

# Double Spike Dirichlet Priors for Structured Weighting

Huiming Lin

Meng Li

*Department of Statistics*

*Rice University*

*Houston, TX 77005, USA*

HL68@RICE.EDU

MENG@RICE.EDU

**Editor:** Ryan Adams

## Abstract

Assigning weights to a large pool of objects is a fundamental task in a wide variety of applications. In this article, we introduce the concept of *structured* high-dimensional probability simplexes, in which most components are zero or near zero and the remaining ones are close to each other. Such structure is well motivated by (i) high-dimensional weights that are common in modern applications, and (ii) ubiquitous examples in which equal weights—despite their simplicity—often achieve favorable or even state-of-the-art predictive performance. This particular structure, however, presents unique challenges partly because, unlike high-dimensional linear regression, the parameter space is a simplex and pattern switching between partial constancy and sparsity is unknown. To address these challenges, we propose a new class of double spike Dirichlet priors to shrink a probability simplex to one with the desired structure. When applied to ensemble learning, such priors lead to a Bayesian method for structured high-dimensional ensembles that is useful for forecast combination and improving random forests, while enabling uncertainty quantification. We design efficient Markov chain Monte Carlo algorithms for implementation. Posterior contraction rates are established to study large sample behaviors of the posterior distribution. We demonstrate the wide applicability and competitive performance of the proposed methods through simulations and two real data applications using the European Central Bank Survey of Professional Forecasters data set and a data set from the UC Irvine Machine Learning Repository (UCI).

**Keywords:** ensemble method, forecast combination puzzle, high-dimensional simplex, posterior contraction, random forests

## 1. Introduction

There is an immense variety of applications in which the key is to assign weights to a large pool of objects, such as in model averaging, forecast combination, and tree ensembles. High-dimensional models and sparse learning methods have gained extensive attention in the literature; however, much less focus has been placed on the simplex setting, where the entries of the unknown parameter or weights are non-negative and sum up to one. In this paper, we are interested in *structured* high-dimensional simplexes with *sparsity* and *partial constancy*, i.e., most of the weights are zero or near zero and the remaining ones are close to each other.

Such structure is well motivated by the recurrent observation in various areas that equal weights often achieve favorable or even state-of-the-art empirical performance in practice, and the weights may be high-dimensional in modern applications.

The first notable example is forecast combination in economics. There is a rich literature on how to optimally combine individual forecasts, and weights on the probability simplex are commonly used as they lead to unbiased estimation if each forecast is unbiased (Bates and Granger, 1969; Granger and Ramanathan, 1984). One remarkable observation is that the simplest combination method using equal weights constantly outperforms more complicated weighting schemes in terms of mean-squared forecast error, a phenomenon that is known as the ‘forecast combination puzzle’ (Clemen, 1989; Stock and Watson, 2004; Smith and Wallis, 2009; Makridakis et al., 2020). On the other hand, the number of forecasters can be large compared to the number of historical records in many applications, pointing to the high-dimensional regime where sparsity may be desired.

Another example that may benefit from such structure is the well-known random forests method (Breiman, 2001), an ensemble method that first builds then averages a large collection of de-correlated trees. Random forests have long been considered as one of the most successful general-purpose supervised machine learning methods (Svetnik et al., 2003; Segal, 2004; Palmer et al., 2007; Biau and Scornet, 2016). It is a special case of high-dimensional forecast combination with each tree regarded as an individual forecast. Breiman (2001) proved that the generalization error converges almost surely to a limit as the number of trees goes to infinity, which implied that above a certain number of trees, adding more trees in the forest does not improve accuracy much. This has invited discussions on whether a subset of individual trees can outperform the whole forest in prediction (Kulkarni and Sinha, 2012), in addition to its apparent advantage of improved parsimony. For example, Zhang and Wang (2009) and Bernard et al. (2009) developed methods to reduce the forest size while maintaining the prediction accuracy comparable to the initial forest. However, they did not incorporate both sparsity and partial constancy into ensemble methods such as random forests, which may lead to substantial predictive gain using real world data. Furthermore, the joint statistical inference on this constrained high-dimensional space with uncertainty quantification has not been studied in the literature.

Such structure, high-dimensional simplex with partial constancy and sparsity, is indeed ubiquitous as equal weights and high-dimensional settings that necessitate sparsity are routinely encountered in modern applications, and convex combination often enjoys stability and theoretical guarantees (Bunea et al., 2007; Polley and Van der Laan, 2010). Throughout this article, we assume the true parameter obeys this structure. While structured simplexes can be integrated into any model or flexible loss functions beyond the likelihood as long as the parameter space is a probability simplex, we will focus on high-dimensional ensemble learning for concreteness. This setting has wide applications; for example, it is a common framework in forecast combination (Granger and Ramanathan, 1984; Diebold and Pauly, 1990), and it serves as a main building block for *super learner* in the model aggregation literature (Van der Laan et al., 2007; Polley and Van der Laan, 2010).

When the high-dimensional parameters are constrained on a simplex, the popular  $\ell_1$  penalty for sparse learning in high-dimensional regression falls short, as the  $\ell_1$  norm of the parameters is equal to one. To circumvent this difficulty and achieve sparsity, one may alternatively penalize certain transformations of the parameters, say their square root.

However, solving such constrained optimization problem may be far from trivial as any regularizer that promotes exact sparsity under simplex constraints cannot be convex (Li et al., 2020). There is recent literature on simplex structure, including sparse projection onto the simplex (Kyriallidis et al., 2013) and constrained optimization (Clarkson, 2010). In the high-dimensional setting, Li et al. (2020) discussed strategies to estimate sparse simplexes, including to couple empirical risk minimization with thresholding or re-weighted  $\ell_1$  regularization. Conflitti et al. (2015) studied the high-dimensional forecast combination problem and found that the probability simplex constraint helped stabilize the determination of the optimal weights. However, the constraint of partial constancy is not studied, and uncertainty quantification remains a challenge.

A natural approach to carry out estimation constrained on the structured simplex space is to adopt a two-step strategy of ‘sparsity first, then partial constancy’: first select a subset of forecasts, then shrink the remaining ones toward a constant while maintaining the simplex constraint. There is a rich menu of methods for the first step, for example, lasso (Tibshirani, 1996); this approach has been implemented in Diebold and Shin (2019), and hereafter we refer to it as two-step lasso. However, such a two-step strategy suffers from the model selection error in the first step, such as false positives of lasso (Bogdan et al., 2015; Su et al., 2017), which is disruptive for estimating structured simplexes in light of the strong structure implied by partial constancy. Moreover, there is no guarantee that the selected forecasters always have non-negative coefficients, leading to ambiguity regarding whether coefficients of opposite signs should be reset to the same constant. Indeed, our simulations show that two-step lasso may yield a non-negligible estimation error even in the most ideal case when the parameter strictly follows the structure, and its performance is highly sensitive to its regularization parameter.

Using Dirichlet priors in the Bayesian paradigm enables a unified approach to learn simplexes with uncertainty quantification. Along the line of weighting objects, Yang and Dunson (2014) considered convex aggregation and assigned symmetric Dirichlet priors to aggregation weights, and Rousseau and Mengersen (2011) studied overfitted mixture models where they advocated using Dirichlet distributions with small concentration parameters as priors for weights of densities. There is also a rich literature on generalizing Dirichlet distributions via mixtures to incorporate richer dependence patterns; see, for example, Ongaro and Migliorati (2013); Migliorati et al. (2017); Ongaro et al. (2020). In a recent work, Heiner et al. (2019) proposed a sparsified Dirichlet mixture prior to retain sparsity in probability simplex vectors. None of the Dirichlet priors in these works account for partial constancy, a key characteristic of our weight parameter space that arises from the aforementioned motivation.

In this article, we propose a Bayesian method to simultaneously accommodate sparsity, partial constancy, and the simplex constraint under the high-dimensional setting, while allowing for inference. This extends the useful structure embedded in the ‘forecast combination puzzle’ to a broader class of applications. At the core of our method lies a new class of *double spike Dirichlet* priors, which leads to joint inference without resorting to two-step alternatives that have pitfalls.

When applied to ensemble learning, the proposed double spike Dirichlet priors lead to a principled Bayesian strategy for structured ensembles, which enables uncertainty quantification and tuning key parameters without relying on cross-validation. We establish a

posterior contraction property by showing the entire posterior distribution concentrates around the desired structure, adaptive to the unknown sparsity level, which to the best of our knowledge is the first of its kind for structured high-dimensional probability simplexes. The proposed method is computationally easy to implement using Metropolis-Hastings algorithms. We illustrate through simulations and two case studies that the proposed method leads to substantial performance gain over alternatives. In particular, the proposed method complements random forecasts with improved prediction, enhanced parsimony, and easy implementation, strongly suggesting its wide applicability building on the success and popularity of random forests in practice. We provide R code at <https://github.com/xylimeng/StructuredEnsemble> for routine implementation.

The rest of the paper is organized as follows. Section 2 presents the proposed double spike Dirichlet priors and compares their properties to that of the symmetric Dirichlet prior. The posterior sampling algorithm is described in Section 3. Section 4 studies properties of the coefficients' posterior distribution and gives a posterior contraction rate. Simulation studies are carried out in Section 5. In Section 6 and Section 7, the method is illustrated by real data applications using the European Central Bank Survey of Professional Forecasters data set and the concrete data set (Yeh, 1998) from the UCI repository (Dua and Graff, 2017), respectively.

## 2. The Double Spike Dirichlet Priors

We consider a high-dimensional ensemble learning problem with model weights constrained on a simplex. Let  $\Theta^{K-1} = \{\beta \mid \beta \in \mathbb{R}^K, \sum_{i=1}^K \beta_i = 1, \beta_i \geq 0, 1 \leq i \leq K\}$  be a probability simplex. Ensemble learning can be cast as a constrained linear regression model

$$Y = X\beta + \varepsilon,$$

where the unknown  $K$ -dimensional weight parameter  $\beta$  is constrained on the probability simplex  $\Theta^{K-1}$ ,  $X$  is the given  $n \times K$  design matrix with each column being the prediction by a forecaster or learner, and  $\varepsilon \sim N(0, \sigma^2 I)$  is the error term with  $\sigma^2$  possibly unknown. We allow the situation  $n \leq K$  and are interested in structured coefficients, i.e., we assume the ground truth  $\beta^*$  belongs to the set

$$\Theta(s, K) = \{\beta \mid \beta \in \Theta^{K-1} \text{ has only } s \geq 1 \text{ nonzero elements} \\ \text{and they all take a common value } 1/s\}, \tag{1}$$

where the number and locations of nonzero elements are both unknown, and  $s$  is considerably smaller than  $K$ . We note that the methods developed in this article can generalize to any model or risk functions in lieu of the Gaussian likelihood, such as generalized linear regression and Gibbs posterior (Jiang and Tanner, 2008), leading to structured counterparts.

