the underlying distribution of the elpd difference point estimates as iid draws from a Gaussian distribution centred on zero (so that none is truly better than the baseline model) with some variance  $\sigma_K^2$ ,  $\widehat{\Delta \text{elpd}}_{\text{LOO}}^{(i)} \sim \text{normal}(0, \sigma_K)$ ,  $i = 1, \dots, K$ , we can leverage order statistics to produce a selection-induced bias estimate.

Blom (1960) posited an approximation to the expected maximum from  $K$  iid samples from a standard normal distribution, denoted  $S^{(K)}$  (later developed by Harter, 1961; Royston, 1982):

$$S^{(K)} = \Phi^{-1} \left( \frac{K - \alpha}{K - 2\alpha + 1} \right). \quad (13)$$

The formulation allows for an approximation parameter,  $\alpha$ . In the limit as  $n \rightarrow \infty$ , the expectation of the maximum order statistic lies within the interval between levels

$\alpha = 0.39$  and  $\alpha = 0.5$ . As such, for maximum safety we will use  $\alpha = 0.5$  herein.

In the presence of  $K$  candidate models, they are practically equivalent if the maximum LOO-CV elpd difference point estimate to the baseline model is less than

$$S^{(K)} \hat{\sigma}_K. \quad (14)$$

We estimate  $\hat{\sigma}_K$  as the maximum likelihood estimator of the standard deviation of a half-normal distribution fit to the upper half-tail of elpd difference point estimates (greater than their median,  $\hat{m}_K$ ),

$$\hat{\sigma}_K = \sqrt{\frac{2}{K} \sum_{k=1}^K \mathbf{1}_{\widehat{\text{elpd}}_{\text{LOO}}^{(k)} \geq \hat{m}_K} \left( \widehat{\Delta \text{elpd}}_{\text{LOO}}^{(k)} \right) \cdot \left( \widehat{\Delta \text{elpd}}_{\text{LOO}}^{(k)} - \hat{m}_K \right)^2}. \quad (15)$$

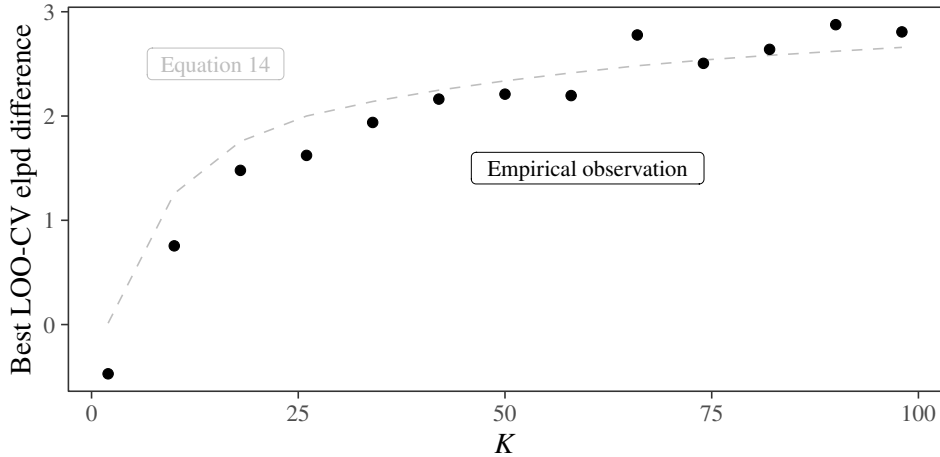
We denote by  $\mathbf{1}_A(x)$  the indicator function over a set  $A \subset X$  at  $x \in X$ . This is not yet an estimate of the selection-induced bias, but the highest LOO-CV elpd difference expected from a collection of  $K$  models, none of them truly more predictive than any other. In particular, this is an estimate of the expected distance between the zero line and the red line in Figure 3. We discuss later in Section 4.1 how one might use this as the basis for such an estimate of the selection-induced bias (the difference between red and black line in Figure 3).

We fit a half-normal distribution to the upper half of the observed elpd difference point estimates instead of a fitting a full normal distribution over all of them because (even if none of them is truly relevant) the distribution of elpd difference point estimates is likely to be slightly skewed. A similar result has previously been noted by [Sivula et al. \(2022\)](#), who showed the *full* elpd difference distribution (not just the point estimates) to be more skewed the more similar the models being compared are. We interest ourselves only in the right tail of the distribution of point estimates as it is the more relevant to the maximum order statistic. We discuss the validity of this stabilisation technique and how to diagnose it in Section 3.3.

In Figure 4, we fix  $\beta_\Delta = 0$  and vary  $K$  in the experiment design previously described by Equations 10–11. For each value of  $K$ , we investigate the best observed LOO-CV elpd difference to the baseline model. Since we have fixed  $\beta_\Delta = 0$ , we know that no model is theoretically better than any other. As such, these empirical LOO-CV elpd differences are just random noise. We repeat the experiment 100 times for each value of  $K$  and show the mean of the empirically observed best LOO-CV elpd difference in each regime compared to the expected best observed LOO-CV elpd difference as computed according to Equation 14. This theoretical boundary grows quickly for  $K$  less than 10, before slowing. Our order statistics-based equation closely matches empirical observations in this example.

A second empirical demonstration of Equation 14 is provided in Appendix B with reference to a real data study on the effect of a poverty reduction intervention on infant brain activity.





**Figure 4:** Simulated nested linear regression example. The mean best LOO-CV elpd difference of  $K$  candidate models from the baseline model (black dots) over 100 data realisations when  $\beta_{\Delta} = 0$ , and the expected best observed LOO-CV elpd difference of  $K$  models as computed according to Equation 14 (dashed grey line).

### 3.3 Diagnosing the theoretic assumptions

We have made a number of assumptions thus far in our analysis. In particular, we have assumed that the elpd difference distributions of models of equivalent predictive performance are:

1. normally distributed around zero (or at least half-normal above zero); and,
2. independent and identically distributed.

It is known that the empirical elpd difference distributions are slightly skewed, with heavier right tails (this was seen in empirical results not shown here, and has previously been discussed by [Sivula et al., 2022](#)). This behaviour becomes more pronounced as models become closer together (through the use of regularising priors for example), and such correlation may appear as models make more similar predictions. It is possible then that the distribution of elpd difference point estimates (which we interest ourselves in) may also be slightly skewed. As such, the normal approximation may not be representative of the underlying distribution, and the half-normal approximation is better able to capture the characteristics of the right tail. Indeed, it is this right tail that is of most interest to us, as it contributes most to the maximum order statistics.

It would be possible to fit a sup-Gaussian distribution to this tail with more parameters, e.g. a half-Student- $t$ , but as shown empirically, the half-normal works sufficiently well. Since it has fewer parameters we expect it to be more robust, and from the half-normal we get information from the bulk of the distribution. For example we could also have fit a Pareto distribution to the tail as done by [Vehtari et al. \(2022\)](#). We would, however, then have to get information from the extreme tail. If the distributions are truly skewed, we expect our normal approximation to inject some bias, but if we use a more flexible approximation there would be higher variance.

We are also able to diagnose when our half-normal approximation is likely to fail:

when the right tail of the empirical distribution has infinite variance. We follow [Vehtari et al. \(2022\)](#) and first identify a tail cut-off point of the empirical elpd difference point estimate distribution in line with [Scarrott and MacDonald \(2012\)](#), and subsequently fit a generalised Pareto distribution to this tail. [Zhang and Stephens \(2009\)](#) present an efficient Bayesian procedure for estimating the parameters of the distribution. In particular, the  $\hat{k}$  parameter of the generalised Pareto distribution is related to its number of fractional moments. We follow the rule of thumb presented by [Vehtari et al. \(2022, Section 3.2.4., Equation 13\)](#) in establishing the minimum required value of  $\hat{k}$  to be confident that the right tail has finite variance when comparing 10 or more models.<sup>6</sup>

When this Pareto  $\hat{k}$  value is too large, it suggests that the elpd difference point estimates are likely from a distribution with infinite variance. It is at this stage that one should fall back on safer, albeit more computationally expensive approaches: the double-cross ([Stone, 1974](#)) or the bootstrap ([Tibshirani and Tibshirani, 2009](#)).

Elsewhere, we know that the distribution of the maximum order statistic (of which we interest ourselves only in the mean) is slightly skewed. Because of this skewed-ness, it is possible that the mean alone (which we have denoted  $S^{(K)}$ ) under-estimates the maximum of LOO-CV elpd difference estimates.

When the model space grows too large to consider each possible model at once, and if the models can be nested, then one might favour some efficient search through the model space. Such searches come with their own risks. We presently move from the realm of single decisions towards sequential model selection.

## 4 Compounding the bias through forward search

While selection-induced bias might remain negligible when we need only select one model from a group of models once, when we compound decisions, the magnitude of the bias is likely to compound too. For instance, in forward search, we begin with a baseline model (usually the intercept-only model) and then consider all models containing one predictor and the intercept term. Having fit all so-called “size-one” candidate models, we choose the one with the highest LOO-CV elpd point estimate. We now perform another forward step, and fit all models containing the intercept, the predictor chosen previously, and a second additional predictor. Once more we select the size-two model based on the highest LOO-CV elpd point estimate. We repeat this either until either some stopping criterion is met, or we have exhausted all available predictors.

We can easily implement various other search heuristics, such as exhaustive search ([Galatenko et al., 2015](#)), backward search ([Nilsson, 1998](#)), and stochastic search ([George and McCulloch, 1993](#); [Ntzoufras et al., 2000](#)). However, we proceed with forward search alone as it affords the possibility of an efficient stopping rule ([Harrell, 2001](#)) and intuitive compounded bias estimation (which we present in [Section 4.1](#)).

