

Intuitive joint priors for Bayesian linear multilevel models: The R2D2M2 prior

Javier Enrique Aguilar¹ and Paul-Christian Bürkner²

¹*Cluster of Excellence SimTech, University of Stuttgart*
e-mail: javier.aguilar-romero@simtech.uni-stuttgart.de
url: <https://jear2412.github.io>

²*Cluster of Excellence SimTech, University of Stuttgart*
e-mail: paul.buerkner@gmail.com
url: <https://paul-buerkner.github.io>

Abstract: The training of high-dimensional regression models on comparably sparse data is an important yet complicated topic, especially when there are many more model parameters than observations in the data. From a Bayesian perspective, inference in such cases can be achieved with the help of shrinkage prior distributions, at least for generalized linear models. However, real-world data usually possess multilevel structures, such as repeated measurements or natural groupings of individuals, which existing shrinkage priors are not built to deal with.

We generalize and extend one of these priors, the R2D2 prior by Zhang et al. (2020), to linear multilevel models leading to what we call the R2D2M2 prior. The proposed prior enables both local and global shrinkage of the model parameters. It comes with interpretable hyperparameters, which we show to be intrinsically related to vital properties of the prior, such as rates of concentration around the origin, tail behavior, and amount of shrinkage the prior exerts.

We offer guidelines on how to select the prior’s hyperparameters by deriving shrinkage factors and measuring the effective number of non-zero model coefficients. Hence, the user can readily evaluate and interpret the amount of shrinkage implied by a specific choice of hyperparameters.

Finally, we perform extensive experiments on simulated and real data, showing that our inference procedure for the prior is well calibrated, has desirable global and local regularization properties and enables the reliable and interpretable estimation of much more complex Bayesian multilevel models than was previously possible.

Keywords and phrases: Bayesian inference, multilevel models, prior specification, shrinkage priors, regularization.

Received August 2022.

1. Introduction

Regression models are ubiquitous in the quantitative sciences and industry, making up a big part of all statistical data analyses. Their success can be explained by a combination of several factors, involving the ease of interpretation of their additive structure, rich mathematical theory, and (relative) simplicity of their estimation, to name a few key aspects (Gelman et al., 2020a; Harrell, 2013).

Despite these advantages, the vast amount of modeling options and required analyst choices render building interpretable, robust, and well-predicting regression models a highly difficult task.

As the complexity of the analyzed data increases, so does the required complexity of the applied regression models and we are bound to encounter multilevel structures which contain overall (population-level or fixed effects) and varying (group-specific, group-level, or random effects) terms. Overall effects influence every member of the population equally, whereas varying effects do so per group. For example, multilevel structures can be found in Psychology, Medicine, or Biology, due to the natural groupings of individuals, or repeated measurement of the same individuals in experiments or longitudinal studies. Multilevel models are designed specifically to account for the nested structure in multilevel data and are a widely applied class of regression models (Bates et al., 2015; Gelman and Hill, 2006; Bürkner, 2021). Their key statistical idea is to assume that a set of regression coefficients, defined according to the multilevel structure, originate from the same underlying distribution whose hyperparameters are subsequently estimated from the data. The implied partial-pooling property increases the robustness of the parameter estimates, helps with uncertainty calibration, and improves out-of-sample predictive performance (Gelman et al., 2013; Gelman and Hill, 2006). While multilevel models can be estimated in both frequentist and Bayesian frameworks (Bates et al., 2015; Bürkner, 2017), we will focus on the latter framework in this work, since it offers considerable flexibility in the specification and regularization of multilevel models (Gelman et al., 2013).

From a Bayesian perspective, the above mentioned distributions of the regression coefficients constitute prior distributions (priors) that describe the uncertainty in the model parameters before seeing the data. The desirable properties of multilevel models can then be explained by the fact that these form joint priors over a set of parameters with shared hyperparameters, rather than separate independent priors for each parameter (Gelman et al., 2013). Joint priors are a Bayesian success story more generally. For example, joint priors can help to improve the predictive performance of individual additive terms parameterized by more than one parameter, for example, splines, spatial, or monotonic effects (Bürkner and Charpentier, 2020; Wood, 2017; Morris et al., 2019). With very few exceptions (Fuglstad et al., 2019; Goodrich et al., 2020; Yanchenko et al., 2021), the development of joint priors for Bayesian multilevel models has been limited to individual additive terms. In contrast, different terms, corresponding to different parameter sets, still receive mutually independent priors, for example, independent (inverse-)gamma, uniform, or half Student- t priors for variance or standard deviation parameters (Browne and Draper, 2006; Bürkner, 2017; Simpson et al., 2017; Depaoli and Clifton, 2015). As more and more terms are being added to the model while the number of observations remains constant, such models will overfit the data, leading to unreliable or uninterpretable estimates as well as bad out-of-sample predictions (Van Erp et al., 2019). Also, overall Type I error rates may be severely inflated if many terms with mutually independent priors are tested (Gelman et al., 2013).

Without advanced regularization methods such as joint priors, reliable and interpretable estimation of the required highly parameterized models is very hard to achieve. This is particularly obvious in the high dimensional case regime (more covariates than observations: $p > N$), where the number of model parameters p becomes larger than the number of observations N in the data and unregularized regression models would not be estimable (Hoerl and Kennard, 1970); but even if $p < N$, regression models often benefit strongly from regularization, for example, to prevent overfitting and Type I error inflation (Hoerl and Kennard, 1970; Gelman et al., 2013; Van Erp et al., 2019). While all of these difficulties apply to standard regression models already, they become even worse for multilevel models as the number of additive terms to explore further increases when considering all valid combinations of predictor coefficients and grouping factors which are permitted by the given data structure (Catalina et al., 2020; Barr et al., 2013; Paananen et al., 2020).

In the context of *single-level* linear models, that is, linear models without multilevel structure, a wide range of joint shrinkage priors have been developed (Park and Casella, 2008; Carvalho et al., 2010; Piironen and Vehtari, 2017; Bhattacharya et al., 2015). These models enable the estimation, interpretation, and selection of predictor terms and prevent overfitting even when the number of predictors is large and/or the data is small (Van Erp et al., 2019; van der Pas, 2021). On an abstract level, these priors can all be described as Global-Local (GL) priors because they have both, global parameters controlling the shrinkage of all terms jointly and local parameters controlling the (relative) shrinkage of individual terms (van der Pas, 2021). Usually, they are combined with a location-scale distribution, such as normal or a double exponential, where the location is fixed to zero and the scale is computed as a product of global and local parameters. While these priors have been very successful within their supported model class (Van Erp et al., 2019; van der Pas, 2021), they remain limited in several key aspects which we aim to address in our work:

1. Most existing joint priors for regression models require the manual specification of hyperparameters with potentially strong impact on the obtained inference (Van Erp et al., 2019). These hyperparameters are often unintuitive and abstract especially for non-statistical experts, making these priors' practical application and communication much more difficult.
2. High dimensionality in multilevel models has so far been studied in cases where the number of overall coefficients p is big while the amount of varying terms q is very small, $p > N, q \ll N$. For example, in their papers, Buehlmann et al. (2014) and Lin et al. (2020) carried out theoretical research and simulation studies of frequentist multilevel models with large p yet $q \leq 2$.
3. Existing joint priors on multilevel models focus only on the group-specific coefficients (Fuglstad et al., 2019; Goodrich et al., 2020), without at the same time considering the overall coefficients. Further, it remains unclear how these priors scale with the number of predictors when their coefficients vary across an increasing number of grouping factors. However, in order to

ensure consistent and robust regularization in multilevel models, we need to define a single global joint prior on all coefficients, and ensure it scales well to large numbers of additive terms and complex multilevel structure.

An initial step towards addressing these issues, at least for single-level linear models, is the R2D2 prior developed by Zhang et al. (2020). The main idea is to specify a prior on the coefficient of determination R^2 – also known as *proportion of explained variance* – and decompose the *explained variance* into individual variance components via Dirichlet Decomposition (abbreviated as D2), thus propagating uncertainty from R^2 on to the regression coefficients. We further motivate the use of priors over R^2 in Section 2. The R2D2 prior shares several desirable properties with other global-local shrinkage priors for single-level linear models (Bhadra et al., 2016; Carvalho et al., 2010; Zhang et al., 2020) and is easy to understand even for non-expert analysts since the only hyperparameters it requires are the shape parameters of the prior on R^2 and a single concentration parameter for the Dirichlet distribution. However, its scope is limited to single-level linear models. These simple models do not live up to the requirements of model-based research, where multilevel data structures are omnipresent, so a much more general solution is needed.

Therefore, the aim of our paper is to fill this gap by generalizing the R2D2 prior to models with increased complexity in terms of multilevel structures. We name our new prior the *R2D2M2 prior*, where M2 stands for *Multilevel Models*. In parallel to our work, related ideas were put forward by Yanchenko et al. (2021) although with a different focus: they concentrate on defining useful R^2 measures for non normal likelihoods without studying the theoretical properties on the resulting priors as we do here.

1.1. Main contributions

Our main contributions are as follows:

- We propose a prior over a global R^2 that jointly regularizes both the overall and varying coefficients at the same time. Our new prior’s hyperparameters are easier to interpret, which facilitates their specification by the user.
- We expand upon the theoretical properties of the original R2D2 prior by considering normal base for distributions for the coefficients. We study concentration properties around the origin and behavior in the tails and show how both are dependent on the hyperparameters selected by the user.
- We propose the use of shrinkage factors to evaluate how the chosen hyperparameters determine the amount of shrinkage for each individual coefficient. We also propose a global measure of shrinkage, which quantifies the effective total amount of non-zero coefficients in the model. Both of these quantities were not considered for the R2D2 prior before.
- We implement both the R2D2 prior and the R2D2M2 prior in the probabilistic programming language Stan (Stan Development Team, 2022; Car-

penter et al., 2017) and demonstrate that inference for our implementations is well calibrated. We also provide an implementation in the R package `brms` (Bürkner, 2017, version 2.19.2 or higher), which provides a high level interface to fit Bayesian generalized (non-)linear multivariate multilevel models using Stan. We further perform extensive simulations to test the capabilities of the R2D2M2 prior and discuss its shrinkage behavior, out-of-sample predictive performance, as well as the relationship of the latter with the amount of shrinkage. We also show that frequentist error metrics such as Type I and Type II errors can be controlled with the help of our prior even in complex scenarios.