Motivated by the special structure in Equation (1) as well as the simplex constraint, we propose the following hierarchical class of double spike Dirichlet priors on  $\beta$ :

$$\beta \mid \gamma, \rho_1, \rho_2 \sim \text{Dir}(\rho_1 \gamma + \rho_2(1 - \gamma)) \\ \gamma_i \mid \theta \overset{\text{independent}}{\sim} \text{Bernoulli}(\theta), \quad i = 1, \dots, K, \tag{2}$$

where  $\rho_1 > \rho_2 > 0$ . We use a conjugate inverse-Gamma prior on  $\sigma^2$  when it is unknown, i.e.,  $\sigma^{-2}|a_1, a_2 \sim \text{Gamma}(a_1, a_2)$ . In Equation (2),  $\gamma_i = 1$  indicates  $\beta_i$  is associated with the larger concentration parameter  $\rho_1$ , and otherwise a smaller concentration parameter  $\rho_2$ . Let  $\gamma = (\gamma_1, \dots, \gamma_K)$  and denote  $|\gamma| = \sum_{i=1}^K \gamma_i$ . Marginalizing out  $\gamma$ , the prior on  $\beta$  becomes

$$g(\beta; \theta, \rho_1, \rho_2) = \sum_{\gamma} \frac{\Gamma(\rho_1|\gamma| + \rho_2(K - |\gamma|))}{\Gamma(\rho_1)^{|\gamma|}\Gamma(\rho_2)^{K-|\gamma|}} \prod_{i=1}^K \beta_i^{\rho_1\gamma_i + \rho_2(1-\gamma_i) - 1} \theta^{|\gamma|} (1 - \theta)^{K-|\gamma|}.$$

The proposed prior is a mixture of Dirichlet distributions, where the concentration parameters are independent Bernoullis supported on  $\{\rho_1, \rho_2\}$ . Through this article, we shall choose a large value for  $\rho_1 \gg 1$  and a small value for  $\rho_2 \ll 1$  to encode the structural information formulated in Equation (1), although the posterior sampling method in Section 3 is applicable for any  $(\rho_1, \rho_2)$ . With such choices, conditional on  $\gamma$ , elementary properties of Dirichlet distributions indicate that those  $\beta_i$  associated with  $\gamma_i = 0$  tend to be near 0 and the remaining ones are nonzero and tend to be near  $1/|\gamma|$ , leading to two spikes around which the values of  $\beta$  concentrate. Hence, we call  $g(\beta; \theta, \rho_1, \rho_2)$  the double spike Dirichlet priors. In the hierarchical representation of  $g(\beta; \theta, \rho_1, \rho_2)$  in (2), the use of Dirichlet distributions enables inference on parameters constrained to  $\Theta(s, K)$  in a unified manner. In contrast, two-step alternatives discussed in the preceding section can be loosely interpreted as learning  $\gamma$  via variable selection methods, followed by post-hoc normalization.

Employing the well-known relationship between Dirichlet and Gamma distributions, an equivalent form of the double spike Dirichlet priors is

$$\beta = \left( \sum_{i=1}^p u_i \right)^{-1} u, \quad u = (u_1, \dots, u_p), \quad u_i | \gamma_i \stackrel{\text{independent}}{\sim} \text{Gamma}(\gamma_i \rho_1 + (1 - \gamma_i) \rho_2, 1), \quad (3)$$

$$\gamma_i | \theta \stackrel{\text{independent}}{\sim} \text{Bernoulli}(\theta), \quad i = 1, \dots, K.$$

Hence, the double spike Dirichlet prior dichotomizes the shape parameter in the Gamma distributions. This representation facilitates posterior sampling as shown in the next section.

It is well known that the concentration parameters in the Dirichlet distribution are instrumental in determining prior properties, the choice of which has invited an extensive discussion. When using Dirichlet priors to learn multinomial distributions, for example, Good (1967) proposed a hierarchical approach by assigning priors to these parameters, and Fienberg and Holland (1973) discussed a variety of choices for the sum of concentration parameters. For weighting objects using Dirichlet priors, the symmetric Dirichlet distribution  $(\pi_1, \dots, \pi_K) \sim \text{Dir}(\alpha, \dots, \alpha)$ , which can be seen as a special case of our prior with  $\rho_1 = \rho_2$ , is still a default prior on the probability simplex. Often a very small  $\alpha$  value is chosen to induce sparsity. However, as indicated by the following proposition, as  $\alpha \rightarrow 0$ , most realizations sampled from the symmetric  $\text{Dir}(\alpha, \dots, \alpha)$  are very close to vectors containing all zeros with a single one at a random location from  $1, \dots, K$ .

**Proposition 1** *Let  $(\pi_{(1)}, \dots, \pi_{(K)})$  be the order statistics of  $(\pi_1, \dots, \pi_K)$ . If  $(\pi_1, \dots, \pi_K) \sim \text{Dir}(\alpha, \dots, \alpha)$ , then for any  $t \in (0, 1)$  and  $\alpha > 0$ , we have*

$$\text{pr}(\pi_{(K-1)} \leq t \pi_{(K)}) \geq t^{\alpha(K-1)}.$$

**Proof** See the Appendix. ■

Proposition 1 indicates that for any  $t \in (0, 1)$  and  $\alpha > 0$ , we have

$$\text{pr}(\pi_{(K-1)} \leq t) \geq \text{pr}(\pi_{(K-1)} \leq t\pi_{(K)}) \geq t^{\alpha(K-1)}.$$

Therefore, if we let  $\alpha \downarrow 0$ , the probability  $\text{pr}(\pi_{(K-1)} \leq t) \rightarrow 1$ , which implies that the second largest probability  $\pi_{(K-1)}$  is bounded from above by any fixed constant  $t$  with probability approaching 1. Furthermore, for small  $\alpha \in (0, 1)$ , if we let  $t = \alpha$ , we obtain that

$$\text{pr}(\pi_{(K-1)} \leq \alpha) \geq \alpha^{\alpha(K-1)}.$$

If  $\alpha = o(1)$ , it follows that

$$\text{pr}(\pi_{(K-1)} \geq \alpha) \leq 1 - \alpha^{\alpha(K-1)} = 1 - e^{\alpha(K-1)\log \alpha} = O(-\alpha \log \alpha) = o(\alpha^c),$$

for any positive constant  $c < 1$ . In contrast, the proposed double spike Dirichlet prior puts approximately at least  $\binom{K}{s}\theta^s(1-\theta)^{K-s}$  mass on probability vectors with any number  $s$  of nonzeros for small  $\rho_2$  and large  $\rho_1$ . This is formally described by the following proposition.

**Proposition 2** *For any  $s \in \{1, \dots, K\}$ , let  $N$  be the number of  $\beta_i$  that are greater than  $1/(2s)$  under the double spike Dirichlet prior. Suppose  $\rho_2 \leq 1/K$  and  $\rho_1 \geq 5/s$ . Then we have*

$$\text{pr}(N \geq s) \geq \binom{K}{s}\theta^s(1-\theta)^{K-s}(1-9s/\rho_1). \quad (4)$$

**Proof** See the Appendix. ■

The tail probability in Equation (4) is close to  $\binom{K}{s}\theta^s(1-\theta)^{K-s}$  for sufficiently large  $\rho_1$ , which is the probability of the event  $\{|\gamma| = s\}$ . The configuration of samples from the proposed prior depends on the choice of  $\theta$ ,  $\rho_1$  and  $\rho_2$ , where a large  $\rho_1$  together with a small  $\rho_2$  helps shrink the nonzeros to a constant. With  $\rho_1 \rightarrow \infty$  and  $\rho_2 \rightarrow 0$ , the proposed prior approaches  $\beta_i = \gamma_i / \sum_{i=1}^K \gamma_i$  conditional on  $\gamma$ , corresponding to the exact structure in  $\Theta(s, K)$ . For other choices of  $(\rho_1, \rho_2)$ , the proposed prior is a continuous relaxation. This flexibility is motivated by the consideration that many weights might be small (but not exactly zero) while the others are close to each other (but not exactly equal to each other), similar to the rich but different literature on continuous shrinkage priors in sparse high-dimensional regression. The developed sampling algorithm and theory in this article are applicable for both the limiting case and continuous relaxation.

### 3. Posterior Sampling and Posterior Summary

We adapt the stochastic search algorithm that used add, delete and swap moves for posterior sampling (Brown et al., 1998). We add a stay step so that the algorithm could keep drawing samples of  $\beta$  to improve mixing if stabilizing around a vector of  $\gamma_i$ 's. As such, we propose a random search algorithm for posterior sampling, which we call Add/Delete/Swap/Stay (ADSS). In the literature, methods based on this random walk type of Markov chain Monte Carlo (MCMC) are commonly used for posterior sampling of multivariate or high

dimensional binary vectors (Chapple et al., 2017). The algorithm is essentially Metropolis-Hastings; it exploits the well-known relationship between the Dirichlet and Gamma distribution as described in (3).

The ADSS algorithm includes the following steps:

1. Set  $t=0$ . Initialize  $(\sigma^{-2})^{(t)}$ ,  $\gamma_i^{(t)}$ ,  $A_i^{(t)}$  for  $i = 1, \dots, K$ . Set  $\beta^{(t)} = A^{(t)} / \|A^{(t)}\|_1$ , where  $A^{(t)} = (A_1^{(t)}, \dots, A_K^{(t)})$ , and  $\|\cdot\|_1$  is the  $\ell_1$  norm of a vector.
2. Set  $t = t + 1$ . Given  $\gamma^{(t-1)}$ , initialize a candidate vector  $\tilde{\gamma} = \gamma^{(t-1)}$ . Proceed to one of the following with equal probability:
  - (a) (add) randomly select a  $j$  from  $J = \{j' | \gamma_{j'}^{(t-1)} = 0\}$ . Set  $\tilde{\gamma}_j = 1$ ;
  - (b) (delete) randomly select a  $j$  from  $J^c = \{j' | \gamma_{j'}^{(t-1)} = 1\}$ . Set  $\tilde{\gamma}_j = 0$ ;
  - (c) (swap) randomly select a  $j_1$  from  $J$  and  $j_2$  from  $J^c$ . Set  $\tilde{\gamma}_{j_1} = 1$ ,  $\tilde{\gamma}_{j_2} = 0$ ;
  - (d) (stay) no actions;

Conditional on  $\tilde{\gamma}$ , for  $i = 1, \dots, K$ , propose a candidate  $\tilde{A}_i \sim \text{Gamma}(\rho_1 \tilde{\gamma}_i + \rho_2(1 - \tilde{\gamma}_i), 1)$ . Set  $\tilde{\beta} = \tilde{A} / \|\tilde{A}\|_1$ , where  $\tilde{A} = (\tilde{A}_1, \dots, \tilde{A}_K)$ .

3. Accept  $\gamma^{(t)} = \tilde{\gamma}$  and  $\beta^{(t)} = \tilde{\beta}$  with probability

$$\min \left( 1, \frac{\{\theta/(1-\theta)\}^{|\tilde{\gamma}|} \exp\{-(\sigma^{-2})^{(t-1)} \sum_{i=1}^n (y_i - x_i^T \tilde{\beta})^2 / 2\}}{\{\theta/(1-\theta)\}^{|\gamma^{(t-1)}|} \exp\{-(\sigma^{-2})^{(t-1)} \sum_{i=1}^n (y_i - x_i^T \beta^{(t-1)})^2 / 2\}} \right);$$

Otherwise, keep  $\gamma^{(t)} = \gamma^{(t-1)}$ ,  $\beta^{(t)} = \beta^{(t-1)}$ .

4. Draw  $(\sigma^{-2})^{(t)} \sim \text{Gamma} \left( a_1 + n/2, a_2 + \sum_{i=1}^n (y_i - x_i^T \beta^{(t)})^2 / 2 \right)$ .
5. Repeat Steps 2–4 for  $niter$  times.

Remark. For the limiting case when  $\rho_1 \rightarrow \infty$  and  $\rho_2 \rightarrow 0$ , which encodes the exact structure  $\beta_i = \gamma_i/|\gamma|$  in the prior, the ADSS algorithm proceeds by setting  $A_i^{(t)} = \gamma_i^{(t)}$  in Step 1 and  $\tilde{A}_i = \tilde{\gamma}_i$  in Step 2.