At the beginning of this search, it is likely that there exist at least some predictors that improve on the predictive performance of the baseline model, but there comes a

---

<sup>6</sup>We require the number of models to be at least 10 since at least 5 tail samples are needed to reliably estimate the  $\hat{k}$  parameter.

point after which the addition of further predictors affords the statistics no theoretical predictive performance. We will refer to this as the “point of predictive saturation” herein. We wish to identify this point, and estimate the bias induced by continuing the search beyond it. Under appropriate prior consideration, the full model is considered the oracle model. If one witnesses a sub-model having better test performance than the full model, then we recommend they reconsider their priors. In practice, we do not need to identify the oracle model size, only a sub-model having negligible excess loss compared to the oracle model.

When we observe a separation of the LOO-CV and test elpd along the search path, we can conclude that there is non-negligible selection induced bias in the performance estimates: purely due to the predictor ordering achieved by LOO-CV based search, we are led to believe that some sub-models can outperform larger models. This is seen by the elpd estimate at the bulge being greater than that of the final full model. In practice, [McLatchie et al. \(2023\)](#) argue that should one observe such a bulge in the forward search path, then one should cross-validate over search paths (something akin to the double-cross), but only up to the model size inducing this bulge. This is because cross-validating over forward search paths can be expensive, and we understand that the model size of maximal selection-induced bias is likely greater than the model size of theoretic predictive saturation.

The dynamics of forward search can be seen in the relationship between the predictive criterion (the elpd, say) and the model sizes along the so-called “search path”. The plot typically shows an increasing elpd curve, indicating that the addition of further predictors improves predictive performance. We expect the curve to flatten after a given model size, indicating that the addition of further predictors does *not* improve the criterion any longer: the point of predictive saturation. However, it is possible that in the finite data regime, and under model misspecification, we instead witness a bulge along this search path, such as that shown in [Figure 5](#). The test curve in [Figure 5](#) reaches the point of predictive equivalence and then dips down with increased model size. This is *not* necessarily selection-induced bias, but could also be over-fitting. Indeed, we aim to correct selection-induced bias to identify non-negligible over-fitting in forward search.

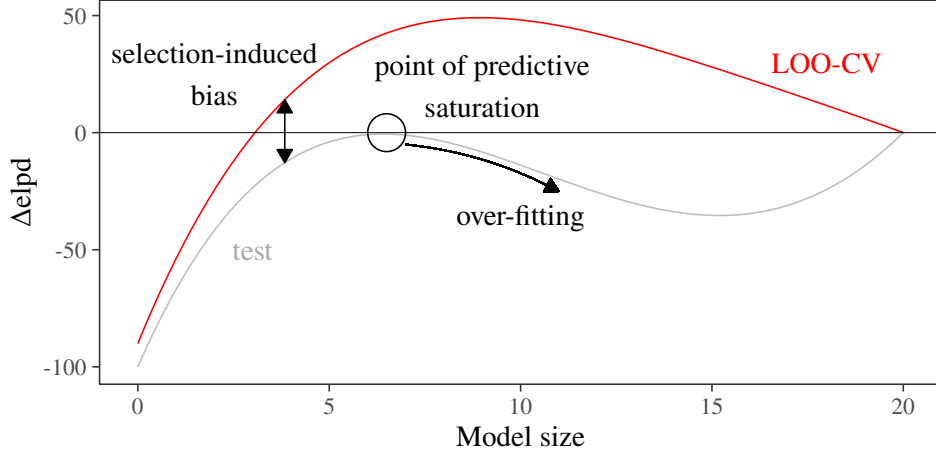
In order to correct selection-induced bias and estimate the location of the point of predictive saturation (and the bias induced by continuing beyond it) we introduce the following application of our proposal from [Section 3.2](#).

#### 4.1 Correcting for selection-induced bias in forward search

We can estimate the selection-induced bias at model size  $k$  in a forward search path, and correct the LOO-CV elpd difference estimate, denoted  $\widehat{\Delta\text{elpd}}_{\text{LOO}}^{(k)}$ , accordingly by

$$\widehat{\Delta\text{elpd}}_{\text{corrected}}^{(k)} = \begin{cases} \widehat{\Delta\text{elpd}}_{\text{LOO}}^{(k)} - \widehat{\text{bias}}^{(k)}, & \text{if } |\widehat{\Delta\text{elpd}}_{\text{LOO}}^{(k)}| < S^{(k)}\hat{\sigma}_k \\ \widehat{\Delta\text{elpd}}_{\text{LOO}}^{(k)}, & \text{otherwise,} \end{cases} \quad (16)$$

where  $S^{(K)}\hat{\sigma}_K$  is as defined in [Section 3.2](#). Intuitively, if we detect non-negligible selection-induced bias, then we correct the LOO-CV elpd difference by an estimate of



**Figure 5:** Illustrative sketch. In red we show the LOO-CV elpd difference to the full (reference) model along growing model size, and in grey the out-of-sample elpd difference. Selection-induced bias is the divergence of the red line from the grey. Over-fitting is the drop in predictive performance with increased model size following the point of predictive saturation.

the selection-induced bias induced by step  $k$ .

We produce this estimate of selection induced bias, denoted  $\widehat{\text{bias}}^{(k)}$ , building on Section 3.2 as  $\widehat{\text{bias}}^{(k)} = 1.5 \times S^{(k)} \hat{\sigma}_k$ . The nature of this estimate stems from Watanabe (2010, Table 3), who discussed how cross-validated scores and the average Bayes generalisation loss can exhibit strong negative correlation (albeit demonstrated with only one example). Theoretically this is due the removal of an (influential) observation from the training set skewing the posterior predictive distribution away from the true data-generating distribution, and its re-appearance in the test set seeming like an outlier. As such, the correction of LOO-CV elpd estimates warrants not only a correction to the median (the difference between red line and zero line in Figure 3), but a negatively-correlated correction (the difference between red and black lines in Figure 3). In a word, the corrected elpd difference value should be a reflection across the median of candidate elpd differences. The choice of magnitude of 1.5 is reflective of the magnitude of the negative correlation of approximately  $-0.8$  noted by Watanabe (2010). However, since this number was achieved only over one example, we understand that there can be instances where this correlation can be closer to 0 or  $-1$ . As such, we show repetitions of some later experiments using values of  $\{1, 1.5, 2\}$  when estimating the bias in Appendix C.

Cumulatively summing these corrected difference estimates,

$$\widehat{\text{elpd}}_{\text{corrected}}^{(k)} = \widehat{\text{elpd}}_{\text{LOO}}^{(0)} + \sum_{i=1}^k \Delta \widehat{\text{elpd}}_{\text{corrected}}^{(i)}, \quad (17)$$

starting at the base LOO-CV elpd estimate at model size 0, denoted  $\widehat{\text{elpd}}_{\text{LOO}}^{(0)}$ , allows us to estimate the bias induced by continuing beyond the point of predictive saturation.

As such, our proposition comprises two parts: an order statistics-based estimation of selection-induced bias (from Section 3.2), and an associated correction term.

## 4.2 A simulated example

We can visualise the performance of our correction through a simple simulated data experiment: we generate data according to

$$x \sim \text{normal}(0, R) \tag{18}$$

$$y \sim \text{normal}(w^T x, \sigma^2), \tag{19}$$

where the number of predictors is set to  $p = 100$ ,  $\text{normal}(\mu, \Sigma)$  denotes a  $p$ -dimensional normal distribution with mean vector  $\mu$  and covariance matrix  $\Sigma$ . The matrix  $R \in \mathbb{R}^{p \times p}$  is block diagonal, with each block having dimension  $5 \times 5$ . Each predictor has mean zero and unit variance and is correlated with the other four predictors in its block with coefficient  $\rho$ , and uncorrelated with the predictors in all other blocks. We will investigate regimes of zero and high correlation:  $\rho = \{0, 0.9\}$ . Further, the weights  $w$  are such that only the first 15 predictors influence the target  $y$  with weights  $(w_{1:5}, w_{6:10}, w_{11:15}) = (\xi, 0.5\xi, 0.25\xi)$ , and zero otherwise. We set  $\xi = 0.59$  and  $\sigma^2 = 1$  to fix  $R^2 = 0.7$ . We simulate  $n = \{100, 200, 400\}$  data points according to this data-generating process (DGP). Across all six  $n$  and  $\rho$  regimes, we perform forward search based on the maximum LOO-CV elpd point estimate of the candidate model, fitting each model using independent Gaussian priors over the regression coefficients (which we understand to be liable to over-fit), and comparing them to a reference model fit with the sparsity-inducing R2D2 prior (which is discussed in more detail in Section 4.3 and Appendix D).

In Figure 6(a) we visualise the performance of our bias correction across these six regimes in terms of the mean log predictive density (mlpd) computed on  $\tilde{n}$  independent test observations,  $\tilde{y}_{1:\tilde{n}}$ , as

$$\text{mlpd}(\mathbf{M}_k | \tilde{y}) = \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} \log p_k(\tilde{y}_i | \theta), \tag{20}$$

which admits a LOO-CV difference estimate similar to the elpd (Pironen and Vehtari, 2017a). We do not correct the LOO-CV differences beyond their LOO-CV bulge since at this point it is already clear that we have over-fit the data.

We observe from these plots two positive results and one negative:

1. the corrected mlpd does not greatly surpass the mlpd of the of the reference model, contrary to the LOO-CV mlpd;
2. the corrected mlpd estimates the point of predictive saturation close to the true point of predictive saturation on test data; however,
3. the corrected mlpd of the search path can under-estimate the true selection-induced bias slightly in cases of severe over-fitting (most notably when  $n = p$ ).

This suggests that a simple order statistic model is able to represent selection-induced bias well enough to be useful. Naturally, the quality of our correction improves with data size and the number of candidate models.

The approximation performs best in the larger data settings. It would further appear that, in this example, once  $n > 4p$  the degree of selection-induced bias is negligible. This could be due in part to the negative correlation between Bayes cross-validation and the average Bayes generalisation loss (Watanabe, 2010): even if we know the underlying distribution of LOO-CV elpd estimates, the asymptotic theoretical elpd distribution is likely negatively correlated with it (and worse in absolute value). Such correlation is naturally exaggerated in smaller data regimes, as the relative influence of data points grows, and our proposal under-estimates the true selection-induced bias. Despite under-estimating the bias slightly in the  $n = p$  case, it appears to correct enough of the selection-induced bias to accurately and efficiently identify over-fitting across all regimes.