- Our joint prior offers the possibility to study cases in which both the number of overall coefficients p and varying terms q are large and potentially both greater than N , that is, high dimensionality in both the overall and varying coefficients. It is to the best of our knowledge the first method that provides this. To demonstrate the prior's applicability and usefulness in such high-dimensional cases not only in simulations but also in the real world, we reanalyse the riboflavin dataset provided by Buehlmann et al. (2014).

The remainder of the paper is organized as follows: Section 2 starts by motivating the use of priors over R^2 (and thus the use of the R2D2M2 prior) by showing how weakly informative priors on the coefficients can translate into very informative priors for R^2 . The construction of the R2D2M2 prior is presented along with insights that lead to better understanding of the prior. Section 3 presents theoretical properties of both the R2D2 and R2D2M2 prior such as marginal distributions and concentration properties as well as guidelines for hyperparameter selection. We introduce the concept of shrinkage factors for both the R2D2 and R2D2M2 prior as well as the relationship the prior's hyperparameters have with the effective number of nonzero coefficients in the model, thus extending the idea of Piironen and Vehtari (2017) to multilevel models. This provides an easy and intuitive way of setting up the prior, since it is possible to visualize how much sparsity will be induced. Section 4 shows the results of testing the R2D2M2 prior in intensive simulations. We make use of simulation based calibration (Talts et al., 2020) and provide evidence that that our implementation and estimation of both the R2D2 and the R2D2M2 prior models are well calibrated. What is more, we simulate data from sparse multilevel models and provide detailed discussions about how the prior behaves with respect to estimation error, out-of-sample predictive performance, credible interval coverage and Type I and Type II errors. Finally, we show how the R2D2M2 prior performs in real life data by testing it on the Riboflavin production data set made publicly available by Buehlmann et al. (2014). The paper ends with a discussion in Section 5 where we summarize our findings and suggest future lines of research. All mathematical proofs can be found in Appendix A.

2. The R2D2M2 prior

In the following, we denote the response variable as y , its n th observed value as y_n where $n \in \{1, \dots, N\}$, and the predictor variables (features or covariates) as x_i with $i \in \{1, \dots, p\}$. Additionally we assume conditional independence of the N observations in the data. We define the overall coefficients as those that affect every member of the general population and the group or varying coefficients as those that do so per group. In the statistical literature, the terms overall and varying coefficients have been known as fixed and random effects, respectively, nevertheless we believe that this nomenclature is problematic in the Bayesian setting, since all of the parameters of interest are treated as (random) variables whose uncertainty is represented via probability distributions (Gelman and Hill, 2006). We denote the overall coefficients by $b_i, i \in \{0, \dots, p\}$ and the group specific or varying coefficients by u_{ig_j} . Here g denotes a categorical grouping variable, encoding the individual groups g_j across which the coefficient of the i th predictor x_i is expected to vary. If we denote G_i as the index set of all grouping variables across which the coefficient of x_i is expected to vary, we can write the observation model of a linear multilevel model as

$$y_n \sim \text{Normal}(\mu_n, \sigma^2) \quad (1)$$

$$\mu_n = b_0 + \sum_{i=1}^p x_{ni} b_i + \sum_{g \in G_0} u_{0g_j[n]} + \sum_{i=1}^p x_{ni} \left(\sum_{g \in G_i} u_{ig_j[n]} \right), \quad (2)$$

for $n \in \{1, \dots, N\}, i \in \{0, \dots, p\}, g \in \{1, \dots, G_i\}$ and $j \in J_g$, where J_g is the index set over the levels of grouping variable g . Here, b_0 and u_{0g_j} denote the overall intercept and varying intercepts, respectively, whereas the former will not be subject to extra prior regularization (see below).

When doing Bayesian inference in multilevel models, it is a common practice to set independent normal priors on all regression coefficients:

$$b_i \sim \text{Normal}(0, \tilde{\lambda}_i^2), \quad u_{ig_j} \sim \text{Normal}(0, \tilde{\lambda}_{ig}^2), \quad (3)$$

where, in current practice, $\tilde{\lambda}_i^2, \tilde{\lambda}_{ig}^2$ are either fixed or are assigned independent priors. The literature shows that either fixing the values of $\tilde{\lambda}$ or assigning independent priors for $\tilde{\lambda}$ is suboptimal, as the total prior variance of the model scales with the number of included terms, which is implausible in reality (Fuglstad et al., 2019). Beyond that, it remains unclear how these priors scale with the number of predictors when their coefficients vary across an increasing number of grouping factors.

2.1. Implied priors on R^2

The R^2 measure expresses the proportion of variance explained by the model in relation to the total variance σ_y^2 of the response y (Gelman et al., 2020a). For

Gaussian models with residual standard deviation σ we define

$$R^2 := \text{corr}^2(y, \mu) = 1 - \frac{\sigma^2}{\sigma_y^2}, \quad (4)$$

as a *global* measure of proportion of variance explained (Rights and Sterba, 2019). Global means that it jointly comprises all overall and group specific coefficients at the same time. This stands in contrast to some common definitions of R^2 in multilevel models, where R^2 is defined for each level of the model separately, rather than jointly across all levels and terms (Nakagawa and Schielzeth, 2013; Rights and Sterba, 2019). Our work emphasizes the use of a global R^2 as our aim is a joint regularization of all regression coefficients. In that sense, we use the above R^2 metric simply as a starting point for a joint regularizing prior for multilevel models, independently of which metric one would prefer to measure explained variances after estimating the model.

Given that R^2 is invariant under changes to the mean of the response y or any of the predictors x_i , we can assume $\mathbb{E}(y) = \mathbb{E}(x_i) = 0, \forall i \in \{1, \dots, p\}$ without loss of generality in the following. This implies that the model's overall intercept b_{00} is zero and we can rewrite R^2 as

$$R^2 = \frac{\text{var}(\mu)}{\text{var}(\mu) + \sigma^2}. \quad (5)$$

Assume a-priori independence of the regression coefficients and consider a prior for b_i and u_{ig_j} that satisfies $\mathbb{E}(b_i) = 0, \mathbb{E}(u_{ig_j}) = 0$. Additionally, denote by $\sigma_{x_i}^2$ the variance of the i th covariate, then we can write

$$\text{var}(\mu) = \sum_{g \in G_0} \tilde{\lambda}_{0g}^2 + \sum_{i=1}^p \sigma_{x_i}^2 \left(\tilde{\lambda}_i^2 + \sum_{g \in G_i} \tilde{\lambda}_{ig}^2 \right). \quad (6)$$

Equation (6) shows that $\text{var}(\mu)$ can be decomposed into the variance components belonging to the overall and varying effects, respectively. Furthermore, the variance corresponding to the varying coefficients can be sub-decomposed into the variance associated to the varying intercepts and varying slopes.

The decomposition of $\text{var}(\mu)$ shown in Equation (6) allows us to investigate the effect the standard deviations $\tilde{\lambda}$ (or variances $\tilde{\lambda}^2$) have on the implied prior for R^2 . As an example, consider a simple linear model without grouping variables $y_n \sim \text{Normal}(\mu_n, \sigma^2)$ where $\mu_n = \sum_{i=1}^p x_{ni} b_i$, with independent $\text{Normal}(0, \tilde{\lambda}_i^2)$ priors on b_i with $\tilde{\lambda}_i = 1$, and an $\text{Exponential}(1)$ prior on σ . These type of priors are usually considered weakly informative when the variables are standardized, however Figure 1 shows that the implied prior on R^2 is actually very informative. The implied prior on R^2 becomes increasingly narrow with a peak close to $R^2 = 1$ as the number of predictors increases. This behavior becomes even stronger for larger values of λ_i and when adding more terms. Therefore, it is also present in the context of multilevel models, where including additional levels per grouping factor and/or additional grouping factors increases the amount of

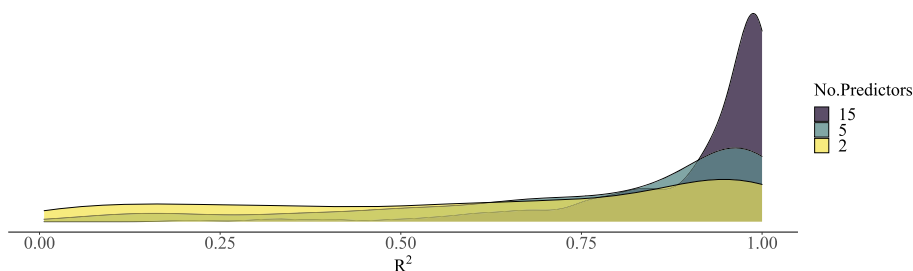


FIG 1. Implied prior on R^2 varying number of predictors in a single level linear model with $\text{Normal}(0, 1)$ priors on the coefficients and an $\text{Exponential}(1)$ prior on the error σ . Notice how weakly informative priors implicate a highly informative prior on R^2 .

parameters substantially. For instance, when considering one grouping factor g with L different levels and all p coefficients varying across the L levels, one would have $(p + 1)L$ additional coefficients (the 1 is for the varying intercept).

An $R^2 \approx 1$ represents almost perfect in-sample predictions. This is usually an indicator of overfitting when relevant amounts of noise are present in the data, such that $R^2 \approx 1$ can only be achieved by partially fitting on noise. Any prior implying such a large R^2 is highly unrealistic and would not be able to regularize sufficiently. Accordingly, it is more sensible to specify a prior on the overall variance explained by the model and then decompose this variance into components assigned to the individual terms, thereby inducing a joint prior over the variances of all terms. For this purpose, we propose (as have others (Goodrich et al., 2020; Zhang et al., 2020)) to specify a prior directly on R^2 as it provides an intuitive and widely understood measure of model fit.

As mentioned before, our main intention is to generalize the R2D2 prior proposed by Zhang et al. (2020) to models with increased complexity in terms of multilevel structures. On account of the inclusion of *Multilevel Models* (M2) in the R2D2 prior, we name our new prior the *R2D2M2 prior*.