A Gibbs sampler is an alternative sampling strategy to ADSS, but it faces increasing challenges partly due to the simplex constraint. In a conventional Gibbs sampler, to update a single  $\gamma_i$ , one may draw a sample from the Bernoulli distribution with success probability  $\frac{\theta A_i^{\rho_1} / \Gamma(\rho_1)}{\theta A_i^{\rho_1} / \Gamma(\rho_1) + (1-\theta) A_i^{\rho_2} / \Gamma(\rho_2)}$ , which is usually close to 0 for all  $i$ . To improve mixing, a more appealing implementation of the Gibbs sampler would integrate out these  $A_i$  and draw directly from  $\gamma_i | Y$ . However, this is challenging when  $\beta$  is constrained on the simplex, unlike in the usual linear regression models where  $\beta$  is unconstrained and the integration can be analytically tractable. Therefore, we do not pursue Gibbs sampling and instead use a proposal distribution and Metropolis-Hastings for posterior sampling.

The ADSS algorithm is amenable to a fully Bayesian approach for  $\theta$  by giving it a prior and sampling it from the posterior distribution. In particular, by specifying a Beta( $b_1, b_2$ ) prior for  $\theta$  and exploiting conjugacy, the Gibbs update formula for  $\theta$  at step  $t$  is to draw

$\theta^{(t)} \sim \text{Beta}(b_1 + |\gamma^{(t)}|, b_2 + K - |\gamma^{(t)}|)$ . For hyperparameters, one may choose  $b_1 = 1$  and  $b_2 = 0.5K$  as suggested by Rockova and George (2018).

The posterior summary depends on the inferential goal. For example, if a point estimate of  $\beta$  is of interest, we may report the posterior mean after the burn-in period. If the uncertainty of the estimated  $\beta$  is also of interest, we may report a region of  $\beta$  based on quantiles of  $\ell_1$  errors. For the task of prediction such as in our real data applications, we mainly use the posterior mean of  $\beta$  to combine forecasts, unless stated otherwise.

#### 4. Posterior Properties

In this section, we study posterior properties of  $\beta$ , which guarantees the entire posterior distribution concentrates around the true structure and leads to insights about choosing  $\rho_1, \rho_2$  and  $\theta$ . For notational simplicity, we assume  $\sigma = 1$ . We first introduce a compatibility condition on the design matrix. Define the compatibility number of the model by

$$\Phi(s) = \inf \left\{ \frac{\|X\beta\|_2 \sqrt{s}}{\|X\| \|\beta\|_1} : \|\beta\|_1 \leq 2 \right\},$$

where  $\|X\| = \max_j \|X_{\cdot,j}\|_2$ , and  $\|\beta\|_q = (\sum_j |\beta_j|^q)^{1/q}$  for  $1 \leq q \leq \infty$ . We say the model satisfies the compatibility condition if  $\Phi(s) > 0$ . The compatibility condition ensures that the parameter  $\beta$  is estimable; in fact, the model may not be identifiable when  $K > n$  if we do not add constraints to the design matrix. Similar conditions are used in the high-dimensional regression literature when  $\beta$  is not constrained on the simplex; for example, see Bühlmann and Van De Geer (2011); Castillo et al. (2015).

The following theorem asserts that under certain conditions, the posterior distribution concentrates around the ground truth  $\beta^*$  at a certain rate.

**Theorem 3** *Suppose  $0 < s = o(K)$ ,  $\rho_1 = K^{\alpha_1}$ ,  $\rho_2 = K^{-\alpha_2}$ ,  $\alpha_1, \alpha_2 > 0$ ,  $\alpha_1/2 + \alpha_2 \geq 1$ . Under compatibility condition  $\Phi(s) > 0$ , if  $\theta$  is chosen such that  $O(1)/K \leq \theta \leq s/K$ , then for sufficiently large  $M > 0$ , we have*

$$\sup_{\beta^* \in \Theta(s,K)} E_{\beta^*} \Pi \left( \|\beta - \beta^*\|_1 > \frac{M}{\Phi(s)} \frac{s \log K}{\min(\|X\|, \sqrt{K^{\alpha_1}})} | Y \right) \rightarrow 0.$$

**Proof** See the Appendix. ■

Remark 1. Castillo et al. (2015) studied full Bayesian procedures for high-dimensional linear regression under sparsity constraints. Under their assumptions, they showed that

$$\sup_{\beta^*} E_{\beta^*} \Pi \left( \|\beta - \beta^*\|_1 > \frac{M}{\psi(s)^2} \frac{s \sqrt{\log K}}{\|X\| \phi(s)^2} | Y \right) \rightarrow 0,$$

where  $\psi(s)$  and  $\phi(s)$  are related to their compatibility conditions. Although our rate has a similar form, the proof is radically different due to the structured simplex constraint under consideration.

Remark 2. Theorem 3 reveals that both  $\|X\|$  and  $K^{\alpha_1}$  affect the posterior contraction rate, which depicts the effect of the likelihood and prior, respectively. Indeed, it can be



deduced from the proof that a larger  $\|X\|$  is linked to a stronger likelihood, while a larger  $K^{\alpha_1}$  leads to a tighter thus stronger prior. In particular, when the likelihood is strong enough to identify the nonzero coefficients, a stronger prior helps enforce the nonzero coefficients to be equal. The parameter  $\alpha_2$  does not affect the contraction rate in Theorem 3 as long as  $\alpha_2 \geq 1 - \alpha_1/2$ .

Remark 3. To achieve the rate in the theorem, the choice of  $\theta$  can be very flexible: instead of requiring  $\theta = s/K$ , letting  $\theta = t/K$  for any  $1 \leq t \leq s$  leads to the established rate. That is, the proposed method is *adaptive* to the unknown sparsity level.

## 5. Simulations

We carry out simulations to investigate the finite-sample performance of the proposed Bayesian method. Section 5.1 focuses on two scenarios, one following the exact partially constant structure and the other with deviation, and assesses estimation accuracy and sensitivity to hyperparameters for both the proposed method and competing methods. Using extended comparisons, Section 5.2 studies the effect of varying levels of deviation from the partially constant structure, different variants of the methods in comparison, and the effect of signal-to-noise ratio.

To implement our method, we use the ADSS algorithm to draw posterior samples. We use the following default settings in our numerical experiments, unless otherwise stated. We choose  $\rho_1 = K^2$  and  $\rho_2 = 1/K$ , following the theoretical results in Section 4 and particularly the condition  $\alpha_1/2 + \alpha_2 \geq 1$ . We put a weak Gamma( $a_1, a_2$ ) prior on  $\sigma^{-2}$  with  $a_1 = a_2 = 0.01$ . For posterior sampling for the proposed method and other Bayesian methods, we run *niter* = 20000 iterations and burn in the first 15000. For the choice of  $\theta$ , we set it to  $1/K$ , which is supported by Theorem 3, but we note that one can also specify other preferred values or a range of values based on their domain knowledge about the sparsity level. Other than the default choice of fixing  $\theta$ , we will also compare a fully Bayesian treatment by placing a Beta prior on  $\theta \sim \text{Beta}(1, 0.5K)$ . In our numerical experiments, we find that the ADSS algorithm is numerically stable without triggering rounding or overflow errors.

### 5.1 Two Scenarios

In the first scenario, the true simplex parameter strictly follows the sparse and partially constant structure, while the second scenario allows some deviation.

We consider two alternative methods for comparison: a Bayesian method using the symmetric Dirichlet prior and the two-step lasso method proposed by Diebold and Shin (2019). For the symmetric Dirichlet method, we use the algorithm in Yang and Dunson (2014). Tuning parameters in these two methods include the concentration parameter  $\rho$  for symmetric Dirichlet, and the regularization parameter  $\lambda$  for two-step lasso. We give them an advantage throughout the numerical experiments by searching a grid of hyperparameters and reporting the best result, as opposed to fixing the hyperparameter values when implementing our method. Details of the grid search are deferred to the sensitivity analysis later in this section.

We use the  $\ell_1$  error of the estimated  $\beta$  as the metric to evaluate the three methods. For the two Bayesian methods, we calculate the  $\ell_1$  error for each of the 5000 MCMC samples,

and report the averaged  $\ell_1$  error. This metric directly points to posterior contraction that is unique to Bayesian methods and has been reflected in the rate calculation in Theorem 3, and utilizes the entire posterior samples without relying on point estimates, eliminating the dependence on posterior summary for which multiple strategies exist. For two-step lasso, we use its point estimate to calculate the  $\ell_1$  error as uncertainty quantification for this non-Bayesian method is not immediately available. According to Jensen’s inequality and the convexity of the  $\ell_1$  norm, the reported  $\ell_1$  error for the two Bayesian methods would be no smaller than the one calculated using the posterior mean; hence, our comparison gives the two-step alternative another advantage as the  $\ell_1$  error of the two Bayesian methods could be further improved if a point Bayes estimator is used instead. Each scenario is repeated 100 times, and the averaged  $\ell_1$  error is compared.

For the first scenario, we generate 80 samples from  $y = x^T \beta^* + \varepsilon$ , where  $x \in \mathbb{R}^{40}$ ,  $\varepsilon \sim N(0, 1.5^2)$ , and  $\beta^* = (1/3, 1/3, 1/3, 0, \dots, 0)$ , i.e.,  $n = 80$ ,  $K = 40$ ,  $s = 3$  in this case. Each column of the design matrix  $X$  is generated independently from  $N(0, 3^2)$ .

The proposed method yields an average  $\ell_1$  error 0.032. The best  $\ell_1$  error over all hyperparameter sets is 0.428 and 0.139 for the symmetric Dirichlet prior and two-step lasso, respectively. This suggests that under this simulation setting the proposed estimates are more accurate than alternatives by a substantial margin; recall that the alternatives have been equipped with the optimal hyperparameter value that achieves the smallest  $\ell_1$  error. The performance gain of our method over symmetric Dirichlet is not surprising as the ground truth does follow the sparse and partially constant pattern as encoded in the double spike Dirichlet prior.

The second scenario mimics the first scenario but allows derivation of  $\beta^*$  from the sparse and partially constant pattern. In particular, the last  $K - s$  entries of  $\beta^*$  equally share a total of 0.05 weight, and the first  $s$  entries are  $(\beta_1, \beta_2, \beta_3) = (0.3089, 0.3672, 0.2739)$ . Its closest projection to the partially constant simplex parameter space is the  $\beta^*$  used in the first scenario, and their  $\ell_1$  distance is 0.1677.

The proposed method yields an average  $\ell_1$  error 0.210. The best  $\ell_1$  error achieved is 0.436 for the symmetric Dirichlet prior and 0.298 for two-step lasso. Not surprisingly, both our method and two-step lasso have a larger  $\ell_1$  error than the first scenario due to the deviation of  $\beta^*$  from the exact structure. However, the proposed method appears to be more flexible and can better adapt to such cases than two-step lasso. On the other hand, the symmetric Dirichlet method continues to show a larger error than our method, perhaps due to the lack of a strong structural prior when the sample size is relatively small compared to the number of predictor variables.

We now turn to the mixing of  $\gamma$  in the proposed method. Sampling of high-dimensional discrete indicators could be challenging if the target distribution is highly multimodal (Schäfer, 2012). Interestingly, the strong structural information that our double spike prior encodes appears to greatly mitigate such concerns. When the underlying  $\beta^*$  follows (or approximately follows) the sparse and partially constant structure, the proposed prior tends to yield a posterior distribution that is less likely to be multi-modal. This is because unlike high-dimensional unconstrained parameter space, a small change in  $\gamma$  could mean a large difference in  $\beta$  and subsequently in the likelihood, when  $\beta$  is constrained to the sparse, partially constant simplex. For example, the  $\ell_1$  distance between any two elements in  $\Theta(s, K)$  as well as that between  $\Theta(s, K)$  and  $\Theta(s - 1, K)$  is lower bounded by

$2/s$ . To see this, we construct elements such that the lower bound is achieved. In particular, we consider  $(1/s, \dots, 1/s, 0, \dots, 0) \in \Theta(s, K)$  versus any other element in  $\Theta(s, K)$  with one of its last  $K - s$  element being  $1/s$ , and  $(1/s, \dots, 1/s, 0, \dots, 0) \in \Theta(s, K)$  versus  $(1/(s - 1), \dots, 1/(s - 1), 0, \dots, 0) \in \Theta(s - 1, K)$ . However, these distances could be arbitrarily small if the parameter  $\beta$  is unconstrained as in traditional high-dimensional linear regression settings.