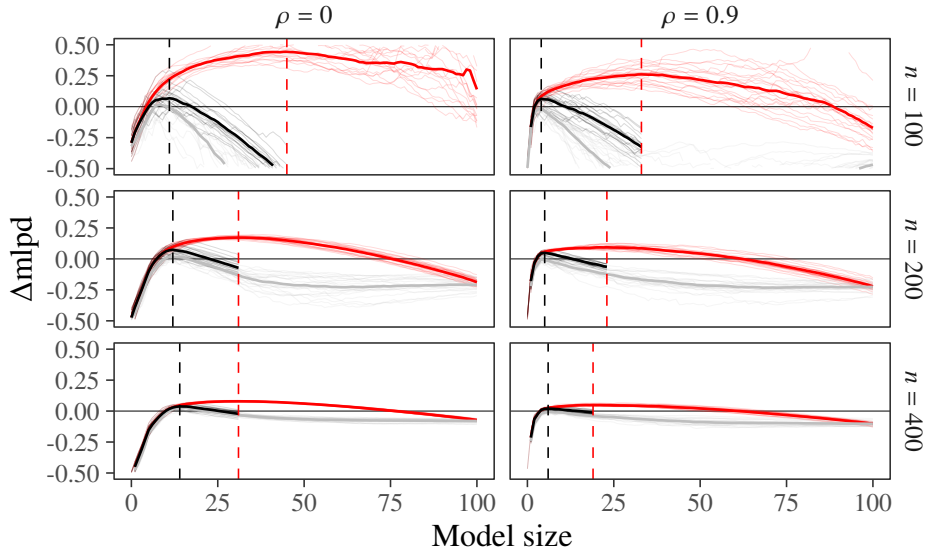
As previously mentioned, it is not clear how one would motivate any particular selection-induced bias correction beyond what we have proposed (see Section 6 for more discussion, and Appendix C for a repetition of the experiments using different levels of correction), and we rather recommend that one endeavours instead to mitigate against non-negligible over-fitting in the first instance. One such way is with a more principled application of the prior.

### 4.3 Reducing selection-induced bias with priors

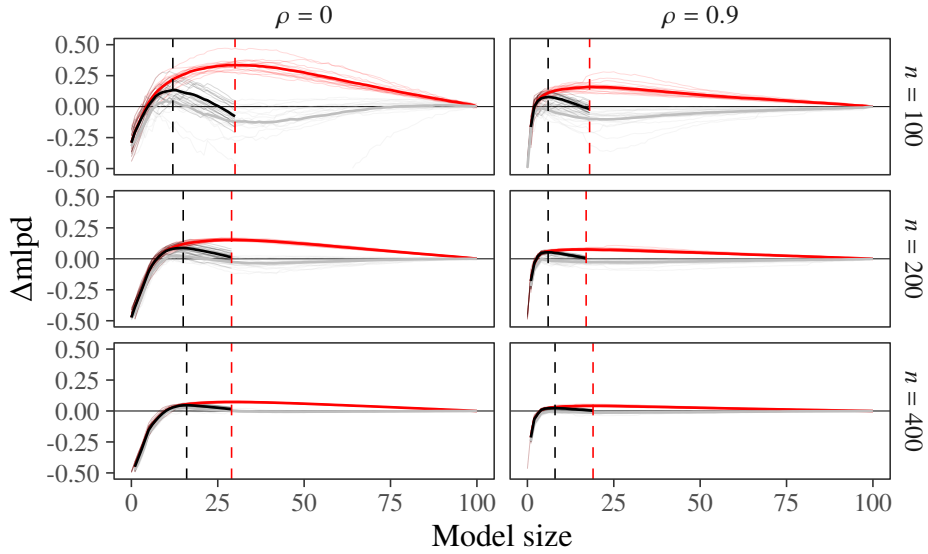
Gelman et al. (2017) reason that prior distributions in Bayesian inference should be considered generatively in terms of the potential observations consistent with the prior, and predictively in terms of how these observations align with the data we truly collect. Cawley and Talbot (2010) also discuss the benefits of using regularising techniques to avoid over-fitting in selection; generative priors are such a regulariser.

Joint shrinkage priors are a modern class of such generative priors which encode the statistician’s belief on a predictive metric and filter this information through the model structure to define prior definitions on the model parameters. For instance, the R2D2(M2) prior (Aguilar and Bürkner, 2023; Zhang et al., 2022) encodes a prior over the model’s  $R^2$  which is then mapped to a joint prior over the regression coefficients and the residual variance. By positing a lower expected predictive performance in terms of explained variance than what is implied by independent Gaussian priors, we shrink irrelevant regression coefficients more. This posterior shrinkage and reduced posterior uncertainty then naturally feeds through to the predictive distribution and thus the LOO-CV elpd estimate in Equation 2. This means that candidate models make more similar predictions, the empirical distribution of elpd difference point estimates is more concentrated, and the magnitude of selection-induced bias is reduced. See Appendix D for more discussion on the specifics of the priors and the hyper-parameters used throughout this paper.

Repeating our example from Section 4.2, but using a more regularising generative prior, we can reduce the degree of selection-induced bias along the search path. In Figure 6(b), we find that once more our heuristic identifies the correct model size after which no more predictors truly aid predictive performance, and recovers the shape of the test curve. We do, however, under-estimate the bias slightly in the  $n = p$  regimes. This is likely due to the models making more similar predictions, and thus our



(a) Using independent Gaussian priors over regression coefficients.



(b) Using the R2D2 prior over regression coefficients.

**Figure 6:** Forward search simulation experiment. The red lines show the LOO-CV mlpd difference to the reference model over 20 iterations, the grey lines show the test mlpd difference along those same paths, and the black show the corrected mlpd difference, produced using our correction from Section 4.1. The vertical dashed red and black lines indicate the point of maximum LOO-CV and corrected LOO-CV mlpd difference respectively.

proposed Equation 14 will likely under-estimate the magnitude of the bias. However, the identified points of inflection are similar under both priors, reassuring us that our selection-induced bias correction is stable across different bias regimes.

Priors are not the only way one might consider reducing selection-induced bias in forward search; projection predictive inference is one information-theoretic approach to the task that has been proven to be very effective.

#### 4.4 Projection predictive inference

We will later compare the model sizes selected following our bias correction to projection predictive inference. The projection approach is known to greatly reduce the variance in the relative performance estimates, leading to much smaller selection-induced bias (see, e.g. [Piironen and Vehtari, 2017a](#)).

Projection predictive inference was first introduced concretely by [Goutis \(1998\)](#) and [Dupuis and Robert \(2003\)](#), but follows the decision-theoretical idea presented previously by [Lindley \(1968\)](#). In this paradigm, one begins with a large reference model one believes to be well-suited to prediction and subsequently constructs small models capable of approximately replicating its predictive performance; essentially, fitting sub-models to the fit of the reference model.

We call the model we project the reference model onto the “restricted model” or “sub-model”, and denote its parameter space by  $\Theta^\perp$  (which may be completely unrelated to the parameter space of the reference model,  $\Theta^*$ ). Further denote the predictor-conditional predictive distribution of the restricted model on  $\tilde{n}$  unobserved data observations  $\tilde{y} = (\tilde{y}_1, \dots, \tilde{y}_{\tilde{n}})^T$  by  $q(\tilde{y}_i | \theta^\perp)$ , for  $\theta^\perp \in \Theta^\perp$  and  $i \in \{1, \dots, \tilde{n}\}$ , and that of the reference model by  $p(\tilde{y}_i | \theta^*)$ . Then we compute the projection draw-wise on  $s \in \{1, \dots, S\}$  posterior draws from the reference model as

$$\theta_s^\perp = \arg \min_{\theta^\perp \in \Theta^\perp} \frac{1}{n} \sum_{i=1}^n \mathbb{KL} \{ p(\tilde{y}_i | \theta_s^*) \parallel q(\tilde{y}_i | \theta^\perp) \}. \quad (21)$$

For implementation details see the papers by [Catalina et al. \(2021, 2020\)](#), [Weber and Vehtari \(2023\)](#), and [Piironen et al. \(2020\)](#). [McLatchie et al. \(2023\)](#) outline further workflow steps to minimise the computational burden while achieving robust results.

Projective forward search begins by projecting the reference model onto the intercept-only model to establish the root of the search tree. Subsequently, the reference model is projected onto size-one models, from which we select the model whose posterior predictive distribution is most similar to the reference model’s posterior predictive distribution by the KL divergence in line with the work of [Leamer \(1979\)](#). The search then continues as before. Because of the use of the KL divergence over LOO-CV, our order statistics-based bias estimation and correction can not be immediately translated to work with projection predictive inference. It is the disjoint between the KL divergence and the predictive performance noted by [Robert \(2014\)](#), however, which is partially responsible for the reduced selection-induced bias in projection predictive inference.

As such, if the reference model (and sub-models) can be constructed in such a way that conforms with the requirements of projection predictive inference, then we recommend one begins with this. For observational families whose KL divergence can not be easily computed, or for model structures that can not be easily decomposed into



additive components, cross-validation forward search (as described in the beginning of Section 4) can be a useful alternative.