2.2. Derivation of the R2D2M2 prior

In the following we present the derivation of our new prior. For this purpose, we begin by decomposing the priors' variances $\tilde{\lambda}_i^2$ and $\tilde{\lambda}_{ig}^2$ as

$$\tilde{\lambda}_i^2 = \frac{\sigma^2}{\sigma_{x_i}^2} \lambda_i^2, \quad \tilde{\lambda}_{ig}^2 = \frac{\sigma^2}{\sigma_{x_i}^2} \lambda_{ig}^2. \quad (7)$$

The factor $\frac{\sigma^2}{\sigma_{x_i}^2}$ accounts for the scales of y and x_i and ensures that λ_i represents prior variances on standardized variables, which are comparable across terms and λ_i are appropriate for any scale of y and x . In practice, σ_{x_i} is unknown and it could be replaced by the sample variance $\hat{\sigma}_{x_i}^2$ obtained from the data.

Plugging (7) into (6) results in

$$\text{var}(\mu) = \sigma^2 \left(\sum_{g \in G_0} \lambda_{0g}^2 + \sum_{i=1}^p \left(\lambda_i^2 + \sum_{g \in G_i} \lambda_{ig}^2 \right) \right) = \sigma^2 \tau^2,$$

where we have taken τ^2 as the sum of standardized prior variances given by

$$\tau^2 = \sum_{g \in G_0} \lambda_{0g}^2 + \sum_{i=1}^p \left(\lambda_i^2 + \sum_{g \in G_i} \lambda_{ig}^2 \right). \quad (8)$$

τ^2 is the (standardized) explained variance in correspondence to the interpretation of R^2 as proportion of explained variance. Using $\text{var}(\mu) = \sigma^2 \tau^2$, R^2 can be written as

$$R^2 = \frac{\text{var}(\mu)}{\text{var}(\mu) + \sigma^2} = \frac{\sigma^2 \tau^2}{\sigma^2 \tau^2 + \sigma^2} = \frac{\tau^2}{\tau^2 + 1}. \quad (9)$$

Following Zhang et al. (2020), we set a Beta distribution on R^2 and we write $R^2 \sim \text{Beta}(\cdot, \cdot)$. The Beta distribution offers a wide number of different parametrizations (Kruschke, 2015). Here we choose to parametrize the Beta distribution on R^2 in terms of the prior mean μ_{R^2} and prior “precision” φ_{R^2} (also known as mean-sample size parameterization), since it provides users with an intuitive way to incorporate prior knowledge into the R^2 prior; Both hyperparameters are understandable and expressible in terms of domain knowledge that represent the existing relationship between the included covariates and the response variable.

Equation (9) implies that, by definition, τ^2 has a Beta-Prime distribution with parameters μ_{R^2}, φ_{R^2} , represented by $\tau^2 \sim \text{BetaPrime}(\mu_{R^2}, \varphi_{R^2})$. Figure 2 illustrates the flexibility that the Beta distribution can offer in expressing prior knowledge about R^2 . The corresponding Beta-Prime prior for τ^2 is also shown. For example, for values of $(\mu_{R^2}, \varphi_{R^2}) = (0.5, 1)$, we get a bathtub-shaped prior on R^2 that places most mass near the two extremes $R^2 = 0$ and $R^2 = 1$. A priori, this indicates that the user expects the model to contain a substantial amount of either noise or signal.

In a consecutive decomposition step, we follow a Global-Local (GL) prior framework and set

$$\lambda_i^2 = \phi_i \tau^2, \quad \lambda_{ig}^2 = \phi_{ig} \tau^2, \quad (10)$$

with a vector ϕ containing all elements $\phi_i, \phi_{ig} \geq 0, \forall i \in \{1, \dots, p\}, g \in G_i$ such that $\sum_i \phi_i + \sum_i \sum_{g \in G_i} \phi_{ig} = 1$, thus rendering ϕ as a simplex. ϕ_i and ϕ_{ig} represent the proportion of standardized explained variance attributed to the i th and ig th term respectively. Decomposition (10) describes the proportion of standardized explained variance attributable to the ig th term, therefore controlling the variability allocated to each individual variance λ_{ig}^2 .




FIG 2. Exemplary densities of the Beta prior for R^2 (left) and corresponding Beta-Prime prior for τ^2 (right) with varying mean and precision parameters μ, φ respectively.

A natural prior for ϕ is a Dirichlet distribution with hyperparameter α chosen by the user, where α is a vector of length equal to the number of variance components that are included in the model and it governs the shape of the Dirichlet distribution. For instance, if the model has p overall coefficients, one grouping factor g and q varying coefficients in the grouping factor, then α is of length $p + q + 1$, due to the inclusion of the varying intercept. The amount of levels L does not affect the length of α according to our prior specification (e.g., see Figure 3 below). We denote the elements of α as α_i and α_{ig} , which represent the a priori expected relevance on each overall and varying regression term, respectively. In our context we will call α a concentration vector, since $\alpha_0 = \sum_i \alpha_i + \sum_{ig} \alpha_{ig}$ determines the narrowness of the distribution, that is, how peaked it is (Lin, 2016).

Naturally, the question arises of how α is to be specified. When $\alpha = (1, \dots, 1)'$ the prior is flat over all possible simplexes, which renders this value a good choice in the absence of additional prior knowledge. More generally, setting $\alpha = (a_\pi, \dots, a_\pi)'$ with a concentration parameter $a_\pi > 0$ (a symmetric Dirichlet distribution) can also be sensible since it drastically reduces the number of hyperparameters to specify, providing us with the ability to globally control the shape of the Dirichlet distribution with a single value. Additionally, the user can control the induced sparsity of the model by use of this single value as we show in Section 3.1. A symmetric Dirichlet distribution is useful when there is no prior preference that favours one component over another. Values for which $a_\pi < 1$ produce simplexes that concentrate their mass on the edges (“bathtub” distribution) and most coefficient values within a single sample will be close to 0. Instead, setting $a_\pi > 1$ results in concentration around the center of the simplex, favouring an even distribution among the different components (Lin, 2016). Finally, the user can also specify asymmetric Dirichlet distributions by designating different values for the components of α , which represent the differently a priori expected importance of the corresponding regression term.

The values of τ^2 and ϕ_{ig} play different roles in the prior. The explained variance τ^2 controls the amount of global shrinkage and is therefore deemed as

a global scale, whereas the i gth attributed variance ϕ_{ig} serves as a local scale that controls the individual shrinkage the prior induces over each specific term. The components of ϕ compete to increase and to acquire a higher proportion from the total variance τ . This can be seen when examining the covariance between two different components, which is given by $\text{cov}(\phi_i, \phi_j) = \frac{-\alpha_i \alpha_j}{\alpha_0^2(\alpha_0 + 1)}$, $i \neq j$. Thus, for a fixed value of τ^2 , as the importance of one term increases the importance of other terms decrease, and vice versa.

In the GL framework, the next step is to assign a base distribution that is able to express the attributed variance to each regression term b_i and u_{igj} . In their original R2D2 paper, Zhang et al. (2020) consider the coefficients to follow double exponential base distributions since they ensure both a higher concentration of mass around a zero and heavier tails than normal distributions. In the present paper, we will instead employ normal base distributions for each regression term for several, not mutually independent reasons:

1. It is customary in the multilevel model literature to consider normal base distributions for the overall and varying coefficients Gelman and Hill (2006); Wakefield (2013). What is more, Global-Local priors with normal base distributions are very flexible and have been the gold standard when proposing shrinkage priors, as can be seen in (Carvalho et al., 2010; Bhattacharya et al., 2015; Van Erp et al., 2019), among others.
2. Posterior approximation via gradient-based MCMC methods encounter issues when sampling from non-differentiable distributions, such as the double exponential (Brooks et al., 2011). Additionally, it is straightforward to reparametrize the posterior distribution to improve the rates of convergence of either Gibbs or Hamiltonian Monte Carlo based MCMC samplers in the context of conditionally normal multilevel models (Bürkner, 2017; Zanella and Roberts, 2021).
3. The use of normal distributions provides us with closed analytical expressions for several quantities of interest as described in Section 3 and shows, for example, good shrinkage and tail behavior.

To ease prior specification and speed up sampling, we will set a prior on the intercept \tilde{b}_{00} implied when all $\mathbb{E}(x_i) = 0$ and then recover the original intercept b_{00} after model fitting using a simple linear transformation (Goodrich et al., 2020; Bürkner, 2017). Common priors on b_{00} are a normal prior with mean $\mathbb{E}(y)$ and a user chosen scale depending on the scale of y as well as Jeffrey's prior which is improper flat in this case (Good, 1962).

What remains to be specified is a prior for the residual variance σ^2 (or equivalently residual standard deviation σ). We follow the recommendations of Gelman (2006) to consider a half Student- t prior on σ with ν degrees of freedom and scale η . We propose setting $\eta \approx \text{sd}(y)$, since both the prior expected mean and variance are proportional to η . To be able to implement a Gibbs sampling approach, an inverse Gamma distribution $\sigma^2 \sim \text{InvGamma}(c, d)$, with shape and scale parameters c, d respectively, could also be considered. We show how to implement a Gibbs sampler for the R2D2M2 prior in Appendix B.




FIG 3. Construction of the R2D2M2 prior schematically. We begin by assigning a distribution to R^2 , afterwards its uncertainty is propagated to the regression terms via a Dirichlet decomposition of the (a priori) explained variance. Finally, we make use of a distribution that can express the amount of variance that has been allocated to each term. Notice how all $u_{ig,l}$ share the same local variance λ_{ig} for a given pair (i, g_j) and varying level l .