To demonstrate the aforementioned tendency toward uni-modality using our numerical experiments, we take the first scenario as an example and assess the relative frequency of two different  $\gamma$ 's by computing an approximated posterior probability ratio  $p(\gamma|Y)/p(\tilde{\gamma}|Y)$  for  $\gamma \neq \tilde{\gamma}$ . It is equivalent to use the unnormalized  $p(\gamma|Y)$ , which, after integrating out  $\sigma^2$ , is approximately  $p(\gamma|Y) \propto \{a_2 + \|Y - X \frac{\gamma}{|\gamma|}\|_2^2/2\}^{-(a_1+n/2)} \theta^{|\gamma|} (1 - \theta)^{K-|\gamma|}$ . This approximation is obtained by considering the limiting case when  $\rho_1 \rightarrow \infty$  and  $\rho_2 \rightarrow 0$ . We look into one generated data set and find that there are 10 unique  $\gamma$  vectors in the posterior samples; after burn-in, there is only one unique  $\gamma$ , which is  $\gamma^* = (1, 1, 1, 0, 0, \dots, 0)$ . For any  $\gamma \neq \gamma^*$  in the 10 unique vectors, the highest posterior probability ratio  $p(\gamma|Y)/p(\gamma^*|Y) = 3.49 \times 10^{-7}$ , which is very small. This provides strong evidence against multi-modality that hurdles mixing of an MCMC sampler. In Section 7, we further assess the convergence of our MCMC samples using real data when the exact structure is not guaranteed to hold.

We next provide sensitivity analyses for the three methods with respect to hyperparameters. For the proposed method, we compare results for  $\rho_1 = \{6.3, 23.6, 88.2, 329.2, 1229.4, 4590.4, 17140.2, 64000\}$  and  $\rho_2 = \{1.6 \times 10^{-5}, 6.8 \times 10^{-5}, 3 \times 10^{-4}, 1.3 \times 10^{-3}, 5.7 \times 10^{-3}, 2.5 \times 10^{-2}\}$ , which are generated by taking  $K$  to the power of an equally-spaced grid on  $[0.5, 3]$  and  $[-3, -1]$ , respectively. For the symmetric Dirichlet with the concentration parameter  $\rho$ , we vary  $\rho$  from the set for  $\rho_2$ . For the regularization parameter  $\lambda$  in two-step lasso, we use a grid of 80 values equally spaced on  $[-8, 8]$  for  $\log \lambda$ , and this gives a list of  $\lambda$  ranging from 0.0003 to 2981.

Figure 1 shows the sensitivity to hyperparameters for each method. The upper panel is for Scenario 1. From the plot, the proposed double spike Dirichlet prior yields uniformly better estimates than symmetric Dirichlet in terms of the  $\ell_1$  error, which is expected since the double spike Dirichlet prior utilizes the structural information of the coefficients. Performances of two-step lasso largely depend on the choice of  $\lambda$ ; in contrast, the proposed method is much less sensitive to the choice of  $\rho_1$  and  $\rho_2$ . We also observe a monotonic pattern in the upper right heatmap of Figure 1: the  $\ell_1$  error tends to decrease with larger  $\rho_1$  values and smaller  $\rho_2$  values. Indeed, large  $\rho_1$  and small  $\rho_2$  values are desired in this scenario as the true  $\beta^*$  does follow the sparse and partially constant structure. A limiting case is  $\rho_1 \rightarrow \infty$  and  $\rho_2 \rightarrow 0$ , which amounts to  $\beta_i = \gamma_i/|\gamma|$ . Samples drawn from the double spike prior with such choices possess the exact structure, which is desirable when the exact structure holds. However, the numerical difference in  $\ell_1$  error becomes indistinguishable for large enough  $\rho_1$ . The lower panel is for Scenario 2. The boxplot continues to show that the proposed method outperforms the others, although the sparse and partially constant assumption does not strictly hold. However, the monotone decreasing pattern in  $\ell_1$  error as  $\rho_1$  gets larger and  $\rho_2$  gets smaller no longer holds (the lower right plot), which is not surprising as the true  $\beta^*$  deviates from the exact structure. In both scenarios, the proposed method is not sensitive to  $\rho_1$  and  $\rho_2$  when  $\rho_1$  exceeds certain thresholds.

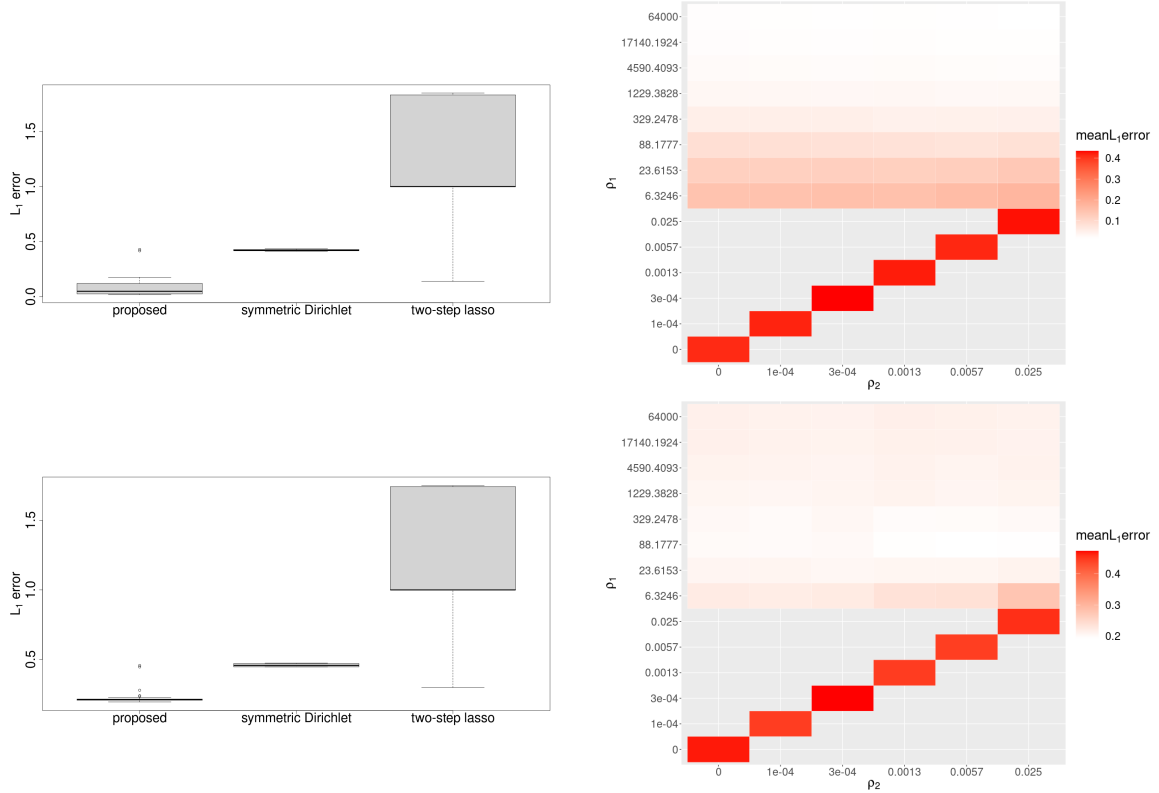


Figure 1: Comparison of the  $\ell_1$  error of coefficient estimates of three methods averaged over 100 simulations for Scenario 1 (the first row) and Scenario 2 (the second row). Left: boxplots of  $\ell_1$  error as hyperparameters vary (the hyperparameters are  $(\rho_1, \rho_2)$  for the proposed method,  $\rho$  for the symmetric Dirichlet prior, and  $\lambda$  for two-step lasso). Right:  $\ell_1$  error of the proposed method as  $(\rho_1, \rho_2)$  vary (the upper half) and the symmetric Dirichlet method as  $\rho = \rho_1 = \rho_2$  varies (anti-diagonal elements in the lower half).

## 5.2 Extended Comparisons

We extend the previous section by varying the deviation level of  $\beta^*$  from the strict sparse and partially constant structure. We also consider variants of methods in comparison. In particular, for the proposed method, instead of fixing  $\theta = 1/K$ , we put a prior  $\text{Beta}(1, 0.5K)$  on  $\theta$  and update it as well; for the two-step approach, instead of using lasso, one may use other adaptive shrinkage methods such as horseshoe (Carvalho et al., 2010), leading to the method of two-step horseshoe. We implement two-step horseshoe by using the default `HS.var.select()` function from the R package `horseshoe`, which selects a variable if its marginal credible interval does not contain 0.

We create 10 different  $\beta^*$  values corresponding to different deviations from partial constancy. Starting with the  $\beta^*$  vector in the second scenario, which uses  $\beta_1 = 0.3089$  and distributes a total of 0.05 weight to the last  $K - 3$  elements, we vary  $\beta_2$  from 0.2 to 0.3. The value of  $\beta_3$  is determined by the simplex constraint. We choose this range for  $\beta_2$  to ensure that their closest point on  $\{\Theta(s, K) : s \leq K\}$  is  $(1/3, 1/3, 1/3, 0, \dots, 0)$ , which can be checked empirically. The  $\ell_1$  distance between the generated  $\beta^*$  and its projection in  $\{\Theta(s, K) : s \leq K\}$  ranges from 0.1244 to 0.3155; this is a wide interval and covers the derivation in Scenario 2 in Section 5.1 (0.1677).

Table 1 shows the average  $\ell_1$  error of  $\beta$  by our methods and competing methods when the level of deviation varies. We observe a consistent pattern that the proposed methods outperform the symmetric Dirichlet and the two-step approaches. For our methods, the two versions with fixed  $\theta$  and random  $\theta$  do not show significant differences relative to standard errors, although the fixed  $\theta$  approach has a slightly smaller  $\ell_1$  error for all cases. For the two-step approaches, two-step horseshoe has smaller  $\ell_1$  error than two-step lasso, but it is still uniformly worse than the proposed methods. This adds to evidence of efficiency gain by adopting the proposed unified approach relative to two-step alternatives. Besides, in our first real data application, we observed that two-step horseshoe failed to select any predictors. We next compare symmetric Dirichlet with the fixed  $\theta$  version of our methods. Symmetric Dirichlet leads to uniformly larger  $\ell_1$  errors, with a widening margin as the deviation turns smaller. It is reassuring that the proposed method continues to outperform alternatives when the deviation is as large as 0.3155. A close inspection of the symmetric Dirichlet method in one simulated data set indicates that it assigns in total more than 0.15 weight to the near-zero components of  $\beta$ , much larger than the true weight of 0.05, while our method assigns 0.007. Therefore, our method appears to better maintain the sparsity than the symmetric Dirichlet. For a non-sparse  $\beta^*$ , which is not particularly interesting in the high-dimensional setting considered here, symmetric Dirichlet may become competitive.

In the appendix, we evaluate whether the signal-to-noise ratio would affect the comparisons among these methods. The results are consistent with what has been reported.