## 4.5 Real world data

To manifest the performance of our selection-induced bias correction beyond simulated settings, we compare it with classical stopping criteria, including projection predictive inference, across four publicly available real world datasets: one regression task (Redmond and Baveja, 2002) and three binary classification (Gorman and Sejnowski, 1988; Krzysztof Cios, 2001; Sigillito et al., 1989) summarised in Table D1 (in Appendix D) after any necessary pre-processing.<sup>7</sup> For the regression task we fit a linear regression reference model with a Gaussian observation family, and for the binary classification a Bernoulli generalised linear model as implemented in `brms` (Bürkner, 2017) using the probit link function. For each dataset, we perform forward search based on LOO-CV elpd, fitting sub-models with independent standard Gaussian priors over the regression coefficients in the first instance, and then again with sparse priors (the R2D2 for regression tasks, and the regularised horseshoe for the classification tasks; see Appendix D for details on the hyper-parameters used for each task). We split all data into 10 cross-validation folds, leaving 10% of the data for testing in each fold. In Figure 7 we show the cross-validated and test search paths for each of these, along with the corrected search paths as computed using Section 4.1.

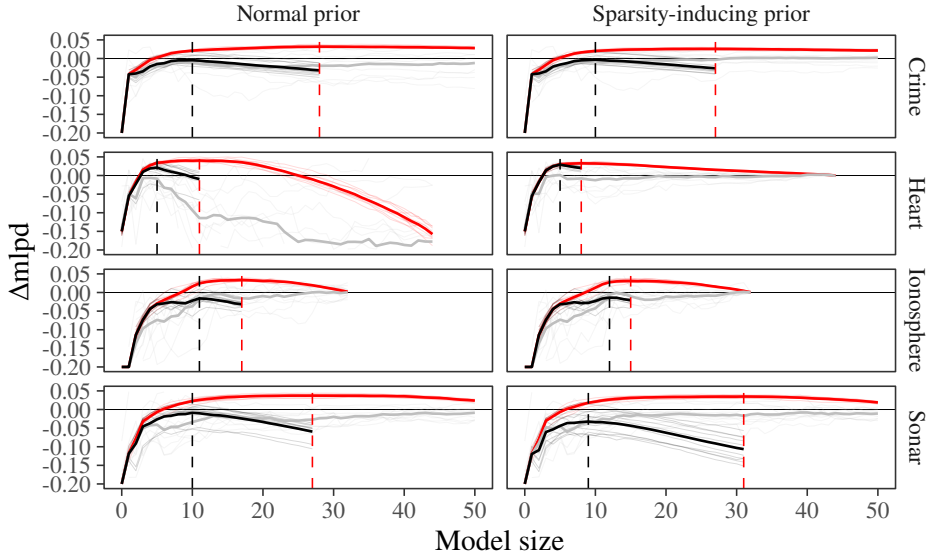
Across all real data experiments, our corrected LOO-CV mlpd did not significantly exceed that of the reference model, except for the heart data. This indicates our correction effectively addresses selection-induced bias without falsely suggesting better performance than the reference model. In some cases, our correction appeared to under-correct for small model sizes but eventually recovered the true test mlpd. In the sonar data we witnessed the opposite: our proposal over-corrected bias. Even when our correction over- or under-estimates bias, it provides valuable insights into test mlpd behaviour.

We presently extend our selection-induced bias correction to produce a stopping criterion: we select the model size inducing the maximum *corrected* mlpd along the search path as computed by Section 4.1. We compare this to simply stopping at the point of maximum LOO-CV (which we call the “bulge” rule), and projection predictive inference. Again, we do not look to identify the oracle model size, but rather we hope to select a model size having negligible excess loss compared to the oracle model. Alongside the bulge, we consider the so-called  $2\sigma$  rule, which proposes that one should select the smallest model that is essentially indistinguishable from the model with the best LOO-CV performance, as judged by the two standard error simultaneous confidence interval.

In projection predictive model selection, we first achieved a search path inline with the previous discussion in Section 4.4. We then choose the smallest model size  $k$

---

<sup>7</sup>The crime, ionosphere, and sonar data are freely available from the UCI machine learning benchmarks database: <https://archive.ics.uci.edu/ml/index.php>; the SPECT heart data are available from: <http://odds.cs.stonybrook.edu/>. Pre-processing is as described by Piironen and Vehtari (2017a).

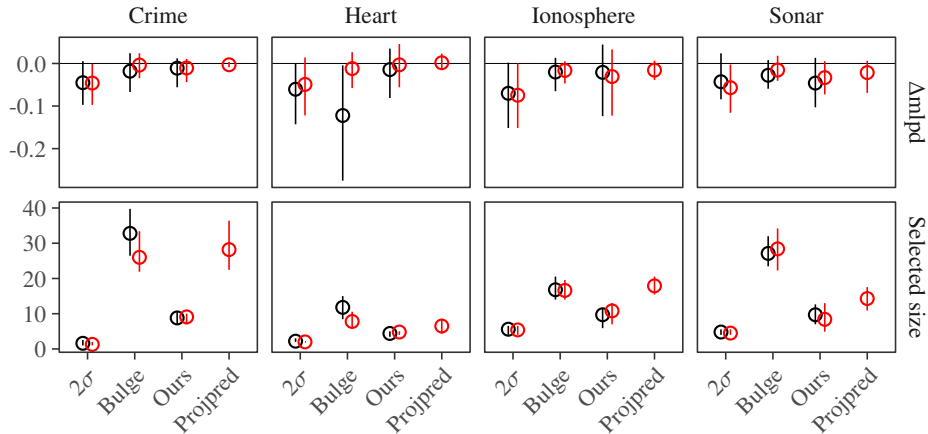


**Figure 7:** Real world datasets. The red lines show the LOO-CV mlpd difference to the reference model over CV folds, while the grey lines show the test mlpd difference along those same folds. In black we see the corrected mlpd difference, produced using our correction from Section 4.1.

whose (projected sub-model) elpd difference point estimate to the reference model (for which we use the full model fitted by standard Bayesian inference) is less than four. This is in line with the recommendation of McLatchie et al. (2023, and discussed more in Appendix A). We compare the model sizes chosen and their stability across forward search cross-validation folds in Figure 8. We make three primary observations:

1. projection predictive inference selects smaller model sizes than the most naïve approaches and does so stably;
2. our correction achieves slightly smaller model sizes than the projective approach with similar mlpd, and requires fewer iterations of forward search than the bulge and  $2\sigma$  rules (while achieving better mlpd than the latter two); and,
3. regardless of model size selection heuristic, the use of shrinkage priors results in smaller and more stable model size selection.

Across all data, the  $2\sigma$  rule chooses the smallest sub-model sizes and the bulge rule chooses the largest. The model sizes selected by our correction are, by contrast, slightly smaller to those selected by projection predictive inference. There is more variation in the model size selected by the bulge rule. In fact, model size selection is generally more variable when using the independent Gaussian priors irrespective of the heuristic. Our correction is similarly stable to projection predictive, which is known to be very stable under the KL forward search heuristic, providing confidence in our own tool. Indeed, the good performance of projection predictive inference has been extensively covered by Piironen and Vehtari (2017a). Our correction identifies models with better predictive performance than the more naïve  $2\sigma$  and bulge rules, and with better stability.



**Figure 8:** Real world datasets. In the bottom row we show the selected model sizes in four real world datasets across different heuristics when using standard independent Gaussian priors (in black), the sparsity-inducing priors (in red). The top row shows the test mlpd difference to the reference model of the model sizes selected by each heuristic.

The model size and mlpd difference estimated by our correction do not seem to be sensitive to the prior, unlike LOO-CV bulge model sizes; this is most prevalent in the heart data where the use of Gaussian priors results in vastly reduced predictive performance when using the bulge rule. Sparsity-inducing priors result in smaller and more stable model sizes being chosen across all data examples. This is likely related to our previous discussion on generative predictive priors: since they regularise in favour of predictive performance, they result in forward search paths that are less variable and less likely to introduce selection-induced bias (although may lead us to slightly under-estimate bias). Their predictive regularisation means that we are likely to more quickly and more stably identify predictive predictors, and ultimately identify more predictive predictor subsets.

The importance of identifying small models is particularly clear in this forward search example: when we must fit all models up to the bulge model size, this requires approximately  $30^2 = 900$  model fits in the sonar and crime case studies. When these models are large, this can mean days of computation time. While the  $2\sigma$  heuristic can identify small models with comparable predictive performance, it still requires fitting up to the bulge model size and then working back. In a word, it does not reduce the computational cost. As such, our bias correction is particularly attractive to those who are able to implement online stopping rules.

It would also have been possible to construct a stopping rule based on the inclusion of zero in the two standard error range of the LOO-CV elpd *difference* estimate. However, such heuristics struggle in the case of many weakly-relevant predictors. In particular, they are liable to identify too-early stopping points and achieve sub-optimal out-of-sample predictive performance when many predictors have weak (but non-zero) predictive information (e.g. in the ionosphere example). This is further discussed in

Appendix [E](#).

While our heuristic was not motivated from the perspective of model size selection, we find that even correcting just some of the selection-induced bias along forward search paths can reveal the behaviour of test performance. Using this corrected LOO-CV mlpd as a proxy for test mlpd is reasonable across these real data examples. As such, model size selection likely does not depend greatly on the assumed asymptotic correlation between LOO-CV and Bayes generalisation loss which we discussed in Section [4.1](#).

## 5 Recommendations

Cross-validation of predictive metrics is a common non-parametric method of making model selection decisions. It is particularly well-suited to cases when robust alternatives such as projection predictive inference are unavailable. Given the uncertainty present in LOO-CV estimates, it is possible that they are not always sufficient to make model selection decisions, or that they might encourage an incorrect decision based purely on noise. We discussed the risk of selecting between two candidate models (Scenario 1), presented a selection-induced bias estimation scheme based on order statistics in the case of selecting many candidate models (Scenario 2), and leveraged it to build a lightweight bias correction method in the forward search setting (Scenario 3).

**Scenario 1: the two model case** In the case where we compare two models only, we can safely make a model selection decision if the elpd difference point estimate between them is greater than four in absolute value (see [Appendix A](#) for more discussion). Should the elpd difference point estimate be less than four in absolute value we have three options: (1) if the models are not nested, combine them by model averaging or stacking; (2) ensure the models' respective priors are reasonable and select the more complex of the two; or, (3) keep them both as a set of best models. This scenario is the least concerning of the three, and in most cases we can safely select the most predictive of the two provided its priors are reasonable.