The full R2D2M2 model can summarized as

$$\begin{aligned}
 y_n &\sim \text{Normal}(\mu_n, \sigma^2) \\
 \mu_n &= b_0 + \sum_{i=1}^p x_{ni} b_i + \sum_{g \in G_0} u_{0g_j[n]} + \sum_{i=1}^p x_{ni} \left(\sum_{g \in G_i} u_{ig_j[n]} \right) \\
 b_0 &\sim p(b_0) \\
 b_i &\sim \text{Normal} \left(0, \frac{\sigma^2}{\sigma_{x_i}^2} \phi_i \tau^2 \right), \quad u_{ig_j} \sim \text{Normal} \left(0, \frac{\sigma^2}{\sigma_{x_i}^2} \phi_{ig} \tau^2 \right) \\
 \tau^2 &= \frac{R^2}{1 - R^2} \\
 R^2 &\sim \text{Beta}(\mu_{R^2}, \varphi_{R^2}), \quad \phi \sim \text{Dirichlet}(\alpha), \quad \sigma \sim p(\sigma).
 \end{aligned} \tag{11}$$

A schematic construction of the R2D2M2 is shown in Figure 3. We do not specify distributions in this diagram, thus emphasizing that there is flexibility in the way uncertainty can be propagated from R^2 to the coefficients (as long as consistent rules that respect the corresponding domains of each parameter are followed). For example, the user might want to use a different distribution for R^2 , such as the Kumaraswamy distribution which offers advantages when using quantile-based approach to statistical modelling (Jones, 2009), or they might desire to model correlations among the components of ϕ explicitly, in which case they could use a distribution that allows it, such as a logistic normal distribution (Aitchison and Shen, 1980).

We close this section commenting on how the intuitiveness of the R2D2M2 prior assists in eliciting prior knowledge or lack of it. Users can readily communicate a prior expected value for R^2 through μ_{R^2} . Given a value for μ_{R^2} , the value selected for φ_{R^2} is an indicator of how confident the user is on how much of the response variability the model is explaining. This falls in line with the recommendations given by Mikkola et al. (2021) on how to improve prior elicitation: the hyperparameters on the prior for R^2 are easy to understand and

the priors for the regression terms regularized jointly (see Figure 3), hence significantly decreasing the number of hyperparameters to the user has to choose. Thus, the R2D2M2 prior brings us a step forward in terms of prior elicitation for multilevel models and provides a valuable tool in the Bayesian workflow for data analysis more generally (Gelman et al., 2020b).

3. Properties of the R2D2M2 prior

In this section, we derive several mathematical properties of the R2D2M2 prior. We begin by presenting the marginal prior distributions of the overall and varying coefficients. We show how the behavior around the origin and the tails of the marginal densities can be controlled by the hyperparameters $\mu_{R^2}, \varphi_{R^2}, a_\pi$. The distinct behaviors induced by hyperparameter choices can be exploited by the user to represent their prior knowledge and desired degree of regularization.

We then present the conditional posterior distributions of the overall and varying coefficients. Specifically, based on their conditional posterior means, we introduce shrinkage factors that allow us to calculate the effective number of coefficients and give further guidelines on how to choose the hyperparameters of the R2D2M2 prior.

3.1. Tail and concentration behaviors

The original R2D2 prior proposed by Zhang et al. (2020) uses a double exponential (Laplace) base distribution for the coefficients b_i . However, as mentioned before, in the multilevel model literature it is customary to treat b_i and in particular the varying coefficients u_{ig_j} as normally distributed (see Section 2 for details). The use of normal base distributions for the coefficients may lead to different marginal prior distributions, different concentration properties around the origin and tail behavior, which we will study in the following.

Consider the hyperparameters of the R2D2 prior μ_{R^2}, φ_{R^2} and let $a_1 = \mu_{R^2}\varphi_{R^2}, a_2 = (1 - \mu_{R^2})\varphi_{R^2}$ be the corresponding shape parameters of the Beta distribution on R^2 in the standard parameterization. In the following choose $\alpha = (a_\pi, \dots, a_\pi)'$ with $a_\pi > 0$. Zhang et al. (2020) showed that, under specific conditions depending on the hyperparameters a_π and a_2 , their prior is able to attain polynomial concentration rates around the origin and polynomial decay rates in the tails, thus allowing for shrinkage of noise towards zero and detection of signals. This shows that the R2D2 prior is able to compete with other well known GL shrinkage priors. We arrive to similar conclusions when using normal base distributions and taking multilevel structure into consideration. In the following propositions, the marginal distributions presented depend on a fixed value of σ . Other authors (Carvalho et al., 2010; Bai and Ghosh, 2019; Zhang et al., 2020) usually fix the value of $\sigma = 1$ for simplicity, however the marginal priors they present are still conditional on σ , as are ours.

Proposition 3.1. *Given the R2D2M2 prior, the marginal prior densities of b_i, u_{ig_j} given σ for any $i = 1, \dots, p, g \in \{1, \dots, G_i\}, j \in J_g$ are*

$$p(b_i|\sigma) = \frac{1}{\sqrt{2\pi q_i^2} B(a_\pi, a_2)} \Gamma(\eta) U(\eta, \nu, z_i) \quad (12a)$$

$$p(u_{ig_j}|\sigma) = \frac{1}{\sqrt{2\pi q_i^2} B(a_\pi, a_2)} \Gamma(\eta) U(\eta, \nu, z_{ig_j}), \quad (12b)$$

where $q_i^2 = \frac{\sigma^2}{\sigma_i^2}, a_2 = (1 - \mu_{R^2})\varphi_{R^2}, \eta = a_2 + 1/2$ and $\nu = 3/2 - a_\pi$. $B(\cdot, \cdot)$ and $\Gamma(\cdot)$ represent the Beta and Gamma functions respectively and $U(\eta, \nu, z_i)$ represents the confluent hypergeometric function of the second kind, as well as $z_i = \frac{|b_i|^2}{2q_i^2}$ and $z_{ig_j} = \frac{|u_{ig_j}|^2}{2q_i^2}$. If there are no varying coefficients, the marginal prior densities of b_i remain unaltered.

Proposition 3.2. *As $|b_i| \rightarrow 0, |u_{ig_j}| \rightarrow 0, 0 < a_\pi \leq 1/2$, and $a_2 > 0$, the marginal prior densities are unbounded with a singularity (i.e., undefined) at zero. Moreover, the marginal priors satisfy*

$$p(b_i|\sigma) \sim \begin{cases} c_1 b_i^{2a_\pi-1} + \mathcal{O}(|b_i|^{2a_\pi+1}), & a_\pi < 1/2 \\ -c_2 \ln(b_i^2/2q_i^2) + \mathcal{O}(b_i^2 \ln(b_i^2/2q_i^2)), & a_\pi = 1/2 \end{cases}$$

$$p(u_{ig_j}|\sigma) \sim \begin{cases} c_1 u_{ig_j}^{2a_\pi-1} + \mathcal{O}(|u_{ig_j}|^{2a_\pi+1}), & a_\pi < 1/2 \\ -c_2 \ln(u_{ig_j}^2/2q_i^2) + \mathcal{O}(u_{ig_j}^2 \ln(u_{ig_j}^2/2q_i^2)), & a_\pi = 1/2 \end{cases}$$

where $c_1 = (2q_i^2)^{1/2-a_\pi} \Gamma(1/2 - a_\pi) / \Gamma(a_2 + 1/2), c_2 = 1/\Gamma(a_2 + 1/2)$. When instead $a_\pi > 1/2$, the marginal prior densities are bounded and continuous at zero. The marginal prior densities are differentiable at $t = 0$ for all values of $a_\pi \geq 1$.

Proposition 3.2 implies that the value of a_π controls the boundedness of the marginal priors near the origin, which in turn influences the level of shrinkage that is enforced upon the coefficients. Sufficient mass near zero allows the prior to shrink small signals sufficiently (Van Erp et al., 2019). As a_π decreases from 1/2 to zero, the R2D2M2 prior shifts its mass even stronger to the origin and we can expect the posterior distribution to be concentrated near zero. This leads to prior concentration around sparse coefficient vectors and favors sparse estimation of the coefficients. On the other hand, an increasing a_π from 1/2 produces a bounded prior and indicates that we are in favor of less sparse coefficient vectors.

We can further understand the behavior of the marginal priors by considering that, when assuming $\phi \sim \text{Dirichlet}(\alpha)$ where $\alpha = (a_\pi, \dots, a_\pi)'$, the value of a_π determines where the density places its mass (Lin, 2016). When $0 < a_\pi < 1$ the density congregates at the edges of the simplex and prefers sparse distributions, i.e., most of the values within a single sample will be close to 0, and the vast majority of the mass will be concentrated in a few of the ϕ values. As a_π increases




FIG 4. Marginal prior densities of the coefficients with $\sigma = 1$. The value of a_π controls the behavior around the origin. If $a_\pi \leq 1/2$ the marginal density increases towards infinity for values close to zero and is unbounded and non-differentiable at the origin. In contrast, if $a_\pi > 1/2$ the marginal density is bounded also at the origin. When $a_\pi \geq 1$ the densities become differentiable at the origin.

to and exceeds 1, the density concentrates near the center of the simplex, with a mode appearing in $\left(\frac{1}{\dim(\alpha)}, \dots, \frac{1}{\dim(\alpha)}\right)'$, where $\dim(\alpha)$ denotes the length of α . Values of $a_\pi > 1$ prefer variates that are dense and evenly distributed, i.e., all the values within a single sample are similar to each other. This property is inherited by the R2D2M2 prior as Proposition 3.2 shows, since selecting $a_\pi \leq 1/2$ will attempt to concentrate the total variance into few regression terms, whereas $a_\pi > 1/2$ will distribute it evenly among the terms. Translation of prior beliefs over sparsity or lack thereof is thus possible by selection of the parameter a_π .

Figure 4 shows the marginal density of $p(b_i|\sigma)$ for different values of a_π with fixed $q_i^2 = 1$, the latter implying an equal scale of predictors and response variable without loss of generality. For the value of $a_\pi = 0.25$, the prior distribution shows high concentration of mass near the origin and is undefined (singular) at zero, which will lead to stronger shrinkage of weak signals towards the origin. On the other hand, when $a_\pi = 0.75$, the singularity is not present anymore and the prior becomes bounded at zero. A case of interest that presents bounded marginal priors is that of the (uniform) flat Dirichlet distribution over the simplex, that is when $a_\pi = 1$, meaning that the use of the flat Dirichlet prior over ϕ enforces less sparsity. Both types of behavior can be beneficial, since they can play in favor of the prior beliefs about sparsity the user might have.