## 6. Application to Survey Forecast Data

We use the European Central Bank’s quarterly Survey of Professional Forecasters (SPF) data set to illustrate our method. The SPF collects information on the expected rates of inflation, real GDP growth, and unemployment in the euro area at several horizons, ranging from the current year to the longer term. Here we focus on making quarterly 1-year-ahead forecasts of euro-area real GDP growth (year-on-year percentage change) using the survey

Deviation	Proposed		Symmetric Dir	Two-step	
	$\theta = 1/K$	$\theta \sim \text{Beta}$		lasso	horseshoe
0.3155	0.420 (0.010)	0.432 (0.009)	0.591 (0.013)	0.464 (0.016)	0.453 (0.017)
0.2943	0.395 (0.011)	0.414 (0.011)	0.592 (0.015)	0.448 (0.020)	0.438 (0.018)
0.2731	0.367 (0.012)	0.387 (0.011)	0.597 (0.016)	0.436 (0.020)	0.420 (0.019)
0.2518	0.329 (0.012)	0.355 (0.012)	0.592 (0.016)	0.409 (0.021)	0.378 (0.020)
0.2306	0.294 (0.012)	0.318 (0.012)	0.581 (0.019)	0.384 (0.021)	0.353 (0.021)
0.2094	0.264 (0.012)	0.292 (0.012)	0.565 (0.020)	0.354 (0.022)	0.313 (0.018)
0.1881	0.237 (0.012)	0.261 (0.012)	0.555 (0.021)	0.328 (0.022)	0.276 (0.017)
0.1669	0.209 (0.012)	0.225 (0.012)	0.531 (0.020)	0.308 (0.022)	0.248 (0.017)
0.1457	0.179 (0.010)	0.195 (0.011)	0.525 (0.021)	0.295 (0.023)	0.231 (0.018)
0.1244	0.149 (0.008)	0.166 (0.009)	0.494 (0.019)	0.277 (0.025)	0.210 (0.018)

Table 1: Comparison of the  $\ell_1$  error of coefficient estimates for various levels of deviation from sparsity and partial constancy. Deviation of  $\beta^*$  is defined as the  $\ell_1$  distance between the generated  $\beta^*$  and its projection in  $\{\Theta(s, K) : s \leq K\}$ . Results are averaged over 100 simulations. Standard errors are reported in parentheses.

data from 1999Q1 to 2016Q2. This data set contains predictions from some forecasters of GDP growth rate for periods 1999Q3–2016Q4. The data preprocessing step, including a preliminary forecaster selection and missing data imputation procedure, follows Diebold and Shin (2019). In total, we have 70 surveys, and each quarterly survey contains forecasts from  $K = 23$  forecasters.

We calculate the out-of-sample root-mean-squared error to measure the performance of the proposed method, the Bayesian procedure using the noninformative symmetric Dirichlet prior, and two-step lasso. We also include the simple average method, which leads to robust and constantly favorable performances hinted by the forecast combination puzzle. For the proposed method, we take the linear combination of individual forecasts as the final prediction, where the weights are determined by the posterior mean of the coefficient. The evaluation period starts from 2000Q1, and the root-mean-squared error is calculated for the forecasts of 68 quarters.

Except for the simple average, all other methods require selection of hyperparameters. For our method, we fix  $\rho_1 = K^2$  and  $\rho_2 = 1/K$ ; we fix  $\theta = 0.05$ , which is roughly  $1/K$ . Similarly to the preceding section, we also implement a random- $\theta$  strategy by giving it the  $\text{Beta}(1, 0.5K)$  prior. For two-step lasso, at each period, we search for the optimal  $\lambda$  among the same grid of values used in Section 5, and choose the one that yields the smallest root-mean-squared error historically.

Starting from period two and rolling forward, we obtain the forecasts for  $t = 3, \dots, 70$  based on which the root-mean-squared error is calculated. For our method, the fixed- $\theta$  strategy yields an out-of-sample root-mean-squared error 1.441. The random- $\theta$  counterpart performs similarly, leading to an error of 1.444; for the proposed method, we focus on the fixed- $\theta$  strategy in the following, unless stated otherwise. The simple average gives 1.502, and two-step lasso gives 1.542. A closer look at the squared prediction errors reveals three outliers for all methods, which occurred in 2008Q2, 2008Q3, and 2008Q4 for the

GDP growth over 2007Q4-2008Q4, 2008Q1-2009Q1, and 2008Q2-2009Q2, respectively. This corresponds to the financial crisis in 2008. After removing the three outliers, the out-of-sample root-mean-squared error is 1.014 for our method, 1.032 for the simple average method, and 1.058 for two-step lasso. The two-step lasso method, which chooses weights adaptively, does not outperform the simple average method. The comparison between the proposed method and two-step lasso confirms the advantage of a joint approach versus a two-step alternative, at least in this data application.

The inferred weights in this data application are very sparse. Taking the prediction at 2011Q2 as an example (i.e., we fit the model using data up to 2011Q1), the posterior mean of  $|\gamma| = \sum_{i=1}^K \gamma_i$  is 1.12, and the posterior mean of  $\beta$  has four non-negligible components: 0.40, 0.26, 0.09, 0.09. This observation is consistent across quarters; for example, the posterior mean of  $|\gamma|$  when averaged over quarters is 1.13. We have also observed that the proposed method can differentiate highly correlated predictors. For example, using data up to 2011Q1 again, the predictor with the most prominent  $\beta$  (whose posterior mean is 0.40) turns out to be highly correlated with the 19 predictors with negligible coefficient estimates: the correlation ranges from 0.83 to 0.93. Therefore, in this data application we do not observe that the proposed method selects an arbitrary candidate from a set of highly correlated predictors. We additionally conduct a simple simulation study in Appendix C to assess the proposed method in the presence of high correlation among predictors, which shows supporting evidence for this interesting observation.

To assess the convergence for  $\gamma$ , we conduct Gelman and Rubin diagnostic (Gelman and Rubin, 1992) for each  $\gamma_i$ ,  $1 \leq i \leq K$ . This procedure runs two parallel Markov chains, each with a different initialization, and calculates a diagnostic statistic. For our method, the largest diagnostic statistic is  $0.9999 < 1.1$ , indicating that the MCMC is run long enough. Note that here we exclude the cases when the diagnostic statistic is reported as ‘NA’, which occurs when  $\gamma_i$ ’s empirical frequency after burn-in is close to the boundary of the parameter space (either 0 or 1). We further visualize the empirical frequency of different  $\gamma$  vectors between the two parallel MCMC chains. Similarly to the logic behind the Gelman and Rubin diagnostic, if the sampler converges and mixes well, then the two parallel chains with different initializations should have similar frequencies for the same  $\gamma$ . In Figure 2, we show the frequencies of the five most frequent  $\gamma$  vectors and group all other  $\gamma$ ’s into a combined category. The two frequency distributions for  $\gamma$ ’s are close, again suggesting good mixing and convergence.

## 7. Improving Random Forests Predictions

In this section, we use the concrete data set from the UCI repository to demonstrate how the proposed method improves random forests. Random forests have been considered one of the most popular supervised learning methods. The high dimension of trees and the simple average structure in random forests (and also bagging) make the proposed method particularly well suited for further improvement. This will lead to structured random forests aided by double spike Dirichlet priors. The induced adaptive weighting effect that takes into account learners’ performance has also been seen in other ensemble methods such as boosting (Freund and Schapire, 1997), but unlike boosting, which considers learners in a




Figure 2: Frequencies of different  $\gamma$  vectors in two parallel MCMC chains. The x-axis represents each  $\gamma$  by its nonzero components; for example, “{22}” means only the 22th element of the corresponding  $\gamma$  vector is nonzero.

sequential manner, our structured weighting jointly assesses all learners with additional shrinkage toward equal weights.

Concrete is a critical material in civil engineering. The concrete compressive strength is a function of age and ingredients. With the concrete data set (Yeh, 1998), we use eight input features, namely cement ( $\text{kg}/\text{m}^3$ ), fly ash ( $\text{kg}/\text{m}^3$ ), blast furnace slag ( $\text{kg}/\text{m}^3$ ), water ( $\text{kg}/\text{m}^3$ ), superplasticizer ( $\text{kg}/\text{m}^3$ ), coarse aggregate ( $\text{kg}/\text{m}^3$ ), fine aggregate ( $\text{kg}/\text{m}^3$ ) and age of testing (days) to predict the compressive strength of concrete. The data set contains 1030 samples. We randomly sample 515 samples as the training set and use the rest for testing. We first construct random forests and obtain individual tree predictions using the **randomForest** package (Liaw and Wiener, 2002) in R with the default settings for all arguments. This corresponds to  $K = 500$  trees. We then treat the 500 trees as 500 predictors and feed them into the proposed Bayesian procedure. Since our simulations suggest the proposed method is not sensitive to hyperparameters, we set  $\rho_1 = K^2$ ,  $\rho_2 = 1/K$ , and  $\theta = 0.2$ . Here we set  $\theta$  larger than  $1/K$  in order to leverage a larger number of trees for variance reduction, while still being more parsimonious than the original forests. We also implement the fully Bayesian version of our method by placing a  $\text{Beta}(1, 0.5K)$  prior on  $\theta$ . Since the number of predictors  $K$  is large, we increase the number of iterations to 30000 (from the default value of 20000) and discard the first 20000 as burn-in. We then calculate the posterior mean weight for each tree and use this weighted sum as our proposed prediction. The out-of-sample root-mean-squared error is calculated on the test set.

Figure 3 shows the out-of-sample root-mean-squared errors of the proposed method and the original random forests method, as well as their differences. The results are based on 100 replicates of the above procedure, each time generating a new train-test split. We can see that our method outperforms random forests in almost every train-test split, leading to a smaller out-of-sample root-mean-squared error. The average out-of-sample root-mean-






Figure 3: Comparison of the out-of-sample root-mean-squared error of the proposed method and the random forests based on 100 replications. Left: difference of prediction error of the proposed method subtracting that of random forests. Right: prediction error of each method.

squared error over 100 replications of our method is 5.870, reducing the original random forests’ 6.332 substantially by 7.30%. For the full Bayesian version of our method, the improvement becomes 5.37%, reducing the error from 6.332 to 5.992.

To gain insights on why the proposed method improves random forests, we extract one random forest consisting of 500 trees on one training-test set and calculate the number of trees that actually contribute to the final prediction, that is, the number of trees associated with  $\gamma_i = 1$ , at each iteration after the burn-in period. It turns out that in this example, the number of contributing trees ranges from 10 to 18 out of 500 trees. In particular, only 13 trees are selected with a frequency higher than 0.2, only 29 trees are selected with a frequency higher than 0.05, and the vast majority of 393 trees are never selected. Thus, the proposed method leads to extreme parsimony when applied to random forests in this data set.

We further term the 18 most frequently selected trees as the ‘selected group’ and the rest as the ‘unselected group’, and compare their performances to learn more insights. Taking the observed responses as the ground truth, for each tree, we calculate its averaged bias and variance over the samples. Figure 4 shows that the selected group generally has smaller variance than the unselected group, and larger bias appears more often in the unselected group than in the selected group. These observations hold for both the training and test sets. Therefore, the proposed method indeed selects a group of better predictors and generalizes well to the test set. Interestingly, the selected group does not necessarily consist of the best 18 individual trees, which are those trees yielding the smallest mean-squared errors on the training set; in fact, the simple average of them gives a worse out-of-sample root-mean-squared error than the proposed method. Looking further into the selected group and the




Figure 4: Comparisons of the selected group and unselected group: biases (left) and variances (right).

best 18 individual tree group suggests that this might be due to the larger correlation among members in the best 18 individual trees, causing them to bias toward the same direction and thus hampering the prediction. Indeed, we summarize tree correlations for each group by averaging all pairwise tree correlations in that group, where the pairwise tree correlation is defined as the correlation between the pair’s prediction biases, and find that the selected group has an average pairwise correlation of 0.16, while the best 18 individual trees group has 0.23. This case study indicates that the proposed method accounts for the dependence across individual trees, which appears to be crucial to improve the prediction.