**Scenario 2: the many model case** In the many model case all previous recommendations apply, and we might further use our proposal from [Section 3.2](#) to determine whether all models are equivalent in terms of predictive performance using order statistics. If we detect non-negligible selection-induced bias, then we do not make a model selection decision and instead perform model averaging or stacking (in the non-nested case). Our estimation procedure was shown to closely match empirical observations, and requires no further computation beyond the sufficient statistics of standard CV estimates. We further proposed a diagnostic to determine when the underlying assumptions of our proposition break down, and more computationally expensive approaches can be used instead.

**Scenario 3: forward search** When the model space is large and we consider nested models, then we should first try projection predictive inference for model size selection. If the reference model is too difficult to fit, or if it belongs to an observation family which can not easily be implemented in the projective framework, then our lightweight selection-induced bias correction can have a direct application: by performing CV forward search and correcting the incremental elpd differences along the search path according to Section 4.1, we can recover an accurate estimate of the out-of-sample performance along the search path. And we do so at a significantly reduced cost to more computationally expensive procedures. In cases of severe over-fitting, and occasionally with shrinkage priors, it under-estimates the true bias slightly. However, we found that we did not need to correct the selection-induced bias perfectly to be able to accurately identify the point of predictive saturation in forward search. Indeed, just recovering the shape of the test elpd was sufficient for stable model size selection. Our procedure was found to effectively and reliably identify the point of over-fitting, and provides an accurate estimate of the bias, closely matching empirical observations.

## 6 Discussion

### 6.1 Limitations

Naturally, our selection-induced bias correction hinges on a number of assumptions which may not always be met in practice. For instance, should models produce correlated predictions then the distributions the elpd point estimates are sampled from may not be independent leading to some error in the order statistic. We already witnessed this to a degree when using shrinkage priors. Likewise, the use of  $\alpha = 0.5$  in the order statistic approximation is known to be conservative, although seemed to fit empirical data well.

Further, while the elpd difference distribution of models with no extra predictive information from the baseline may theoretically be centred on zero, this is unlikely to be the case in practice. Indeed, this centering may in fact be slightly below zero. The choice of selection-induced bias correction based on the negative correlation results from [Watanabe \(2010\)](#) is also heuristic, and may not be representative of the general case. It is not immediately clear how the correction could be alternatively motivated in a robust manner, and we invite future work on the subject given the positive results we were able to achieve using a very simple model. Despite these limitations, our correction performed well across all simulated and real data exercise.

### 6.2 Future directions

It might be possible to improve on our simple correction by modelling the candidate models' LOO-CV elpd differences with a hierarchical model, and making decisions based on the posterior of the meta-model rather than the raw LOO-CV estimates ([Gelman et al., 2012](#); [Schmitt et al., 2023](#)). As well as trying to collapse all models into one group of in-differentiable models, as is the aim of partially-pooled models, one

might produce a correlation-aware hierarchical model capable of identifying clusters of models exhibiting similar predictive performance. For example, one might consider a multivariate model with a similarity-based covariance matrix.

A more attainable extension is perhaps soft-thresholding the point at which we correct the LOO-CV elpd along the search path. The Fisher–Tippett–Gnedenko theorem tells us that the limiting distribution of properly-normalised maxima from a sequence of iid standard Gaussian observations follows a generalised extreme value distribution. Using this, or posterior weights from some meta-model, we might achieve some probability of having observed a truly predictive model from a set of candidates. This, as opposed to using the order statistic’s mean as a hard threshold, would allow us to weight the correction we apply. While we did not observe any issue with the hard threshold we proposed, it is possible that in instances of many weakly-relevant predictors it is harder to separate between only negligibly relevant models, and our correction rule may be too rigid.

## References

- Aguilar, J. E. and Bürkner, P.-C. (2023). Intuitive joint priors for Bayesian linear multilevel models: The R2D2M2 prior. *Electronic Journal of Statistics*, 17(1).
- Ambroise, C. and McLachlan, G. J. (2002). Selection bias in gene extraction on the basis of microarray gene-expression data. *Proceedings of the National Academy of Sciences*, 99(10):6562–6566.
- Arlot, S. and Celisse, A. (2010). A survey of cross-validation procedures for model selection. *Statistics surveys*, 4:40–79.
- Barbieri, M. M. and Berger, J. O. (2004). Optimal predictive model selection. *The Annals of Statistics*, 32(3).
- Bernardo, J. M. and Smith, A. F. M. (1994). *Bayesian Theory*. John Wiley & Sons.
- Blom, G. (1960). Statistical estimates and transformed beta-variables. *Biometrika*, 47(1/2):210.
- Brown, P. J., Vannucci, M., and Fearn, T. (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 60(3):627–641.
- Burnham, K. P. and Anderson, D. R. (2002). *Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach*. Springer, 2nd edition.
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1):1–28. tex.encoding: UTF-8.
- Bürkner, P.-C., Gabry, J., and Vehtari, A. (2020). Approximate leave-future-out cross-validation for Bayesian time series models. *Journal of Statistical Computation and Simulation*, 90(14):2499–2523. arXiv:1902.06281 [stat].

- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2009). Handling sparsity via the horseshoe. In van Dyk, D. and Welling, M., editors, *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 73–80, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA. PMLR.
- Catalina, A., Bürkner, P., and Vehtari, A. (2021). Latent space projection predictive inference. arXiv:2109.04702 [stat].
- Catalina, A., Bürkner, P.-C., and Vehtari, A. (2020). Projection predictive inference for generalized linear and additive multilevel models. arXiv:2010.06994 [stat].
- Cawley, G. C. and Talbot, N. L. C. (2010). On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11:2079–2107.
- Cooper, A., Simpson, D., Kennedy, L., Forbes, C., and Vehtari, A. (2023). Cross-validatory model selection for Bayesian autoregressions with exogenous regressors. arXiv:2301.08276 [stat].
- Dupuis, J. A. and Robert, C. P. (2003). Variable selection in qualitative models via an entropic explanatory power. *Journal of Statistical Planning and Inference*, 111(1-2):77–94.
- Efron, B. and Tibshirani, R. (1993). *An introduction to the bootstrap*. Number 57 in Monographs on statistics and applied probability. Chapman & Hall, New York.
- Galatenko, V. V., Shkurnikov, M. Y., Samatov, T. R., Galatenko, A. V., Mityakina, I. A., Kaprin, A. D., Schumacher, U., and Tonevitsky, A. G. (2015). Highly informative marker sets consisting of genes with low individual degree of differential expression. *Scientific Reports*, 5(1):14967.
- Garthwaite, P. H. and Mubwandarikwa, E. (2010). Selection of weights for weighted model averaging: Prior weights for weighted model averaging. *Australian & New Zealand Journal of Statistics*, 52(4):363–382.
- Geisser, S. and Eddy, W. F. (1979). A predictive approach to model selection. *Journal of the American Statistical Association*, 74(365):153–160.
- Gelfand, A. and Ghosh, S. K. (1998). Model choice: a minimum posterior predictive loss approach. *Biometrika*, 85(1):1–11.
- Gelfand, A. E. (1996). Model determination using sampling-based methods. *Markov chain Monte Carlo in practice*, 4:145–161. Publisher: London.
- Gelfand, A. E., Dey, D. K., and Chang, H. (1992). Model determination using predictive distributions with implementation via sampling-based methods. Technical report, Stanford Univ CA Dept of Statistics.

- Gelman, A. (2022). I'm skeptical of that claim that "Cash aid to poor mothers increases brain activity in babies".
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. Chapman and Hall/CRC, 3rd edition.
- Gelman, A., Hill, J., and Yajima, M. (2012). Why we (usually) don't have to worry about multiple comparisons. *Journal of Research on Educational Effectiveness*, 5(2):189–211. Publisher: Routledge \_eprint: <https://doi.org/10.1080/19345747.2011.618213>.
- Gelman, A., Hwang, J., and Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24(6):997–1016.
- Gelman, A., Simpson, D., and Betancourt, M. (2017). The prior can often only be understood in the context of the likelihood. *Entropy*, 19(10):555.
- Gelman, A., Xiao-Li, M., and Stern, H. S. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, 6(4):733–760.
- George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889.
- Gorman, R. and Sejnowski, T. J. (1988). Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks*, 1(1):75–89.
- Goutis, C. (1998). Model choice in generalised linear models: a Bayesian approach via Kullback-Leibler projections. *Biometrika*, 85(1):29–37.
- Han, C. and Carlin, B. P. (2001). Markov chain Monte Carlo methods for computing Bayes factors: A comparative review. *Journal of the American Statistical Association*, 96(455):1122–1132.
- Harrell, F. E. (2001). *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. Springer Series in Statistics. Springer New York, New York, NY.
- Harter, H. L. (1961). Expected values of normal order statistics. *Biometrika*, 48(1/2):151.
- Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1999). Bayesian model averaging: a tutorial (with comments by M. Clyde, David Draper and E. I. George, and a rejoinder by the authors. *Statistical Science*, 14(4).
- Jeffreys, H. (1998). *Theory of probability*. Oxford classic texts in the physical sciences. Clarendon Press ; Oxford University Press, Oxford [Oxfordshire] : New York, 3rd ed edition.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.