The tails of the prior dictate whether or not it is possible that large signals can be detected (even when specific levels of sparsity are assumed) and how much these are shrunken towards zero (Carvalho et al., 2010; Bhattacharya et al., 2015; van der Pas, 2021). It is in principle desirable that heavy tails are present in order to prevent over-shrinkage of true signals (Carvalho et al., 2010; Piironen and Vehtari, 2017), which in turn can result in priors that are of bounded influence, i.e., sufficiently strong signals are left unshrunk by the prior. This behavior, known as tail robustness (Carvalho et al., 2010; van der Pas, 2021), is vital in sparse settings to be able to shrink coefficients near zero

much more forcefully than those far from it. However, we should keep in mind that in practice, this behavior does not always play in our favour, especially when parameters are weakly identified by the data. To circumvent this issue, authors such as Piironen and Vehtari (2017); Nishimura and Suchard (2022) have proposed *regularized* versions of shrinkage priors that also regularize very strong signals to some degree.

Our proposed R2D2M2 prior is able to attain heavier tails than the Cauchy distribution and, when no varying terms are present (the R2D2 prior with normal base distributions), it is of bounded influence, as is shown in the following propositions.

Proposition 3.3. *Given $|b_i| \rightarrow \infty, |u_{ig_j}| \rightarrow \infty$ for any $a_\pi > 0$ and $a_2 > 0$ the marginal prior densities satisfy*

$$p(b_i|\sigma) \sim \mathcal{O}(1/|b_i|^{2a_2+1}), \quad p(u_{ig_j}) \sim \mathcal{O}(1/|u_{ig_j}|^{2a_2+1}).$$

Moreover, when $0 < a_2 \leq 1/2$ the R2D2M2 marginal priors have heavier tails than the Cauchy distribution.

Proposition 3.4. *Assume there are no varying coefficients u_{ig_j} in the model, then the R2D2M2 prior is of bounded influence, i.e., sufficiently strong signals are not shrunken.*

While we were not able to prove that bounded influence also holds for the R2D2M2 prior, we have no reason to believe that it does not hold. Indeed, the results from the simulations presented in Section 4 show the bounded influence for the R2D2M2 prior empirically, at least for a range of different hyperparameter and data generating setups.

Propositions 3.2 and 3.3 show the adaptability of the R2D2M2 prior, in the sense that the user can tune the hyperparameters according to their prior beliefs on the degree of sparsity and signal strength. However, one should keep in mind that the complexity of the problem rapidly increases when considering multilevel structures and care should be taken when selecting the hyperparameters as to neither overfit (too little sparsity) nor over-shrink (too much sparsity).

3.2. Shrinkage factors

When performing Bayesian inference, the prior can be seen as a regularizer. This is clearly exemplified when using the Bayesian LASSO or other Bayesian shrinkage priors (Park and Casella, 2008; Van Erp et al., 2019), since the posterior modes under these priors correspond to certain frequentist regularizers (Bhattacharya et al., 2015). However, usual frequentist properties are not precisely equivalent; for instance, the posterior distribution for a continuous random variable is not able to produce exact sparse results, as the latter would require a point mass at zero. When using shrinkage priors, some authors have been able to express the posterior means as shrunken versions of (unshrunk) frequentist estimators (Piironen and Vehtari, 2017), which in practice leads to a better

out-of-sample predictive performance and a decrease in the variance of the estimators (Hoerl and Kennard, 1970). Several authors have used the concept of *shrinkage factors* to study the amount of shrinkage the prior induces and the effect the different values of the hyperparameters (Carvalho et al., 2010; Piironen and Vehtari, 2017; Bai and Ghosh, 2019; Polson and Scott, 2012). Here, we present the theory of shrinkage factors for the R2D2M2 prior as well as for the original R2D2 prior, since it has not been done before.

It can be shown that the conditional posterior distributions of the coefficients b_i and u_{ig_j} given $b_{-i}, u_{-ig_j}, \phi, \sigma, \tau, y$ in (11) are normal with means given by

$$\begin{aligned}\mathbb{E}(b_i | b_{-i}, u, \sigma, \phi, \tau^2, y) &= \left(\frac{\sigma_{x_i}^2}{\phi_i \tau^2} + SS_i \right)^{-1} \sum_{n=1}^N e_n^{(-i)} x_{ni} \\ \mathbb{E}(u_{ig_j} | b, u_{-ig_j}, \sigma, \phi, \tau^2, y) &= \left(\frac{\sigma_{x_i}^2}{\phi_{ig} \tau^2} + SS_{ig_j} \right)^{-1} \sum_{n \in L_{g_j}} e_n^{(-ig_j)} x_{ni},\end{aligned}\tag{13}$$

where b_{-i} and u_{-ig_j} represent the exclusion of the i th overall coefficient and the ig_j varying coefficient, respectively, $SS_i = \sum_n x_{ni}^2$ and $SS_{ig_j} = \sum_{n \in L_{g_j}} x_{ni}^2$. We use $e_n^{(-i)}$ and $e_n^{(-ig_j)}$ to represent the errors that would be obtained when excluding b_i and u_{ig_j} when trying to predict the n th observation.

$$\begin{aligned}e_n^{(-i)} &= y_n - \sum_{i' \neq i} x_{ni'} b_{i'} - \sum_{i=0}^p x_{ni} \left(\sum_{g \in G_i} u_{ig_{j[n]}} \right) \\ e_n^{(-ig_j)} &= y_n - \sum_{i=1}^p x_{ni} b_i - \sum_{i' \neq i} \sum_{g \in G_{i'}} x_{ni'} u_{i'g_{j[n]}} - x_{ni} \sum_{g \in G_i, g' \neq g} u_{ig'_{j[n]}}.\end{aligned}$$

If we assume there are no varying terms in the model, i.e., we are considering the original R2D2 prior, and $\sum_n x_{ni} x_{nj} = 0$ for $i \neq j$ then the conditional posterior mean of b_i is given by

$$\mathbb{E}(b_i | \sigma, \phi, \tau^2, y) = \left(\frac{\sigma_{x_i}^2}{\phi_i \tau^2} + SS_i \right)^{-1} SS_i \hat{b}_i,\tag{14}$$

where \hat{b}_i is the maximum likelihood estimate (MLE) of b_i . Thus, the posterior mean under the R2D2 prior is a shrunken version of the MLE, a common behavior that is encountered when regularizing (Hoerl and Kennard, 1970) or performing Bayesian inference with shrinkage priors such as the regularized horseshoe (Piironen and Vehtari, 2017). We are able to quantify the amount of shrinkage of b_i from their MLE towards zero by introducing shrinkage factors denoted as κ_i , where $0 \leq \kappa_i \leq 1$. Let

$$\kappa_i = \frac{1}{1 + \frac{SS_i}{\sigma_{x_i}^2} \phi_i \tau^2}\tag{15}$$

then Equation (14) becomes

$$\mathbb{E}(b_i|\sigma, \kappa_i, y) = (1 - \kappa_i)\hat{b}_i. \quad (16)$$

The shrinkage factor κ_i quantifies the amount of shrinkage that is exerted by the use of the R2D2 prior, relative to the non-regularized MLE. Shrinkage values close to 1 indicate full shrinkage from the MLE towards 0 and vice-versa, values close to 0 indicate no or only minor shrinkage.

When considering the presence of varying terms in the model, frequentist inference for b_i is typically performed by first integrating out u_{ig_j} (Bates et al., 2015). From a practical perspective, this is necessary as it is not possible to fully identify the overall and varying coefficients simultaneously when solving the optimization problem related to maximum likelihood. Thus, classical inference chooses to predict, rather than estimate, the varying terms u_{ig_j} after having obtained estimates \hat{b}_i of the overall coefficients b_i (Bates et al., 2015; Wakefield, 2013). In Bayesian inference on the other hand, both unknown quantities b_i and u_{ig_j} are considered random variables and are present in the full posterior distribution. Considering this, the generalization of shrinkage factors from an MLE when including varying terms is not immediate, because the joint MLE of the b_i and u_{ig_j} does not exist in this case. However, Equations (13) still show shrinkage is being carried out from some fixed quantities that involve the conditioned parameters. The approach to define shrinkage from quantities different than the MLE has been considered by other authors too (Bai and Ghosh, 2019; Polson and Scott, 2012), although not in a multilevel context. Doing so here opens up the way to define shrinkage factors κ_i and κ_{ig_j} for overall and varying coefficients respectively by

$$\kappa_i = \frac{1}{1 + r_i\phi_i\tau^2}, \quad \kappa_{ig_j} = \frac{1}{1 + r_{ig_j}\phi_{ig_j}\tau^2}, \quad (17)$$

where $r_i = \frac{SS_i}{\sigma_{x_i}^2}$ and $r_{ig_j} = \frac{SS_{ig_j}}{\sigma_{x_i}^2}$. Substituting in Equations (13) results in

$$\begin{aligned} \mathbb{E}(b_i|b_{-i}, \sigma, \kappa_i, y) &= (1 - \kappa_i)(SS_i)^{-1} \sum_{n=1}^N e_n^{(-i)} x_{ni} \\ \mathbb{E}(u_{ig_j}|b, u_{-ig_j}, \sigma, \kappa_{ig_j}, y) &= (1 - \kappa_{ig_j})(SS_{ig_j})^{-1} \sum_{n \in L_{g_j}} e_n^{(-ig_j)} x_{ni}. \end{aligned} \quad (18)$$

Shrinkage factors are useful in providing insight on the influence of the prior hyperparameters μ_{R^2} , φ_{R^2} , α , and a_π on the posterior distribution of each term. For instance, notice that in the extreme cases that α is chosen such that for a fixed i , $\phi_i \rightarrow 0$, then $\kappa_i \rightarrow 1$. If $\phi_i \rightarrow 1$ we have $\kappa_i \rightarrow (1 + r_i\tau^2)^{-1}$, so regularization takes place even if a priori it is believed that all the explainable variance is explained by a single component b_i . Additionally, by comparing κ_i and κ_{ig_j} in Equation (17) we can see that if $\phi_i \approx \phi_{ig_j}$ (which can be expected a-priori if $\alpha = a_\pi(1, \dots, 1)'$ when $a_\pi > 1$ (Lin, 2016)), then the varying coefficients will exhibit stronger shrinkage since $SS_{ig_j} \leq SS_i$ naturally, due to observations being

partitioned into different levels. It is also worth mentioning that the shrinkage phenomenon present in the posterior of u_{ig_j} is two-fold: (1) resulting from the usual shrinkage of varying coefficients that hierarchical models exhibit (i.e., partial pooling, Gelman and Hill (2006)) and (2) resulting from the local and global scales that are part of the R2D2M2 prior.