## 8. Conclusions

In this paper, we propose a novel class of double spike Dirichlet priors for structured high-dimensional probability simplexes. In particular, we accommodate the structure of sparsity and partial constancy, abstracting and extending a useful concept that emerged as the so-called ‘forecast combination puzzle’ in economics to a broader class of applications. The proposed priors lead to a Bayesian procedure for estimating weight parameters. We design an efficient MCMC algorithm for posterior sampling, enabling inference and uncertainty quantification. We establish posterior concentration properties for the proposed priors under mild regularity conditions and proper choice of hyperparameters, which are adaptive to the unknown sparsity level. Simulations show that our method achieves a smaller estimation error and is more robust and flexible than its competitors in both ideal and less ideal scenarios. We resort to two applications to assess the performance of the proposed method using real world data. In the context of forecast combination, the proposed method exhibits superior predictive performance in an application to the European Central Bank Survey of Professional Forecasters data set; it further improves the prediction of random forecasts by more than 7% using a UCI data set, suggesting its wide applicability building on the success and popularity of random forests.

There are several interesting directions building on the proposed method. Firstly, our method can be extended to incorporate group structures. This can be achieved by placing

a hierarchical prior on the binary vector  $\gamma$  in (2):

$$\begin{aligned} \gamma_i &= \gamma_{g(i),1} \cdot \gamma_{i,2}, \\ \gamma_{g(i),1} &\stackrel{\text{independent}}{\sim} \text{Bernoulli}(\theta_G), \\ \gamma_{i,2} &\stackrel{\text{independent}}{\sim} \text{Bernoulli}(\theta), \end{aligned}$$

where  $g_i$  is the group that  $\beta_i$  falls into. Here  $\theta_G$  represents one’s prior belief on the group sparsity, and  $\theta$  represents the prior belief on the within-group sparsity when group  $g_i$  receives nonzero weights. In this formulation,  $\gamma_i \neq 0$  if and only if it receives nonzero weight ( $\gamma_{i,2} \neq 0$ ) within a nonzero group ( $\gamma_{g(i),1} \neq 0$ ). Such spike-and-slab priors on variable groups are commonly used to achieve group sparsity in other settings (Bai et al., 2022).

Secondly, the sampling of high-dimensional binary vectors using ADSS may have the potential pitfall of slow mixing (Schäfer, 2012), particularly when the number of predictors is large. This is mitigated by the strong structure of the considered parameter space as discussed in preceding sections, and one can additionally increase the number of MCMC iterations for larger  $K$  as in our example to improve random forests. In conventional sparse Bayesian variable selection, Yang et al. (2016) showed under mild conditions certain random walk type of Metropolis-Hastings algorithm can mix rapidly. While our posterior sampling algorithm seems to work well in practice, it is interesting to theoretically investigate whether or not it enjoys such rapid mixing in the presence of structured simplexes.

Finally, non-negative weights on the simplex that our priors are supported on are widely adopted in many applications, such as forecast combination in economics and model selection. For example, the induced convex combination is shown to be more stable than its non-constrained counterpart (Conflitti et al., 2015) and remains a popular choice in model aggregation (Bunea et al., 2007; Polley and Van der Laan, 2010). Nevertheless, an interesting future direction would be to relax the non-negative constraint of a simplex to mixed signs.

## Acknowledgments

We thank David Dunson and Marina Vannucci for their helpful comments on an earlier version of this paper, and Minchul Shin for sharing the preprocessed SPF data used in Section 6. We thank the action editor and three anonymous reviewers for constructive comments that helped to improve the paper. This work was partially supported by the grant DMS-2015569 and DMS/NIGMS-2153704 from the National Science Foundation, and a Ken Kennedy Institute 2020/21 Oil & Gas High Performance Computing Conference Graduate Fellowship.

## Appendix A.

This section contains the proofs of Proposition 1, Proposition 2 and Theorem 3.

**Proof** (Proof of Proposition 1) Let  $A_i$  be drawn independently from  $\text{Gamma}(\alpha, 1)$  and  $\pi_i = A_i / \sum_{i=1}^K A_i$ , then it is well known that  $(\pi_1, \dots, \pi_K) \sim \text{Dir}(\alpha, \dots, \alpha)$ . Note that  $\pi_{(K-1)}/\pi_{(K)} = A_{(K-1)}/A_{(K)}$ , and the event  $\{A_{(K-1)} \leq tA_{(K)}\}$  for given  $t \in (0, 1)$  can be

partitioned into  $K$  disjoint equal-probability events as follows:

$$\{A_{(K-1)} \leq tA_{(K)}\} = \bigcup_{k=1}^K \{A_i \leq tA_k, 1 \leq i \leq K \text{ and } i \neq k\}.$$

Therefore,

$$\begin{aligned} \text{pr}(\pi_{(K-1)} \leq t\pi_{(K)}) &= \text{pr}(A_{(K-1)} \leq tA_{(K)}) \\ &= K \text{pr}(A_1 \leq tA_K, A_2 \leq tA_K, \dots, A_{K-1} \leq tA_K) \\ &= KE \{ \text{pr}(A_1 \leq tA_K, A_2 \leq tA_K, \dots, A_{K-1} \leq tA_K | A_K) \} \\ &= KE \{ F^{K-1}(tA_K) \}, \end{aligned} \tag{5}$$

where the expectation in the last two steps is with respect to  $A_K \sim \text{Gamma}(\alpha, 1)$ , and  $F(\cdot)$  is the cumulative distribution function of  $\text{Gamma}(\alpha, 1)$ . When  $t \in (0, 1)$ , it follows that  $F(tx) \geq t^\alpha F(x)$  for any  $x > 0$ . This is because

$$\begin{aligned} F(tx) &= \int_0^{tx} \frac{1}{\Gamma(\alpha)} z^{\alpha-1} e^{-z} dz \\ &\stackrel{y=z/t}{=} t^\alpha \int_0^x \frac{1}{\Gamma(\alpha)} y^{\alpha-1} e^{-ty} dy \geq t^\alpha \int_0^x \frac{1}{\Gamma(\alpha)} y^{\alpha-1} e^{-y} dy \\ &= t^\alpha F(x). \end{aligned}$$

Combining this inequality with (5), we obtain that

$$\text{pr}(\pi_{(K-1)} \leq t\pi_{(K)}) \geq K t^{\alpha(K-1)} E \{ F^{K-1}(A_K) \}.$$

This completes the proof by noting that  $E [F^{K-1}(A_K)] = \int_0^\infty F^{K-1}(x) dF(x) = 1/K$ .  $\blacksquare$

**Proof** (Proof of Proposition 2) Under the proposed prior, a lower bound on the probability that  $N \geq s$  is the probability that  $\gamma$  has  $s$  nonzero components and  $\beta$  has  $s$  components greater than  $1/(2s)$ . Due to exchangeability, this is simply

$$\binom{K}{s} \theta^s (1-\theta)^{K-s} \cdot \text{pr}(\beta_i > 1/(2s), \text{ for all } i \leq s | \gamma_i = 1 \text{ if } i \leq s \text{ and } 0 \text{ otherwise}).$$

Given that  $\gamma_i = 1$  if  $i \leq s$  and 0 otherwise, for  $i \leq s$ ,  $E(\beta_i) = \rho_1 / (s\rho_1 + (K-s)\rho_2) \in (1/(s + \rho_1^{-1}), 1/s)$  if  $\rho_2 \leq 1/K$ , and  $\text{Var}(\beta_i) < (s^2\rho_1)^{-1}$ . Under the condition  $\rho_1 \geq 5/s$ , we have  $1/(s + \rho_1^{-1}) \geq 5/(6s)$ . By union bound and Chebyshev inequality,

$$\begin{aligned} \text{pr}(\beta_i < 1/(2s), \text{ for some } i \leq s) &\leq s \cdot \text{pr}(\beta_i - E(\beta_i) < 1/(2s) - 1/(s + \rho_1^{-1})) \\ &\leq s \cdot \text{pr}(|\beta_i - E(\beta_i)| > 1/(s + \rho_1^{-1}) - 1/(2s)) \\ &\leq 9s/\rho_1. \end{aligned}$$

Therefore,

$$\begin{aligned} \text{pr}(N \geq s) &\geq \binom{K}{s} \theta^s (1-\theta)^{K-s} \cdot \text{pr}(\beta_i \geq 1/(2s) \text{ for all } i \leq s) \\ &\geq \binom{K}{s} \theta^s (1-\theta)^{K-s} (1 - 9s/\rho_1). \end{aligned}$$

This completes the proof. ■

**Proof** (Proof of Theorem 3) Denote by  $p_\beta$  the density of  $N(X\beta, I)$  distribution, and the corresponding likelihood ratio by

$$\Lambda_{\beta, \beta^*}(Y) = \frac{p_\beta}{p_{\beta^*}}(Y) = \exp \left\{ -\frac{1}{2} \|X(\beta - \beta^*)\|_2^2 + (Y - X\beta^*)^\top X(\beta - \beta^*) \right\}.$$

For any subset  $B \subset \Theta$ , by the Bayes formula, the posterior probability mass of its complement  $B^c$  is

$$\begin{aligned} \Pi(B^c|Y) &= \frac{\int_{B^c} \Lambda_{\beta, \beta^*}(Y) g(\beta; \theta, \rho_1, \rho_2) d\beta}{\int_{\Theta} \Lambda_{\beta, \beta^*}(Y) g(\beta; \theta, \rho_1, \rho_2) d\beta} \\ &= \frac{\int_{B^c} \Lambda_{\beta, \beta^*}(Y) g(\beta; \theta, \rho_1, \rho_2) d\beta}{\int_{B^c} \Lambda_{\beta, \beta^*}(Y) g(\beta; \theta, \rho_1, \rho_2) d\beta + \int_B \Lambda_{\beta, \beta^*}(Y) g(\beta; \theta, \rho_1, \rho_2) d\beta} \\ &=: \frac{N_1}{N_1 + N_2}. \end{aligned}$$

Take  $B = \{\beta \in \Theta : \|\beta - \beta^*\|_1 \leq R\}$  for some  $0 < R < 2$  that will be specified later. Without loss of generality, henceforth we assume that the first  $s$  elements of  $\beta^*$  are equal to  $1/s$  and the rest are zero, while noting that the same argument applies for any  $\beta^* \in \Theta(s, K)$ .