- Key, J., Pericchi, L., and Smith, A. F. M. (1999). Bayesian model choice: What and why? *Bayesian Statistics*.
- Krzysztof Cios, Lukasz Kurgan, L. G. (2001). SPECTF heart data. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5N015>.
- Laud, P. W. and Ibrahim, J. G. (1995). Predictive model selection. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):247–262.
- Le, T. M. and Clarke, B. S. (2022). Model averaging is asymptotically better than model selection for prediction. *Journal of Machine Learning Research*, 23(33):1–53.
- Leamer, E. E. (1979). Information criteria for choice of regression models: A comment. *Econometrica*, 47(2):507.
- Lindley, D. V. (1968). The choice of variables in multiple regression. *Journal of the Royal Statistical Society: Series B (Methodological)*, 30(1):31–53.
- Marriott, J. M., Spencer, N. M., and Pettitt, A. N. (2001). A Bayesian approach to selecting covariates for prediction. *Scandinavian Journal of Statistics*, 28(1):87–97.
- McLatchie, Y., Rögnvaldsson, S., Weber, F., and Vehtari, A. (2023). Robust and efficient projection predictive inference. arXiv:2306.15581 [stat].
- Merkle, E. C., Furr, D., and Rabe-Hesketh, S. (2019). Bayesian Comparison of Latent Variable Models: Conditional Versus Marginal Likelihoods. *Psychometrika*, 84(3):802–829.
- Narisetty, N. N. and He, X. (2014). Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics*, 42(2).
- Nilsson, N. J. (1998). *Artificial Intelligence: A New Synthesis*. Elsevier.
- Ntzoufras, I., Forster, J. J., and Dellaportas, P. (2000). Stochastic search variable selection for log-linear models. *Journal of Statistical Computation and Simulation*, 68(1):23–37.
- Oelrich, O., Ding, S., Magnusson, M., Vehtari, A., and Villani, M. (2020). When are Bayesian model probabilities overconfident? arXiv:2003.04026 [math, stat].
- O’Hara, R. B. and Sillanpää, M. J. (2009). A review of Bayesian variable selection methods: what, how and which. *Bayesian Analysis*, 4(1).
- Piironen, J., Paasiniemi, M., and Vehtari, A. (2020). Projective inference in high-dimensional problems: Prediction and feature selection. *Electronic Journal of Statistics*, 14(1). arXiv:1810.02406 [cs, stat].
- Piironen, J. and Vehtari, A. (2017a). Comparison of Bayesian predictive methods for model selection. *Statistics and Computing*, 27(3):711–735.

- Piironen, J. and Vehtari, A. (2017b). Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics*, 11(2).
- Raftery, A. E. and Zheng, Y. (2003). Discussion: Performance of Bayesian model averaging. *Journal of the American Statistical Association*, 98(464):931–938.
- Redmond, M. and Baveja, A. (2002). A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research*, 141(3):660–678.
- Reunanen, J. (2003). Overfitting in making comparisons between variable selection methods. *Journal of Machine Learning Research*, 3:1371–1382.
- Robert, C. (2014). Projective covariate selection.
- Royston, J. P. (1982). Algorithm AS 177: Expected normal order statistics (exact and approximate). *Applied Statistics*, 31(2):161.
- Scarrott, C. and MacDonald, A. (2012). Review of extreme value threshold estimation and uncertainty quantification. *REVSTAT-Statistical Journal*, pages 33–60 Pages. Artwork Size: 33–60 Pages Publisher: REVSTAT-Statistical Journal.
- Schmitt, M., Radev, S. T., and Bürkner, P.-C. (2023). Meta-uncertainty in bayesian model comparison. In Ruiz, F., Dy, J., and van de Meent, J.-W., editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 11–29. PMLR.
- Scholz, M. and Bürkner, P.-C. (2022). Prediction can be safely used as a proxy for explanation in causally consistent Bayesian generalized linear models. arXiv:2210.06927 [stat].
- Shao, J. (1993). Linear model selection by cross-validation. *Journal of the American Statistical Association*, 88(422):486–494.
- Sigillito, V., Wing, S., Hutton, L. V., and Baker, K. (1989). Classification of radar returns from the ionosphere using neural networks. *Johns Hopkins APL Technical Digest*, 10:262–266.
- Sivula, T., Magnusson, M., Matamoros, A. A., and Vehtari, A. (2022). Uncertainty in Bayesian leave-one-out cross-validation based model comparison. arXiv:2008.10296 [stat].
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):583–639.
- Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133.

- Tibshirani, R. J. and Tibshirani, R. (2009). A bias correction for the minimum error rate in cross-validation. *The Annals of Applied Statistics*, 3(2).
- Troller-Renfree, S. V., Costanzo, M. A., Duncan, G. J., Magnuson, K., Gennetian, L. A., Yoshikawa, H., Halpern-Meekin, S., Fox, N. A., and Noble, K. G. (2022). The impact of a poverty reduction intervention on infant brain activity. *Proceedings of the National Academy of Sciences*, 119(5):e2115649119.
- Vehtari, A., Gelman, A., and Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5):1413–1432.
- Vehtari, A. and Lampinen, J. (2002). Bayesian model assessment and comparison using cross-validation predictive densities. *Neural Computation*, 14(10):2439–2468.
- Vehtari, A. and 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., and Gabry, J. (2022). Pareto Smoothed Importance Sampling. arXiv:1507.02646 [stat].
- Wang, W. and Gelman, A. (2015). Difficulty of selecting among multilevel models using predictive accuracy. *Statistics and Its Interface*, 8(2):153–160.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. arXiv:1004.2316 [cs].
- Watanabe, S. (2013). A widely applicable Bayesian information criterion. *Journal of Machine Learning Research*, page 31.
- Weber, F. and Vehtari, A. (2023). Projection predictive variable selection for discrete response families with finite support. arXiv:2301.01660 [stat].
- Yanchenko, E., Bondell, H. D., and Reich, B. J. (2023a). R2D2 goes to space! A principled approach to setting prior distributions on spatial parameters. arXiv:2301.09951 [stat].
- Yanchenko, E., Bondell, H. D., and Reich, B. J. (2023b). The R2D2 prior for generalized linear mixed models. arXiv:2111.10718 [stat].
- Yao, Y., Carvalho, L. M., Mesquita, D., and McLatchie, Y. (2023). Locking and quacking: Stacking Bayesian model predictions by log-pooling and superposition. arXiv:2305.07334 [cs, stat].
- Yao, Y., Pirš, G., Vehtari, A., and Gelman, A. (2021). Bayesian hierarchical stacking: Some models are (somewhere) useful. *Bayesian Analysis*, -1(-1). arXiv:2101.08954 [cs, stat].

- Yao, Y., Vehtari, A., Simpson, D., and Gelman, A. (2018). Using stacking to average Bayesian predictive distributions (with discussion). *Bayesian Analysis*, 13(3):917–1007.
- Yates, L. A., Richards, S. A., and Brook, B. W. (2021). Parsimonious model selection using information theory: a modified selection rule. *Ecology*, 102(10).
- Zhang, J. and Stephens, M. A. (2009). A new and efficient estimation method for the generalized Pareto distribution. *Technometrics*, 51(3):316–325.
- Zhang, Y. D., Naughton, B. P., Bondell, H. D., and Reich, B. J. (2022). Bayesian regression using a prior on the model fit: The R2-D2 shrinkage prior. *Journal of the American Statistical Association*, 117(538):862–874.

## A Connection between elpd difference and model averaging

Choosing a model based on elpd point estimate, as we have proposed throughout this paper, has a deeper connection to Bayesian model averaging. Suppose we are in the  $K = 2$  case, we might say that one model is sufficiently better than another to feel safe making a model selection decision if

$$\left| \widehat{\Delta \text{elpd}}_{\text{LOO}}(M_a, M_b | y) \right| < 4. \quad (\text{A1})$$

This was initially considered by [Sivula et al. \(2022\)](#) in terms of the reliability of the standard error estimate: when the elpd difference point estimate is this low, it is likely that the elpd difference *standard error* estimate is too small. [McLatchie et al. \(2023\)](#) later used this flavour of selection heuristic in the context of projection predictive inference.

Suppose we have once more the situation where we have two models  $M_a$  and  $M_b$ . If we consider a normal approximation of the LOO-CV elpd difference estimate, then the probability that the more predictive model, say  $M_a$ , is truly better than the other is

$$\mathbb{P}(\text{elpd}(M_a | y) \geq \text{elpd}(M_b | y)) = \int_{-\infty}^0 \text{normal}\left(z \mid \widehat{\Delta \text{elpd}}_{\text{LOO}}, \text{se}\left(\widehat{\Delta \text{elpd}}_{\text{LOO}}\right)\right) dz. \quad (\text{A2})$$

Suppose the standard error is two for demonstration purposes, then when the point estimate is 4, the probability of one model being better than another is approximately 98%. This is shown in [Figure A1](#).

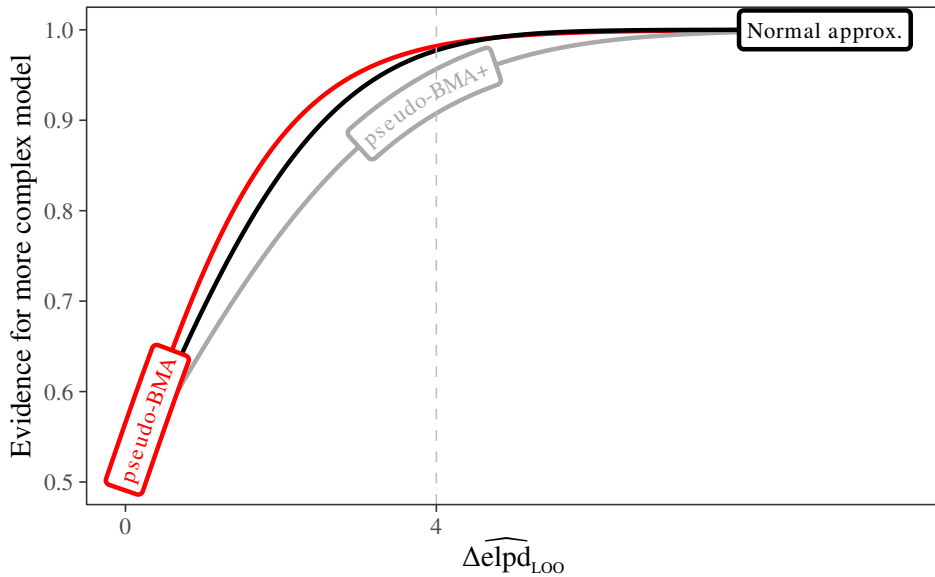
Likewise, [Yao et al. \(2021\)](#) present an AIC-type model weighting score which uses LOO-CV predictive densities which they call pseudo-Bayesian model averaging (pseudo-BMA) weights. The weight of the more predictive model of the two compared is

$$w_k = \frac{1}{1 + \exp\left(-\widehat{\Delta \text{elpd}}_{\text{LOO}}\right)}. \quad (\text{A3})$$