It is possible to find a closed expression for the densities of κ_i and κ_{ig_j} , as well as for their moments. Consider Equations (17) with ϕ fixed and τ^2 random. Using the fact that $\tau^2 \sim \text{BetaPrime}(\mu_{R^2}, \varphi_{R^2})$, the prior densities of the shrinkage factors, conditional on ϕ are given by

$$\begin{aligned} p(\kappa_i | \phi_i) &= \frac{(r_i \phi_i)^{a_2}}{\text{B}(a_1, a_2)} (1 - \kappa_i)^{a_1 - 1} \kappa_i^{a_2 - 1} ((r_i \phi_i - 1) \kappa_i + 1)^{-a_1 - a_2}, \\ p(\kappa_{ig_j} | \phi_{ig_j}) &= \frac{(r_{ig_j} \phi_{ig_j})^{a_2}}{\text{B}(a_1, a_2)} (1 - \kappa_{ig_j})^{a_1 - 1} \kappa_{ig_j}^{a_2 - 1} ((r_{ig_j} \phi_{ig_j} - 1) \kappa_{ig_j} + 1)^{-a_1 - a_2}. \end{aligned} \quad (19)$$

The m -th moments of $\kappa_i | \phi_i$ and $\kappa_{ig_j} | \phi_{ig_j}$ with $m \in \mathbb{Z}^+$ are given respectively by

$$\begin{aligned} \mathbb{E}(\kappa_i^m | \phi_i) &= \frac{(r_i \phi_i)^{a_2}}{\text{B}(a_1, a_2)} \text{B}(a_2 + m, a_1) {}_2F_1(\xi, \beta, \gamma, z_i), \\ \mathbb{E}(\kappa_{ig_j}^m | \phi_{ig_j}) &= \frac{(r_{ig_j} \phi_{ig_j})^{a_2}}{\text{B}(a_1, a_2)} \text{B}(a_2 + m, a_1) {}_2F_1(\xi, \beta, \gamma, z_{ig_j}), \end{aligned}$$

where ${}_2F_1(\xi, \beta, \gamma, z)$ is the hypergeometric function (Jeffrey et al., 2007) with $\xi = a_1 + a_2$, $\beta = a_2 + m$, $\gamma = a_1 + a_2 + m$, $z_i = 1 - r_i \phi_i$ and $z_{ig_j} = 1 - r_{ig_j} \phi_{ig_j}$. In the important case of $m = 1$, we have

$$\begin{aligned} \mathbb{E}(\kappa_i | \phi_i) &= (r_i \phi_i)^{a_2} (1 - \mu_{R^2}) {}_2F_1(\xi, \beta, \gamma, z_i), \\ \mathbb{E}(\kappa_{ig_j} | \phi_{ig_j}) &= (r_{ig_j} \phi_{ig_j})^{a_2} (1 - \mu_{R^2}) {}_2F_1(\xi, \beta, \gamma, z_{ig_j}). \end{aligned} \quad (20)$$

The distribution of the shrinkage factors κ depends directly on the hyperparameters μ_{R^2}, φ_{R^2} specified for the prior on R^2 and is sufficiently flexible to attain different forms that represent a diversity of a priori beliefs as we show in Figure 5. For example, the prior density of κ can become unbounded near 0 or 1 (or both at the same time), favoring limited or full shrinkage respectively. It is also possible to obtain bounded densities, which can represent mild shrinkage. These differences influence how noise and large signals are treated, which is reflected in the values of posterior means (Carvalho et al., 2010). Equations (20) show the effect of μ_{R^2} on the amount of shrinkage and how the prior is able to relate prior belief on R^2 with expected shrinkage on the coefficients. The expected shrinkage of each individual coefficient in the model is proportional to $1 - \mu_{R^2}$, implying that μ_{R^2} acts in a global manner. This is expected due to the relationship between the model's global scale τ^2 and R^2 established in Equation (9). Prior beliefs about R^2 are straightforwardly inherited to the amount of shrinkage the prior exhibits. If the user considers R^2 is low and represents this via a low value of μ_{R^2} , then a noticeable amount of shrinkage will take place. Analogously, a high value of μ_{R^2} results in less shrinkage.




FIG 5. Densities of shrinkage factors κ for several hyperparameter choices. The distribution offers flexibility to express different shrinkage behaviors. Notice that when $\mu = 0.5$ and $\varphi = 1$ the distribution is symmetric, categorizing the coefficient as either complete noise ($k \approx 1$) or as a strong signal ($k \approx 0$).

Figure 5 shows that the density of κ is able to reflect a horseshoe like prior for the amount of shrinkage when $(\mu_{R^2}, \varphi_{R^2}) = (0.5, 1)$. This type of prior can be considered fairly non-informative on the κ scale, since it places 1/3 of its mass on $1/4 \leq \kappa \leq 3/4$, with the rest of the density equally distributed to the left of 1/4 and to the right of 3/4. This shape indicates that we expect to see both relevant strong signals ($\kappa \approx 0$) and irrelevant variables a priori ($\kappa \approx 1$) (Carvalho et al., 2010; Piironen and Vehtari, 2017).

3.3. Effective number of non-zero coefficients

Piironen and Vehtari (2017) proposed that the prior’s hyperparameters can be understood intuitively by analyzing the imposed prior on the *effective number of coefficients*. In the case of single level models they define the effective number of overall coefficients as

$$m_{\text{eff}_O} = \sum_{i=1}^p (1 - \kappa_i). \quad (21)$$

We generalize Equation (21) to multilevel models (11) by also including the varying coefficients:

$$m_{\text{eff}} = \sum_{i=1}^p (1 - \kappa_i) + \sum_{i \in \{0, \dots, p\}} \sum_{g \in G_i} \sum_{j \in J_g} (1 - \kappa_{ig_j}). \quad (22)$$

When the shrinkage factors κ are close to 0 or 1, resulting in no shrinkage and total shrinkage respectively, Equation (22) can help us understand how specific values for μ_{R^2}, φ_{R^2} determine the amount of unshrunk (“active”) variables present in the model and can serve as an indicator of the effective model complexity.

For a given fixed dataset, we can visualize the prior imposed on m_{eff} for different hyperparameter values of μ_{R^2}, φ_{R^2} . This can be done by simulating

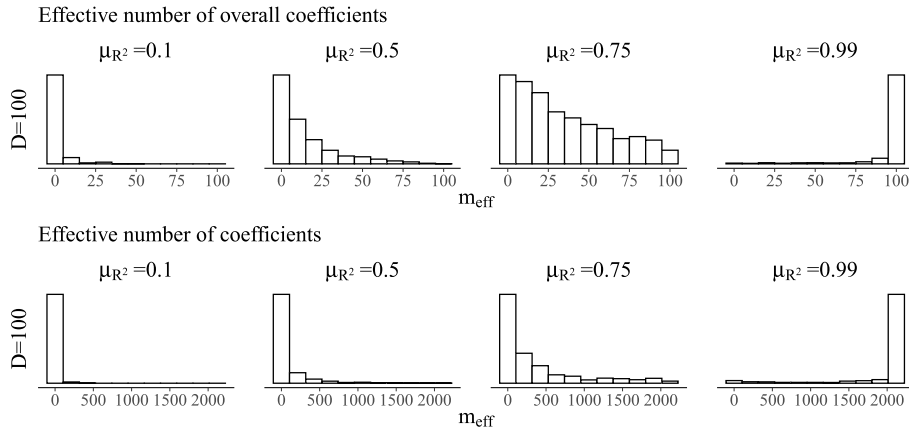


FIG 6. Prior densities for the effective number of coefficients under different conditions for μ_{R^2} with fixed $\varphi_{R^2} = 1$, $a_\pi = 0.5$ and $p = 100$ overall coefficients, one grouping factor and 20 levels. Decreasing μ_{R^2} implies less signal expected a priori, which results in increased a priori shrinkage of the coefficients. These plots serve as an intuitive way to understand the effect the hyperparameters have on the amount of shrinkage.

observations from m_{eff} directly. To do so, we first generate random variates τ, ϕ independently from their prior distributions, then we compute κ_i, κ_{ig_j} from Equations (17) and finally calculate m_{eff} . Figure 6 shows the histograms of the effective number of overall coefficients and total effective number of coefficients (which include the overall plus varying coefficients). In both cases, the behavior is as previously described; as μ_{R^2} increases the effective number of overall coefficients increases too and vice versa.

We recommend that the user visualizes the prior of m_{eff} for different values of μ_{R^2}, φ_{R^2} to get a better idea of how the global shrinkage is affected by the selection of the hyperparameters; and to verify if this choice represents their prior expectations on the number of non-zero coefficients. The prior plots of m_{eff} provide an intuitive way of understanding the induced overall shrinkage and can even become a valuable tool of communication between the user and non-specialized audiences.

4. Numerical experiments and case studies

We have implemented the R2D2M2 model in the probabilistic programming language Stan (Carpenter et al., 2017; Stan Development Team, 2022). The corresponding code can be found in Appendix B as well as on <https://osf.io/wgsth/>. We have also implemented the model in the R package `brms` (Bürkner, 2017), which provides a high level interface to fit Bayesian (non-)linear multivariate multilevel models using Stan. Based on the pure Stan implementation, we performed two large-scale simulations studies and analysed a high-dimensional real-world data set using the R2D2M2 prior as detailed below.

4.1. Simulation based calibration

First, we demonstrate that inference for our implementation of the R2D2M2 prior is probabilistically well calibrated. For this purpose, we applied simulation based calibration (SBC) (Talts et al., 2020). SBC is a general procedure for validating inferences from Bayesian algorithms capable of generating samples from posterior distributions. The procedure gives us the ability to identify inaccurate computations as well as inconsistencies in model implementations. It works by first sampling true parameter values $\tilde{\theta}$ from the prior, $\tilde{\theta} \sim p(\theta)$ and subsequently sampling data \tilde{y} from the likelihood $\tilde{y} \sim p(y|\tilde{\theta})$. Afterwards, using the algorithm that needs to be validated, we obtain S samples $\{\theta_1, \dots, \theta_S\} \sim p(\theta|\tilde{y})$ from the (approximated) posterior. Finally, for a quantity of interest $f(\theta)$, we calculate a rank statistic of how many values in $\{f(\theta_1), \dots, f(\theta_n)\}$ fall below the true value $f(\tilde{\theta})$. If the algorithm is well calibrated, this rank statistic of the prior sample $f(\tilde{\theta})$ in relation to the posterior draws is uniformly distributed (Talts et al., 2020). We will use the terms posterior draws and samples interchangeably.