We first consider to bound  $N_2$ . For any  $0 < r < R$ , there holds

$$N_2 \geq \int_{\|\beta - \beta^*\|_1 < r} \Lambda_{\beta, \beta^*}(Y) D(\beta; \rho^*) \theta^s (1 - \theta)^{K-s} d\beta,$$

where  $D(\beta; \rho^*)$  is the Dirichlet distribution density with concentration parameter  $\rho^*$  whose first  $s$  elements are  $\rho_1$  and the rest are  $\rho_2$ . Now under  $\beta \sim \text{Dir}(\rho^*)$ , for  $1 \leq i \leq s$ , we have

$$E(\beta_i) = \frac{\rho_1}{\rho_1 s + \rho_2 (K - s)} = \frac{1}{s + (K - s) \rho_2 / \rho_1} \in \left( 1/s - s^{-2} K^{1-(\alpha_1 + \alpha_2)}, 1/s \right).$$

We discuss the case for  $s = 1$  and  $s \geq 2$  separately. By Chebyshev's inequality, when  $s \geq 2$ , we have for any  $\varepsilon_1 > 0$ ,

$$\text{pr}(|\beta_i - E(\beta_i)| > \varepsilon_1) \leq \frac{s^{-1}(1 - s^{-1})}{sK^{\alpha_1} + (K - s)K^{-\alpha_2} + 1} \frac{1}{\varepsilon_1^2}.$$

In view of the union bound, with probability at least  $1 - \{[sK^{\alpha_1} + (K - s)K^{-\alpha_2} + 1]\varepsilon_1^2\}^{-1}$ , we have  $\sum_{i=1}^s |\beta_i - E(\beta_i)| \leq s\varepsilon_1$ . Combining this with the triangle inequality yields

$$\sum_{i=1}^s |\beta_i - \beta_i^*| \leq \sum_{i=1}^s |\beta_i - E(\beta_i)| + \sum_{i=1}^s |E(\beta_i) - \beta_i^*| \leq s\varepsilon_1 + s^{-1} K^{1-(\alpha_1 + \alpha_2)}. \quad (6)$$

For  $i > s$ , under  $\beta \sim \text{Dir}(\rho^*)$ , the marginal distribution of  $\sum_{i>s} \beta_i$  is Beta  $(\rho_2(K - s), \rho_1 s)$ . Therefore,

$$E \left( \sum_{i>s} \beta_i \right) = \frac{\rho_2(K - s)}{\rho_2(K - s) + \rho_1 s} \in \left( 0, s^{-1} K^{1-(\alpha_1 + \alpha_2)} \right).$$

Again, by Chebyshev's inequality, for any  $\varepsilon_2 > 0$ ,

$$\text{pr} \left( \left| \sum_{i>s} \beta_i - E \left( \sum_{i>s} \beta_i \right) \right| > \varepsilon_2 \right) \leq s^{-2} K^{1-(2\alpha_1+\alpha_2)} \varepsilon_2^{-2}.$$

Following a similar argument as used to derive (6), we obtain that with probability at least  $1 - s^{-2} K^{1-(2\alpha_1+\alpha_2)} \varepsilon_2^{-2}$ ,

$$\sum_{i>s} |\beta_i - \beta_i^*| = \sum_{i>s} \beta_i \leq \left| \sum_{i>s} \beta_i - E \left( \sum_{i>s} \beta_i \right) \right| + E \left( \sum_{i>s} \beta_i \right) \leq \varepsilon_2 + s^{-1} K^{1-(\alpha_1+\alpha_2)}. \quad (7)$$

Combining (6) and (7), with probability at least

$$1 - [sK^{\alpha_1} + (K-s)K^{-\alpha_2} + 1]^{-1} \varepsilon_1^{-2} - s^{-2} K^{1-(2\alpha_1+\alpha_2)} \varepsilon_2^{-2},$$

we have

$$\|\beta - \beta^*\|_1 \leq s\varepsilon_1 + \varepsilon_2 + 2s^{-1} K^{1-(\alpha_1+\alpha_2)}. \quad (8)$$

Take  $r = s\varepsilon_1 + \varepsilon_2 + 2s^{-1} K^{1-(\alpha_1+\alpha_2)} < R$ . Define the event  $\tau_0 = \{\|X^T(Y - X\beta^*)\|_\infty \leq \delta\}$ . Conditional on  $\tau_0$ ,  $N_2$  is lower bounded by a constant:

$$\begin{aligned} N_2 &\geq \int_{\|\beta - \beta^*\|_1 < r} \exp \left\{ -\frac{1}{2} \|X(\beta - \beta^*)\|_2^2 - \delta \|\beta - \beta^*\|_1 \right\} D(\beta; \rho^*) \theta^s (1 - \theta)^{K-s} d\beta \\ &\geq \exp \left( -\frac{1}{2} \|X\|^2 r^2 - \delta r \right) \theta^s (1 - \theta)^{K-s} \\ &\quad \cdot \left\{ 1 - [sK^{\alpha_1} + (K-s)K^{-\alpha_2} + 1]^{-1} \varepsilon_1^{-2} - s^{-2} K^{1-(2\alpha_1+\alpha_2)} \varepsilon_2^{-2} \right\} \\ &=: L_{N_2}. \end{aligned} \quad (9)$$

Hence, conditional on  $\tau_0$ , it holds that

$$\Pi(B^c|Y) \leq \frac{N_1}{N_1 + L_{N_2}}. \quad (10)$$

Let  $\mathbb{1}\tau_0$  and  $\mathbb{1}\tau_0^c$  indicate the occurrence of events  $\tau_0$  and  $\tau_0^c$ , respectively. Note that

$$E_{\beta^*} \{ \Pi(B^c|Y) \} = E_{\beta^*} \{ \Pi(B^c|Y) \mathbb{1}\tau_0 \} + E_{\beta^*} \{ \Pi(B^c|Y) \mathbb{1}\tau_0^c \} \leq E_{\beta^*} \{ \Pi(B^c|Y) \mathbb{1}\tau_0 \} + E_{\beta^*} (\mathbb{1}\tau_0^c). \quad (11)$$

It follows from (10) that

$$\begin{aligned} E_{\beta^*} \{ \Pi(B^c|Y) \mathbb{1}\tau_0 \} &\leq E_{\beta^*} \left[ \left( 1 - \frac{L_{N_2}}{N_1 + L_{N_2}} \right) \mathbb{1}\tau_0 \right] \\ &\leq 1 - L_{N_2} \left\{ E_{\beta^*} \left( \frac{1}{N_1 + L_{N_2}} \right) - E_{\beta^*} \left( \frac{\mathbb{1}\tau_0^c}{N_1 + L_{N_2}} \right) \right\} \\ &\leq 1 - \frac{L_{N_2}}{E_{\beta^*}(N_1) + L_{N_2}} + E_{\beta^*}(\mathbb{1}\tau_0^c), \end{aligned} \quad (12)$$

where the last line is a direct application of Jensen's inequality. Take  $\delta = 2\|X\|\sqrt{\log K}$ . By Lemma 4 in Castillo et al. (2015),  $E_{\beta^*}(\mathbb{1}_{\tau_0^c}) \leq 2/K$ . This combined with (11) and (12) yields

$$E_{\beta^*} \{ \Pi(B^c|Y) \} \leq 1 - \frac{L_{N_2}}{E_{\beta^*}(N_1) + L_{N_2}} + \frac{4}{K}.$$

It remains to show  $E_{\beta^*}(N_1) = o(L_{N_2})$ . Take  $\varepsilon_1 = (\log K/sK^{\alpha_1})^{1/2}$ ,  $\varepsilon_2 = s\varepsilon_1$ , and  $R = s \log K / \{\Phi(s) \min(\|X\|, K^{\alpha_1/2})\}$ . If  $\alpha_1 + \alpha_2 \geq 1$ ,

$$sK^{\alpha_1}\varepsilon_1^2 = \log(K), \quad s^2K^{-1+(2\alpha_1+\alpha_2)}\varepsilon_2^2 = s^3K^{\alpha_1+\alpha_2-1}\log(K) > \log(K).$$

Therefore,  $L_{N_2}$  defined in (9) satisfies that

$$L_{N_2} \geq \exp\left(-\frac{1}{2}\|X\|^2 r^2 - \delta r\right) \theta^s (1-\theta)^{K-s} \frac{1}{2}$$

for sufficiently large  $K$ . On the other hand, by the definition of  $\Phi(s)$ ,  $E_{\beta^*}(N_1)$  is upper bounded by

$$E_{\beta^*}(N_1) = \int_B \exp\left(-\|X(\beta - \beta^*)\|_2^2\right) g(\beta; \rho_1, \rho_2, \theta) d\beta \leq \exp\left(-\Phi^2(s)\|X\|^2 R^2 s^{-1}\right).$$

Combining these two inequalities, a sufficient condition for  $E_{\beta^*}(N_1) = o(L_{N_2})$  is

$$\exp\left(-\Phi^2(s)\|X\|^2 R^2 s^{-1}\right) \ll \exp\left(-\frac{1}{2}\|X\|^2 r^2 - \delta r - \left\{s \log\left(\frac{1-\theta}{\theta}\right) - K \log(1-\theta)\right\}\right). \quad (13)$$

We next show that the left side of (13) is bounded by each component on the right side. With our specific choice of  $\varepsilon_1$  and  $R$ , we obtain that

$$\begin{aligned} \frac{\|X\|^2 (s\varepsilon_1)^2}{\Phi^2(s)\|X\|^2 R^2 s^{-1}} &= \frac{\min(\|X\|^2, K^{\alpha_1})}{K^{\alpha_1} \log(K)} \leq \frac{1}{\log(K)} \\ \frac{\delta s \varepsilon_1}{\Phi^2(s)\|X\|^2 R^2 s^{-1}} &= \frac{2 \min(\|X\|^2, K^{\alpha_1})}{\log(K)\|X\|\sqrt{K^{\alpha_1}}\sqrt{s}} \leq \frac{2}{\log(K)\sqrt{s}} \\ \frac{s \log(K)}{\Phi^2(s)\|X\|^2 R^2 s^{-1}} &= \frac{\min(\|X\|^2, K^{\alpha_1})}{\|X\|^2 \log(K)} \leq \frac{1}{\log(K)}. \end{aligned} \quad (14)$$

When  $\alpha_1/2 + \alpha_2 \geq 1$ , it can be proved that  $s^{-1}K^{1-(\alpha_1+\alpha_2)} \leq s\varepsilon_1$ , and thus  $r \leq 4s\varepsilon_1$  in view of (8). Choosing  $\theta = t/K$  for  $1 \leq t \leq s$ , as  $K \rightarrow \infty$ , it holds that

$$s \log\{(1-\theta)/\theta\} - K \log(1-\theta) \rightarrow s \log(K-t) - s \log(t) + t \leq s \log(K).$$

Combining this with (13) and (14), we arrive at  $E_{\beta^*}(N_1) = o(L_{N_2})$ .

The proof for the  $s = 1$  case is almost identical to the  $s \geq 2$  case, with minor differences noted below. Under  $\beta \sim \text{Dir}(\rho^*)$ ,  $E(\beta_1) \in (1 - K^{1-(\alpha_1+\alpha_2)}, 1)$ . By Chebyshev's inequality, we have for any  $\varepsilon_1 > 0$ ,

$$\text{pr}(|\beta_1 - E(\beta_1)| > \varepsilon_1) \leq \frac{K^{1-(\alpha_1+\alpha_2)}}{K^{\alpha_1} + K^{1-\alpha_2}} \frac{1}{\varepsilon_1^2}.$$

For  $1 < i < K$  under  $\beta \sim \text{Dir}(\rho^*)$ , the marginal distribution of  $\sum_{i>1} \beta_i$  is  $\text{Beta}(\rho_2(K-1), \rho_1)$ . Using the same arguments for establishing (7) and (8), we obtain that with probability at least  $1 - K^{1-(2\alpha_1+\alpha_2)}\varepsilon_2^{-2}$ ,

$$\sum_{i>1} |\beta_i - \beta_i^*| \leq \varepsilon_2 + K^{1-(\alpha_1+\alpha_2)},$$

and that with probability at least  $1 - K^{1-(2\alpha_1+\alpha_2)}(\varepsilon_1^{-2} + \varepsilon_2^{-2})$ ,

$$\|\beta - \beta^*\|_1 \leq \varepsilon_1 + \varepsilon_2 + 2K^{1-(\alpha_1+\alpha_2)},$$

for large enough  $K$  and any  $\varepsilon_2 > 0$ . Take  $r = \varepsilon_1 + \varepsilon_2 + 2K^{1-(\alpha_1+\alpha_2)} < R$ . We also have (10) but with a different  $L_{N_2} = \exp(-\|X\|^2 r^2 / 2 - \delta r) \theta(1-\theta)^{K-1} \{1 - K^{1-(2\alpha_1+\alpha_2)}(\varepsilon_1^{-2} + \varepsilon_2^{-2})\}$ . As (11)-(12) remain unchanged, to prove the theorem, we only need to show  $E_{\beta^*}(N_1) = o(L_{N_2})$ .