Once more, we find that should we have two models with an elpd difference point estimate of 4, then the psuedo-BMA weight of the more predictive model is again approximately 0.98. This weighting system, however, ignores the uncertainty associated with the point estimate. To remedy this, [Yao et al. \(2021\)](#) present what they call pseudo-BMA+ weights that integrate over this uncertainty, where the weight of the more predictive model is now

$$w_k^+ = \int_{-\infty}^{\infty} \frac{\text{normal}\left(z \mid 0, \text{se}\left(\widehat{\Delta \text{elpd}}_{\text{LOO}}\right)\right)}{1 + \exp\left(-\widehat{\Delta \text{elpd}}_{\text{LOO}} - z\right)} dz. \quad (\text{A4})$$

Assuming a standard error of two as before, the weight of the more predictive model is again over 90%, albeit lower than with pseudo-BMA and normal approximation



**Figure A1:** Theoretical connection between elpd difference and BMA. The relationship between pseudo-BMA weights (Equation A3), the pseudo-BMA+ weights (Equation A4), and the probability of one model being better than another based on the normal approximation of the elpd difference distribution (Equation A2). We fix the standard error of the elpd difference normal approximation to be  $se(\widehat{\Delta elpd}_{LOO}) = 2$  for demonstration purposes. For this choice of standard error, Equation A2 acts very similarly to the probit function (naturally, as it is equivalent to the Gaussian cumulative density function computed up to zero). This probit is precisely the pseudo-BMA weight of the more complex model, and mildly altered becomes the pseudo-BMA+ weights. At the point  $\widehat{\Delta elpd}_{LOO} \approx 4$ , and the probability of it being better than the less complex model is approximately 98%.

calculations. This is again shown in Figure A1. Yao et al. (2021) show that their pseudo-BMA+ weightings perform better than standard pseudo-BMA, and even with this stricter measurement, looking at just the point estimate of the elpd difference is safe when making model selection decisions in the two-model case.

Now we can reflect selection-induced bias in terms of the value 4: if the bias is big enough to impact a decision made based on this *ad hoc* rule-of-four, then we may start to falsely give some models zero weight in pseudo-BMA. For instance, if the bias is much less than 4, then we could say it is negligible. But if the bias grows above 4 then we have problems using the elpd estimates for model averaging. We do not, however, propose this number as a hard threshold.

## B Empirical demonstration of order statistics

We presently consider data from a study on the effect of poverty reduction on infant brain activity (Troller-Renfree et al., 2022). In this study, the authors report the impact

of a cash incentive to mothers on the brain activity in the first year of life of their children. The mothers were randomised to receive an monthly, unconditional cash grant of either “significant” or “nominal” size. The children’s brain activities were then measured by resting electroencephalography (EEG) after a year to ascertain whether these cash transfers changed infant brain activity. We build a collection of different models on the “alpha brain band” data, varying the number of predictors, the observation family, and the data pre-processing.

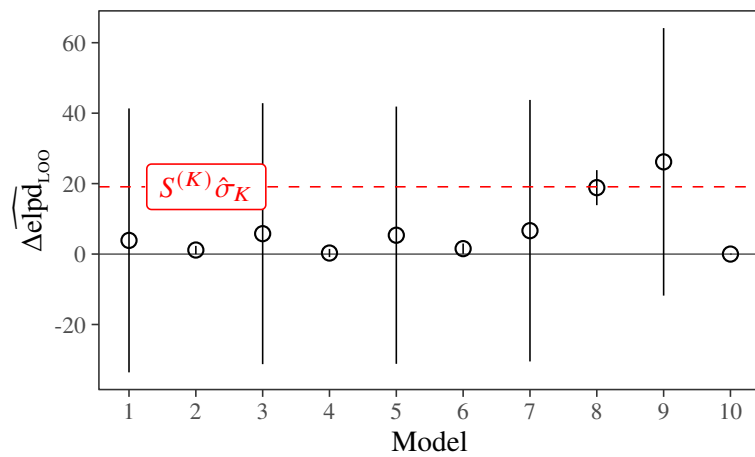
In particular, we create a pre-treatment score similarly to [Gelman \(2022\)](#) by combining birth weight, gestational age, the mother’s health, and amount of smoking and drinking during pregnancy. These predictors are scaled to have zero mean and unit standard deviation, and the signs of smoking and drinking predictors are switched before being summed. We construct generalised linear models (GLMs) with just the treatment effect in a first instance, before adding the pre-treatment score, household income, birth weight, and mother’s race incrementally. We then also investigate GLMs with a Gaussian observation family, and a log-normal observation family. Having fit these 10 models, we then identify the worst-performing model in terms of LOO-CV elpd point estimate, and use this as the baseline model.

In [Figure B1](#) we show the LOO-CV elpd differences of the remaining nine models to the baseline (shown in the plain dot at index 10). We compute the order statistic-based bias estimate,  $S^{(K)}\hat{\sigma}_K$  with  $K = 10$ , as previously described and plot this next to the model LOO-CV elpd differences. These results would have us believe that at least one of the models (specifically the ninth) is significantly better than the baseline, and that in this case we can safely select it. The eighth model lies very close to the border marked out the order statistic; if we were concerned about selection-induced bias in this case, we could average or stack these two seemingly more predictive models rather than select only one.

## C Bias estimation under different assumed asymptotic cross-validated correlation regimes

As previously mentioned, the origin of our bias correction lies in the work of [Watanabe \(2010, Table 3\)](#). They discuss a notable negative correlation between cross-validated scores and average Bayes generalisation loss, observed in just one data instance. This arises due to excluding an influential observation from training, leading to a skewed posterior predictive distribution and the observation’s seemingly anomalous appearance in the test set. The correction of LOO-CV elpd estimates thus necessitates a negatively-correlated correction. In essence, the adjusted elpd difference should be mirrored across the median of candidate differences. The choice of magnitude 1.5 in our bias correction reflects a negative correlation around  $-0.8$  noted by [Watanabe \(2010\)](#), albeit from a single data instance. Recognising this limitation, we repeat the correction in the simulated forward search experiments with values  $\{1, 1.5, 2\}$ , corresponding to correlations  $\{0, -0.5, -1\}$ .

In [Figure C1](#) we see how these different values perform across different  $n$  and  $\rho$



**Figure B1:** EEG cash incentive data. We show in the point ranges the LOO-CV elpd differences and difference standard errors of 10 models with respect to the baseline model (chosen as the model with the lowest LOO-CV elpd, and located at index 10). We plot the expected highest elpd difference under the assumption that no model is better than the baseline with order statistics in the red dashed line. This identifies at least one of the models (the ninth) to be significantly better than the baseline model.



regimes within our simulated forward search experiment. In the case of independent Gaussian priors over the regression coefficients, and in this specific case, it seems the asymptotic correlation may be closer to  $-1$ , as using a bias correction with magnitude 2 fits the data best across all  $n$  and  $\rho$  regimes. Using smaller values (including the compromise value 1.5 we eventually chose) are liable to under-estimate the bias in the case  $n = p$ . The difference between magnitudes lessens as  $n$  grows, and we could already be comfortable using any of them once  $n > 4p$  (again, in this lone data example). When the R2D2 prior is used the results are similar. The main difference we observe is that the difference between the three magnitudes is much smaller, regardless of the  $n$  regime. The use of sparsity-inducing priors remedies selection-induced bias, and we do not have to be as concerned as in the high-risk regimes of independent Gaussian priors.

It is not clear how the magnitude in our bias estimate can be constructed generally in a principled manner. Equivalently, it is not clear how to estimate the asymptotic correlation between LOO-CV estimates and Bayes generalisation loss for a given data task, and we leave this as a promising area of future work.

## D Sparsity-inducing priors

Fitting high-dimensional models to sparse data is generally challenging. When using joint shrinkage priors, practitioners explicit a prior belief that applies to both the parameter space and the predictive space. One such prior is the R2D2 prior (Aguilar and Bürkner, 2023; Yanchenko et al., 2023a,b; Zhang et al., 2022), which allows statisticians to express prior beliefs about the model’s  $R^2$ , and the number of predictors required to achieve it. By doing so, the risk of over-fitting the data is reduced compared to using simpler independent Gaussians for regression coefficients, which implicitly encode a (surprisingly) high *a priori*  $R^2$ . The R2D2 model for a linear regression is:

$$y_i \sim \text{normal}(\beta_0 + \sum_{k=1}^p x_{k,i}\beta_k, \sigma^2) \quad (\text{D1})$$