To test uniformity of the rank statistics, we applied an intuitive graphical test proposed by Säilynoja et al. (2021), which provides simultaneous confidence bands of the empirical cumulative distribution function (ECDF) that satisfy a pre-specified joint coverage assuming uniformity. For ease of illustration, the plots show differences between the perfectly uniform CDF (a diagonal line) and the ECDF. In other words, it rotates a regular ECDF plot by 45 degrees to the right to make the uniform CDF a flat line and make the relevant part of the plot stand out more clearly. The graphical test is straightforward to use and is implemented in the R package `SBC` (Kim et al., 2022).

To sample from the posterior distribution we use Stan (Carpenter et al., 2017; Stan Development Team, 2022), a probabilistic programming language that provides the user with a (now substantially extended) implementation of the No-U-Turn Sampler (NUTS) from Hoffman and Gelman (2014), an adaptive form of Hamiltonian Monte Carlo (HMC) sampling (Brooks et al., 2011). To avoid sampling issues often encountered in hierarchical models we use non-centered parametrizations of the model (see <https://osf.io/wgsth/> and Gorinova et al. (2019) for details).

Calibration is a function of both the probabilistic model under consideration and the posterior approximation algorithm. Here we have tested calibration of our Stan implementations of both the R2D2 and the R2D2M2 prior models, since both of them are highly practically relevant. Note that calibration tests for previous implementations of the R2D2 prior model have not been reported beforehand by other authors.

There are two quantities that play a key role in how computationally intensive SBC can turn out to be: the number of simulation trials per configuration (i.e., the number of fitted models per configuration) and how many posterior draws to extract per simulation (i.e., per fitted model). We have fixed these to $T = 100$ and $S = 3000$, respectively. For MCMC samplers such as random walk Metropolis-Hastings or Gibbs sampling, this amount of posterior draws might seem as a bit small, however for adaptive Hamiltonian Monte Carlo estimates,

a few thousand samples is sufficient for most models (Bürkner, 2017; Hoffman and Gelman, 2014; Stan Development Team, 2022); and, as we show below, is indeed sufficient for the here-considered models.

We have tested calibration for a total of 96 different simulation configurations as detailed below. This amount comes from fully crossing all the factors that are varied in the simulations. An overview of the different configurations can be found in Appendix C in Table 5. The different conditions are chosen to represent several realistic scenarios that analysts might encounter in a similar fashion in real-life data analysis.

To sample prior values $\tilde{\theta}$ from the R2D2 or the R2D2M2 prior we need to specify several quantities: the prior mean μ_{R^2} and prior precision φ_{R^2} of the prior on R^2 as well as the concentration parameter a_π of the Dirichlet prior on ϕ . We chose $\mu_{R^2} \in \{0.1, 0.5\}$, $\varphi_{R^2} \in \{0.5, 1\}$, and $a_\pi \in \{0.5, 1\}$. The number of grouping factors K was chosen as $K \in \{0, 1\}$, implying either an R2D2 or an R2D2M2 prior. The number of covariates p was varied in $p \in \{10, 100, 300\}$. When considering the R2D2M2 prior, we used $q = p + 1$ varying terms per grouping factor K , that is, varying coefficients for each covariate plus a varying intercept. The number of levels per grouping factor was held constant with a value of $L = 20$. Finally, the residual standard deviation σ was sampled from a half Student- t with 3 degrees of freedom and scale of 1 (Gelman, 2006) and the overall intercept b_0 was sampled from a centered normal distribution with standard deviation of 5.

The covariates $x_i, i = 1, \dots, p$ were sampled independently of $\tilde{\theta}$ from a multivariate normal distribution centered at the origin with a covariance matrix Σ_x , formed from a correlation matrix ρ_x which has an autoregressive structure AR(1). The correlations considered here were $\rho \in \{0, 0.5\}$. Setting $\rho = 0$ results in $\Sigma_x = I_p$ such that predictor variables are sampled independently. Finally, the outcome data \tilde{y} was sampled as per Equation (1).

The obtained SBC results indicate that inference for our implementations of the R2D2 and R2D2M2 prior models is well calibrated in all considered configurations (see <https://osf.io/wgsth/> for the full results). For brevity, we only showcase a few selected results below. Specifically, we consider $(p, \mu_{R^2}, \varphi_{R^2}, a_\pi) = (100, 0.5, 1, 0.5)$ and $K \in \{0, 1\}$. The absence of grouping factors, $K = 0$, implies the use of the R2D2 prior.

Figures 7 and 8 show ECDF difference plots of several model parameters for the R2D2 and R2D2M2 models, respectively. For the R2D2 model, we show results for R^2 , σ , two overall coefficients and two components of ϕ chosen randomly. For the R2D2M2 model, we additionally show two varying coefficients chosen randomly too. Selection of the presented terms can be arbitrary, since the coefficients are exchangeable in the simulation process of all configurations. If ranks were perfectly uniform, the plots would display a horizontal line, but deviations from exact uniformity are to be expected because of the simulation process. All trajectories inside the shaded blue area indicate good calibration. These results provide evidence that inference for our implementation is correctly calibrated for the quantities of interest we have investigated (i.e., all model parameters). However, as Kim et al. (2022) mention, SBC provides only a necessary

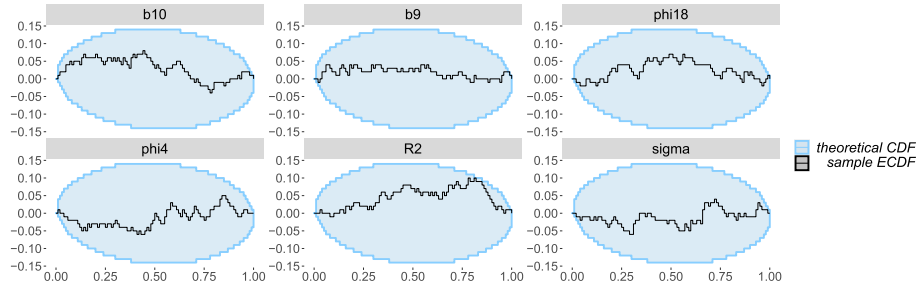


FIG 7. ECDF difference plots for the distributions of randomly chosen parameters in the R2D2 model. The blue areas in the ECDF difference plots indicate 95%-confidence intervals under the assumptions of uniformity and thus allow for a null-hypothesis significance test of self-consistent calibration. The plots show that proper calibration has been achieved.




FIG 8. ECDF difference plots for the distributions of randomly chosen parameters in the R2D2M2 model. The plots show that proper calibration has been achieved.

but not sufficient condition for correct calibration. It is possible the incorrect calibration still occurred for some pushforward quantities of the model, although unlikely given the good calibration of all model parameters (Talts et al., 2020).

4.2. Simulations from sparse multilevel models

4.2.1. Data generation of multilevel models

We conducted simulation studies to analyze the performance of the R2D2M2 prior under different scenarios. To generate multilevel data we follow the approach of Catalina et al. (2020). Their simulation approach is sufficiently general to encompass the different forms (sparse or non-sparse) multilevel data that might be encountered in practice. A graphical model that represents the data generating process is shown in Figure 9 (see Laura Dietz (2010) for more details on this type of diagrams). The different data generating conditions are summarized in Table 6 in Appendix C.

Again, we included all possible $q = p + 1$ varying terms per grouping factor $K \in \{1, 3\}$, with $p \in \{10, 100, 300\}$. The number of levels per grouping factor

remains constant with a value of $L = 20$. The covariates $x_i, i = 1, \dots, p$ are generated from a multivariate normal distribution centered at the origin with a covariance matrix Σ_x , formed from a correlation matrix ρ_x which has an autoregressive structure AR(1). The correlations considered were $\rho \in \{0, 0.5\}$. The overall intercept b_0 and the overall coefficients $b_i, i = 1, \dots, p$ are simulated independently from a normal distribution with mean zero and variances $\sigma_l^2 = 4, \sigma_b^2 = 9$, respectively. For a fixed covariate i , group g and level l , the varying coefficients u_{ig_j} are simulated from a normal distribution with mean zero and variance $\sigma_g^2 = 4$.

To induce sparsity in the data generating process we set each coefficient b_i, u_{ig_j} to zero with probabilities v, z respectively. If an overall coefficient is set to zero, then we also set the corresponding varying counterparts to zero (i.e if $b_i = 0$ then $u_{ig_j} = 0, \forall g, j$). The values for v were set to $\{0.5, 0.95\}$ and z was always set equal to v . The value of the residual standard deviation σ is adjusted in each simulation to maintain the prespecified value for the true proportion of explained variance $R_0^2 \in \{0.25, 0.75\}$, to do so we used Equations (5) and (6).

4.2.2. Prior hyperparameters

In the model, we set the R^2 prior hyperparameters to $\mu_{R^2} \in \{0.1, 0.5\}, \varphi_{R^2} \in \{0.5, 1\}$, which leads to $a_2 \in \{0.45, 0.25, 0.9, 0.5\}$. Values for which $a_2 \leq 1/2$ allow the marginal prior distributions of the regression terms to have heavier tails than the Cauchy distribution (see Proposition 3.3) and vice-versa. The hyperparameter α is set to $\alpha = (a_\pi, \dots, a_\pi)'$ with $a_\pi \in \{0.5, 1\}$ to test the different conditions near the origin, since values of $a_\pi \leq 1/2$ lead to unbounded marginal priors and vice-versa (see Proposition 3.2). Therefore, after fully crossing all these factors, we have a total of 8 different hyperparameter setups, which are summarized in Table 1. For brevity in the upcoming discussion of results, we will only refer to the indexes of these setups rather than to the specific hyperparameter values. The selected values of the hyperparameters were chosen to

TABLE 1
Configurations of the prior hyperparameters that were selected to test the R2D2M2 model under a sparse multilevel simulation setup.