Take  $\varepsilon_1 = \varepsilon_2 = (2 \log K / K^{\alpha_1})^{1/2}$  and  $R = \log K / \{\Phi(1) \min(\|X\|, K^{\alpha_1/2})\}$ . Since  $\alpha_1 + \alpha_2 > 1$ , there holds  $K^{1-(2\alpha_1+\alpha_2)}(\varepsilon_1^{-2} + \varepsilon_2^{-2}) \leq 1 / \log K$ . Therefore, for large enough  $K$ , the new  $L_{N_2}$  is also lower bounded by

$$L_{N_2} \geq \exp\left(-\frac{1}{2} \|X\|^2 r^2 - \delta r\right) \theta(1-\theta)^{K-1} \frac{1}{2}.$$

Noting that  $r \leq 3\varepsilon_1$  and choosing  $\theta = 1/K$ , the rest of the proof remains exactly the same. This completes the proof.  $\blacksquare$

## Appendix B. Additional Simulations with Varying Signal-to-noise Ratio

The signal-to-noise ratio (SNR) in Scenario 1 in Section 5.1 is 1.3. Here we conduct experiments with SNR varying from 0.6 to 2. This is done by adjusting the noise variance while fixing the design matrix and  $\beta$ . Similarly to Section 5.1, we use default hyperparameter values for the proposed method, and give an advantage to the symmetric Dirichlet approach by reporting its best average  $\ell_1$  error from a grid search for its hyperparameter. The two-step approaches, namely two-step lasso and two-step horseshoe, are also evaluated. Table 2 shows that all methods benefit from a stronger signal, and the proposed method leads to the smallest  $\ell_1$  error among all methods for all SNRs considered here.

## Appendix C. Additional Simulations with Correlated Predictors

We set  $n = 80$ ,  $K = 40$ ,  $s = 2$ , and assign nonzero weights to only the first two columns  $\beta_1 = \beta_2 = 0.5$ . All columns of  $X$  are independently drawn from  $N(0, 3^2)$ , with the exception that the second and the third columns are highly correlated with correlation 0.9. We apply the proposed method and calculate the posterior mean of  $\beta$ . This procedure is replicated for 100 times. Figure 5 shows the boxplot of the posterior mean  $\hat{\beta}_i$  for  $i = 1, 2, 3$ . We can see that  $\hat{\beta}_1$  concentrates tightly around 0.5, indicating good signal recovery for the first column. For the two correlated predictors,  $\hat{\beta}_2$  concentrates around 0.5 and  $\hat{\beta}_3$  around 0, both with a few outliers that come from the same replications. Therefore, our method is



SNR	Proposed	Symmetric Dirichlet	Two-step lasso	Two-step horseshoe
0.6	0.366 (0.0261)	0.879 (0.0263)	0.443(0.0421)	0.500 (0.0458)
0.8	0.181 (0.0215)	0.720 (0.0253)	0.303(0.0334)	0.275 (0.0380)
1.0	0.090 (0.0154)	0.581 (0.0228)	0.201(0.0316)	0.157 (0.0263)
1.2	0.047 (0.0098)	0.483 (0.0195)	0.156(0.0291)	0.105 (0.0221)
1.4	0.029 (0.0058)	0.414 (0.0159)	0.129(0.0253)	0.077 (0.0184)
1.6	0.019 (0.0015)	0.370 (0.0126)	0.089(0.0222)	0.075 (0.0179)
1.8	0.018 (0.0006)	0.334 (0.0091)	0.069(0.0193)	0.070 (0.0174)
2.0	0.017 (0.0001)	0.314 (0.0088)	0.053(0.0173)	0.070 (0.0174)

Table 2: Comparison of the mean  $\ell_1$  error of coefficient estimates of the proposed method, symmetric Dirichlet, and two-step methods at different signal-to-noise ratios. Results are based on 100 simulations. Standard errors are reported in parentheses.




Figure 5: Boxplot of posterior mean  $\hat{\beta}_1, \hat{\beta}_2, \hat{\beta}_3$  in 100 replications.

able to differentiate the two predictors most of the time, and is not arbitrarily selecting one out of the two highly correlated variables in this experiment. This is not surprising as under the exact structure the considered parameter space leads to a significant gap  $1/s$  between the weights for true signals and those for null signals due to partial constancy, which eases the identification of true signals even under high correlation.

**References**

Ray Bai, Gemma E Moran, Joseph L Antonelli, Yong Chen, and Mary R Boland. Spike-and-slab group lassos for grouped regression and sparse generalized additive models. *Journal of the American Statistical Association*, 117(537):184–197, 2022.

John M Bates and Clive WJ Granger. The combination of forecasts. *Journal of the Operational Research Society*, 20(4):451–468, 1969.

Simon Bernard, Laurent Heutte, and Sebastien Adam. On the selection of decision trees in random forests. In *International Joint Conference on Neural Networks*, pages 302–307.

- IEEE, 2009.
- G erard Biau and Erwan Scornet. A random forest guided tour. *Test*, 25(2):197–227, 2016.
- Małgorzata Bogdan, Ewout Van Den Berg, Chiara Sabatti, Weijie Su, and Emmanuel J Cand es. Slope—adaptive variable selection via convex optimization. *The Annals of Applied Statistics*, 9(3):1103, 2015.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Philip J Brown, Marina Vannucci, and Tom Fearn. Bayesian wavelength selection in multicomponent analysis. *Journal of Chemometrics*, 12(3):173–182, 1998.
- Peter B uhlmann and Sara Van De Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Science & Business Media, 2011.
- Florentina Bunea, Alexandre B Tsybakov, and Marten H Wegkamp. Aggregation for Gaussian regression. *The Annals of Statistics*, 35(4):1674–1697, 2007.
- Carlos M Carvalho, Nicholas G Polson, and James G Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010.
- Isma el Castillo, Johannes Schmidt-Hieber, and Aad Van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986–2018, 2015.
- Andrew G Chapple, Marina Vannucci, Peter F Thall, and Steven Lin. Bayesian variable selection for a semi-competing risks model with three hazard functions. *Computational Statistics   Data Analysis*, 112:170–185, 2017.
- Kenneth L Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms*, 6(4):1–30, 2010.
- Robert T Clemen. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5(4):559–583, 1989.
- Cristina Conflitti, Christine De Mol, and Domenico Giannone. Optimal combination of survey forecasts. *International Journal of Forecasting*, 31(4):1096–1103, 2015.
- Francis X Diebold and Peter Pauly. The use of prior information in forecast combination. *International Journal of Forecasting*, 6(4):503–508, 1990.
- Francis X Diebold and Minchul Shin. Machine learning for regularized survey forecast combination: Partially-egalitarian lasso and its derivatives. *International Journal of Forecasting*, 35(4):1679–1691, 2019.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Stephen E Fienberg and Paul W Holland. Simultaneous estimation of multinomial cell probabilities. *Journal of the American Statistical Association*, 68(343):683–691, 1973.

- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- Andrew Gelman and Donald B Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992.
- Irving John Good. A Bayesian significance test for multinomial distributions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 29(3):399–418, 1967.
- Clive WJ Granger and Ramu Ramanathan. Improved methods of combining forecasts. *Journal of Forecasting*, 3(2):197–204, 1984.
- Matthew Heiner, Athanasios Kottas, and Stephan Munch. Structured priors for sparse probability vectors with application to model selection in Markov chains. *Statistics and Computing*, 29(5):1077–1093, 2019.
- Wenxin Jiang and Martin A Tanner. Gibbs posterior for variable selection in high-dimensional classification and data mining. *The Annals of Statistics*, 36(5):2207–2231, 2008.
- Vrushali Y Kulkarni and Pradeep K Sinha. Pruning of random forest classifiers: A survey and future directions. In *International Conference on Data Science & Engineering (ICDSE)*, pages 64–68. IEEE, 2012.
- Anastasios Kyrillidis, Stephen Becker, Volkan Cevher, and Christoph Koch. Sparse projections onto the simplex. In *International Conference on Machine Learning*, pages 235–243, 2013.
- Ping Li, Syama Sundar Rangapuram, and Martin Slawski. Methods for sparse and low-rank recovery under simplex constraints. *Statistica Sinica*, 30(2):557–577, 2020.
- Andy Liaw and Matthew Wiener. Classification and regression by randomforest. *R News*, 2(3):18–22, 2002. URL <https://CRAN.R-project.org/doc/Rnews/>.
- Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36(1):54–74, 2020.
- Sonia Migliorati, Andrea Ongaro, and Gianna S Monti. A structured Dirichlet mixture model for compositional data: inferential and applicative issues. *Statistics and Computing*, 27(4):963–983, 2017.
- Andrea Ongaro and Sonia Migliorati. A generalization of the Dirichlet distribution. *Journal of Multivariate Analysis*, 114:412–426, 2013.
- Andrea Ongaro, Sonia Migliorati, and Roberto Ascari. A new mixture model on the simplex. *Statistics and Computing*, 30(4):749–770, 2020.

- David S Palmer, Noel M O’Boyle, Robert C Glen, and John BO Mitchell. Random forest models to predict aqueous solubility. *Journal of Chemical Information and Modeling*, 47(1):150–158, 2007.
- Eric C Polley and Mark J Van der Laan. Super learner in prediction. *Berkeley Division of Biostatistics Working Paper Series*, 2010.
- Veronika Rockova and Edward I George. The spike-and-slab lasso. *Journal of the American Statistical Association*, 113(521):431–444, 2018.
- Judith Rousseau and Kerrie Mengersen. Asymptotic behaviour of the posterior distribution in overfitted mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(5):689–710, 2011.
- Christian Schäfer. *Monte Carlo methods for sampling high-dimensional binary vectors*. PhD thesis, Université Paris Dauphine-Paris IX, 2012.
- Mark R Segal. Machine learning benchmarks and random forest regression. *UCSF: Center for Bioinformatics and Molecular Biostatistics*, 2004.
- Jeremy Smith and Kenneth F Wallis. A simple explanation of the forecast combination puzzle. *Oxford Bulletin of Economics and Statistics*, 71(3):331–355, 2009.
- James H Stock and Mark W Watson. Combination forecasts of output growth in a seven-country data set. *Journal of Forecasting*, 23(6):405–430, 2004.
- Weijie Su, Małgorzata Bogdan, and Emmanuel Candes. False discoveries occur early on the lasso path. *The Annals of Statistics*, 45(5):2133–2150, 2017.
- Vladimir Svetnik, Andy Liaw, Christopher Tong, J Christopher Culberson, Robert P Sheridan, and Bradley P Feuston. Random forest: A classification and regression tool for compound classification and QSAR modeling. *Journal of Chemical Information and Computer Sciences*, 43(6):1947–1958, 2003.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Mark J Van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1), 2007.
- Yun Yang and David B Dunson. Minimax optimal Bayesian aggregation. *arXiv preprint arXiv:1403.1345*, 2014.
- Yun Yang, Martin J Wainwright, and Michael I Jordan. On the computational complexity of high-dimensional Bayesian variable selection. *The Annals of Statistics*, 44(6):2497–2532, 2016.
- I-C Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998.
- Heping Zhang and Minghui Wang. Search for the smallest random forest. *Statistics and its Interface*, 2(3):381, 2009.