$$\beta_0 \sim \text{Student-}t_3(0, 2.5) \quad (\text{D2})$$

$$\sigma \sim \text{Student-}t_3^+(0, 2.5) \quad (\text{D3})$$

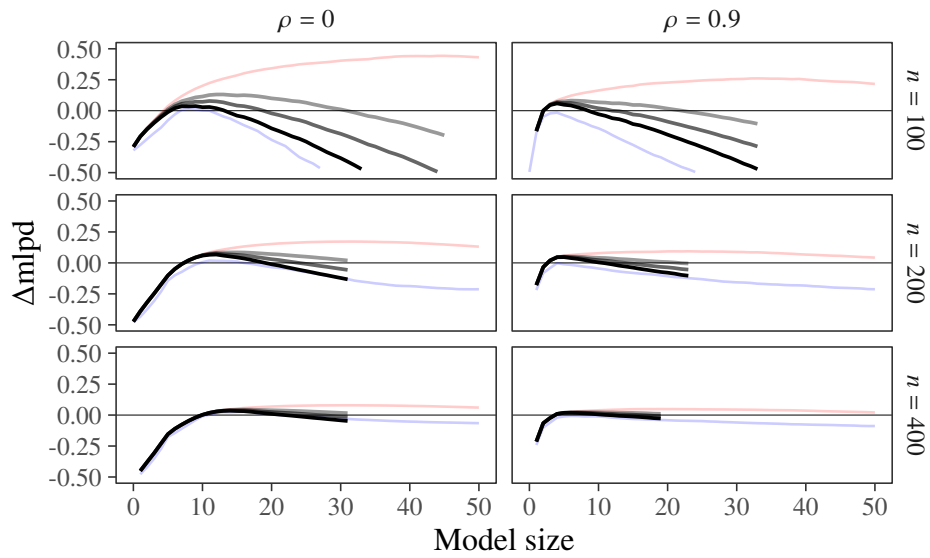
$$\beta_k \sim \text{normal}(0, \sigma^2\tau^2\phi_k) \quad (\text{D4})$$

$$\tau^2 = \frac{R^2}{1 - R^2} \quad (\text{D5})$$

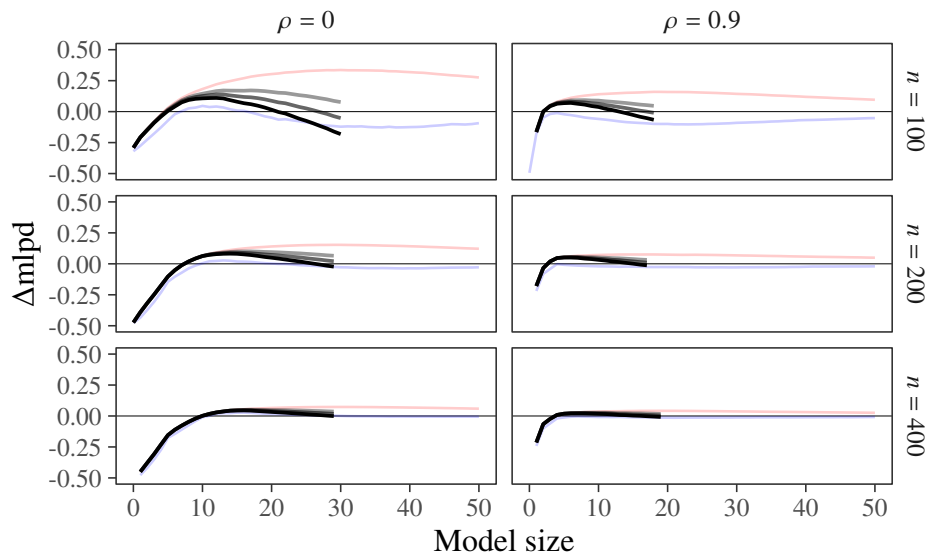
$$R^2 \sim \text{beta}(\mu_{R^2}, \varphi_{R^2}) \quad (\text{D6})$$

$$\phi \sim \text{Dirichlet}(\xi, \dots, \xi). \quad (\text{D7})$$

We employ a mean and a pseudo-precision parameterisation of the beta distribution as opposed to shape parameters  $a, b > 0$ . The Dirichlet concentration parameter  $\xi$ , which is constant across all predictors, controls the sparsity in regression coefficients. Larger values of  $\xi$  result in more uniformly distributed regression coefficients (in terms of magnitude), usually with more predictors explaining the response’s variance.



(a) Independent Gaussian priors over regression coefficients.



(b) R2D2 prior over regression coefficients.

**Figure C1:** Forward search simulation experiment across different priors, and under different assumed asymptotic correlations between LOO-CV estimate and average Bayes generalisation loss. The red line show the LOO-CV mlpd difference to the reference model, and the blue lines show the test mlpd difference. In black we show the mlpd difference corrected according to our proposal with assumed correlation  $-1$ , in lighter grey we assume correlation  $-0.5$ , and in the lightest grey an assumed  $0$  correlation. Higher assumed correlation results in more correction while assuming, for example, a zero correlation under-estimates the selection-induced in most cases.

Data	Task	$n$	$p$	Prior	Hyper-parameters
Crime	Regression	1992	102	R2D2	$\xi = 1; \mu_{R^2} = 0.2; \varphi_{R^2} = 5$
Simulated	Regression	varies	100	R2D2	$\xi = 0.5; \mu_{R^2} = 0.5; \varphi_{R^2} = 4$
Heart	Classification	224	44	RHS	$m_{\text{eff}}/p = 0.3$
Ionosphere	Classification	351	33	RHS	$m_{\text{eff}}/p = 0.6$
Sonar	Classification	208	60	RHS	$m_{\text{eff}}/p = 0.3$

**Table D1:** Specification of datasets, priors, and hyper-parameters used to produce the reference models throughout the paper.

Alongside the R2D2 prior, [Piironen and Vehtari \(2017b\)](#) propose the regularised horseshoe (RHS) prior. Here, we define the model as

$$\beta_k \sim \text{normal}(0, \tau^2 \tilde{\lambda}_k^2) \quad (\text{D8})$$

$$\tilde{\lambda}_k^2 = \frac{c^2 \lambda_k^2}{c^2 + \tau^2 \lambda_k^2} \quad (\text{D9})$$

$$\lambda_k^2 \sim \text{Cauchy}^+(0, 1). \quad (\text{D10})$$

Across all binary classification case studies in this paper, we set the degrees of freedom of Student- $t$  prior of the global and local shrinkage parameters to be one, and the degrees of freedom of the slab to be 100 so as to approach the horseshoe prior of [Carvalho et al. \(2009\)](#). Rather than define our prior over  $\tau$  and  $c$ , we instead prefer to consider the number of effective non-zero parameters, denoted  $m_{\text{eff}}$ , or rather the proportion of effectively non-zero parameters,  $m_{\text{eff}}/p$ . Table D1 presents the hyper-parameters (for both the R2D2 and RHS priors) used in the real world case studies.<sup>8</sup>

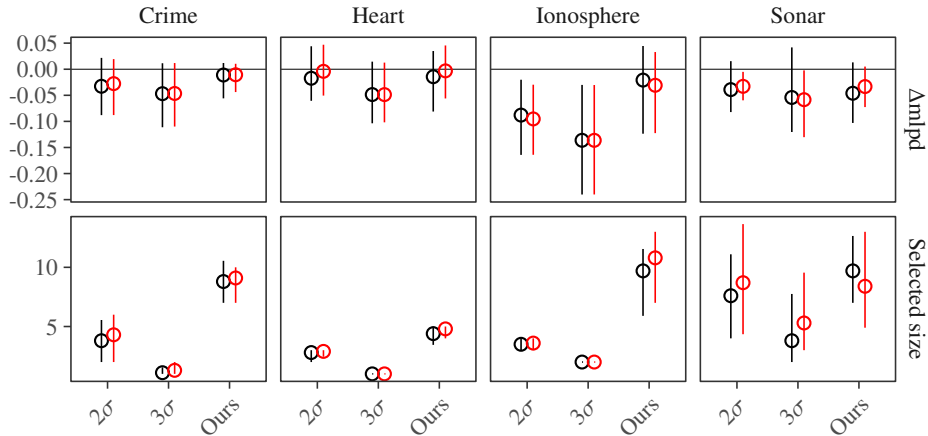
## E Alternative stopping rules

The previously presented stopping rules which involve choosing the ‘‘bulge’’ model, its associated  $2\sigma$  rule, and projection predictive inference all suffer from the same flaw: they require running the forward search up to (at least) the bulge model. A simple alternative would be to run the forward search until none of the candidates improves predictive performance as judged by the two or three standard errors of their elpd difference estimates. Formally, we stop the search if there is *no* candidate model,  $M_k$ , such that

$$\widehat{\Delta \text{elpd}}_{\text{LOO}}(M_k, M_{\text{previous}} \mid y) - 2 \times \text{se} \left( \widehat{\Delta \text{elpd}}_{\text{LOO}}(M_k, M_{\text{previous}} \mid y) \right) \geq 0, \quad (\text{E1})$$

where  $M_{\text{previous}}$  denotes the model chosen in the previous step. Or using a magnitude of 3 rather than 2 in Equation E1 to be even more conservative. Let us call these two rules  $2\sigma_{\Delta}$  and  $3\sigma_{\Delta}$  respectively.

<sup>8</sup>An interactive visualisation of the R2D2 prior is freely available at [https://solviro.shinyapps.io/R2D2\\_shiny/](https://solviro.shinyapps.io/R2D2_shiny/).



**Figure E1:** Real world datasets. The selected model sizes in three real world datasets across different heuristics when using standard independent Gaussian priors (in black) and the R2D2 prior (in red). Our correction identifies larger model sizes than the  $2\sigma_{\Delta}$  and  $3\sigma_{\Delta}$ , but with higher mlpd point estimates. The incremental stopping rules are liable to select too-small models, resulting in sub-optimal out-of-sample performance and reduced intuition.

In Figure E1 we show the model sizes they select, along with the test mlpd of those models selected. We find that these rules identify much smaller model sizes than our proposal, although do so at the cost of predictive performance. Our proposal finds models with higher mlpd point estimates across all datasets, and with sizes only slightly larger. Further, we find that the  $2\sigma_{\Delta}$  and  $3\sigma_{\Delta}$  rules can be volatile, and less stable than our proposal.

It is known that there types of heuristics (the  $2\sigma_{\Delta}$  and  $3\sigma_{\Delta}$  rules) struggle under collinearity of predictors, and in the presence of weakly (but truly) relevant predictors. This is seen in these real data cases, where the heuristics choose model sizes that, while being computationally beneficial, may be overly-sparse. For these reasons, we do not consider them to sufficiently robust (in a qualitative sense) for use in real data experiments where the structure of the predictors is unknown.