Hyperparameter	1	2	3	4	5	6	7	8
a_π	0.5	1.0	0.5	1.0	0.5	1.0	0.5	1.0
μ_{R^2}	0.1	0.1	0.5	0.5	0.1	0.1	0.5	0.5
φ_{R^2}	0.5	0.5	0.5	0.5	1.0	1.0	1.0	1.0

create scenarios in which the marginal priors of the regression terms exhibit (un)boundedness near the origin and heavier (lighter) tails than the Cauchy distribution. For instance, consider setup (5) in Table 1, when $a_\pi = 0.5$ and $(\mu_{R^2}, \varphi_{R^2}) = (0.1, 1)$ (leading to $a_2 = 0.9$) we can expect that the prior will exert high shrinkage while presenting lighter tails. In this case, the prior is unbounded near the origin and there is major concentration of mass near the origin. In contrast, we can expect setup (4) to exert comparably little shrinkage, since the combination of boundedness ($a_\pi = 1$) at the origin and heavy

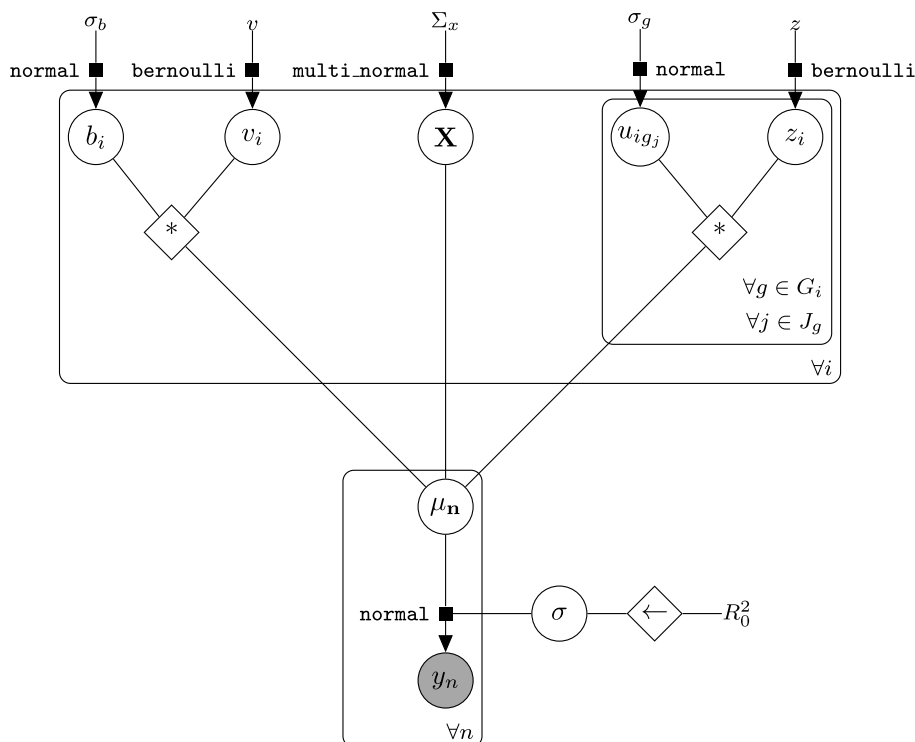


FIG 9. A graphical model showing the data generating process for the simulations. The $*$ symbol denotes the usual product and \leftarrow represents the use of Equations (5) and (6) to find the value of σ to maintain the pre-specified value of R_0^2 . The normal distributions considered for the coefficients b_i, u_{igj}, Σ_x have an expected value of 0. The mean of y_n , given by μ_n , is calculated using Equation (1).

tails ($a_2 = 0.25$) of the marginal priors will lead the R2D2M2 prior to shift its mass to the tails, thus focusing even more on detecting relevant signals. These different behaviors show the flexibility of the R2D2M2 prior. Similar analyses can be conducted for the other hyperparameter setups. The prior for the residual standard deviation σ was chosen as a half Student- t distribution with scale $\eta = \text{sd}(y)$ and $\nu = 3$ degrees of freedom, allowing for a finite variance but still with a heavy right tail.

The simulation setup provides an amount of 48 different data generation configurations, on which each of the different 8 hyperparameter setups will be tested. After fully crossing these conditions we have a total of 384 different simulation configurations. Table 6 in Appendix C summarizes the different conditions we have considered. For each configuration, we generated $T = 40$ datasets consisting of $N = 200$ training observations each. Out-of-sample predictive metrics were computed based on $N_{\text{test}} = 200$ test observations independent of the training data. The test set was designed to include all current grouping factors and

levels that are already in the training test, i.e., we are not performing prediction of new levels. To study the performance of the Bayesian models with the R2D2M2 prior, we have made use of several metrics. The complete list of metrics recorded for the experiments can be found in our OSF repository ¹. The primarily important ones are explained as they come up below.

4.2.3. Analysis of results

We show an analysis for three out of the total 48 different data generation configurations, whose results are representative of the overall results. They describe how the R2D2M2 prior performs in these scenarios with the different hyperparameter setups shown in Table 1. The complexity of the different scenarios considered increases in terms of number of coefficients included in the model. For the selected configurations $K = 1, 3, \rho = 0, v = z = 0.95, R_0^2 = 0.75$ and $p \in \{10, 100, 300\}$. For the different values of p , the total number of coefficients included in the linear predictor are $\{230, 2120, 6320\}$ when $K = 1$ and $\{670, 6160, 18360\}$ when $K = 3$, respectively. The $K = 1$ and $K = 3$ scenarios paint qualitatively the same picture. Of course, model performance in $K = 3$ scenarios was uniformly worse as it is considerably harder, but no new patterns (e.g., in terms of prior hyperparameter choice) emerged. For brevity, the results for $K = 3$ are only shown in Appendix C. The combination of $R_0^2 = 0.75$ and $p_i = 0.95$ implies that few (strong) signals account for most of the variability of y , a scenario which represents that only a handful of covariates have a substantial effect on the response. The R2D2M2 prior performs similarly for the other data generation setups and the results can be found on <https://osf.io/wgsth/>.

Estimation error We present summaries for several metrics in Table 2 considering $K = 1$. Similar results can be found when $K = 3$ in Table 7 in Appendix C. The Root Mean Square Error (RMSE) of parameter estimates (as represented by the parameter's posteriors) was calculated for all the parameters using the posterior mean as an estimator across the T replications. The RMSE is an indicator of estimation error and balances bias and variance. A low value is desirable with zero indicating the ideal case (perfectly precise and unbiased estimation). In this sense, it can be seen as criterion for proper parameter recovery. As p increases, the task of parameter recovery rises in complexity; however, RMSE values are low even in the most difficult scenarios, when $p = 300$.

Figure 10 shows the posterior shrinkage towards zero on the overall coefficients for the case in which $K = 1, p = 300$ when pooling over the different hyperparameter combinations and all simulations together. To show this, we compare the true value of the coefficient b with its posterior mean. The red diagonal line is the identity function and points that fall exactly on it indicate perfect posterior point estimation and no shrinkage. If the true b is positive (negative) and the corresponding dot is below (above) the diagonal line, then

¹<https://osf.io/wgsth/>

TABLE 2
Predictive Table results with $K = 1$.

Scenario $K = 1, \rho = 0, p_i = 0.95, R_0^2 = 0.75, N = 200, \alpha = 0.05$									
p	Metric	1	2	3	4	5	6	7	8
10	RMSE	0.02	0.03	0.02	0.02	0.02	0.02	0.01	0.02
	elpd	-434	-427	-441	-449	-430	-433	-430	-438
	m_{eff}	32	38	33	44	31	41	31	45
100	RMSE	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.01
	elpd	-1271	-1712	-1319	-1654	-1029	-1567	-1388	-1292
	m_{eff}	334	485	333	548	289	481	382	440
300	RMSE	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.02
	elpd	-3848	-3572	-4219	-3340	-3234	-3335	-3708	-3367
	m_{eff}	1412	1662	1560	1756	1111	1208	1330	1652

Results are shown for the Root Mean Squared Error (RMSE), expected log-pointwise predictive density (elpd) on the test set and the effective number of nonzero coefficients m_{eff} . The numbers 1-8 indicate the hyperparameter setup considered.

the posterior is shrinking the term towards zero. This ‘chair-like’ appearance is usually exhibited by shrinkage priors (see Carvalho et al., 2010; Van Erp et al., 2019, for more details) and is present under the different setups we have considered. This behavior indicates that the prior is exerting more shrinkage on weak signals and is consistent with the bounded influence and tail robustness properties the R2D2M2 prior possesses. Sufficiently strong signals are correctly detected and deemed as important, hence being shrunken less by the prior.

We see that, as the complexity of the problem increases (e.g., going from $K = 1$ to $K = 3$), the model responds by performing more shrinkage and by increasing the threshold a true coefficient needs to exceed to remain relatively unshrunk. This is especially clear for the $p = 300$ case since, as K increases, a flatter and wider region appears in the center of the chair-like plot, thus representing that the threshold for detection and bounded influence has increased and everything below it will be highly shrunken. This is an expected behavior of robust shrinkage priors (Carvalho et al., 2010). As complexity increases, they react by shrinking more and thus being more strict as to what they consider as a relevant signal.

Out-sample-predictive performance and shrinkage To assess predictive performance we estimate the expected log-pointwise predictive density (elpd) defined by Vehtari and Ojanen (2012) as

$$\text{elpd} = \sum_{i=1}^N \int p_t(\tilde{y}) \log p(\tilde{y}_i|y) d\tilde{y}_i,$$

where $p(\tilde{y}_i|y)$ is the posterior predictive distribution and $p_t(\tilde{y}_i)$ is the distribution representing the true data generating process. The elpd is as measure of predictive accuracy for the N data points taken one at a time and measures the goodness of the whole predictive distribution. We compute and report an estimator for elpd, denoted as $\widehat{\text{elpd}}$ on the test set to evaluate predictive performance. Denote the draws from the posterior distribution by $\theta^{(s)}$ where $s = 1, \dots, S$, then

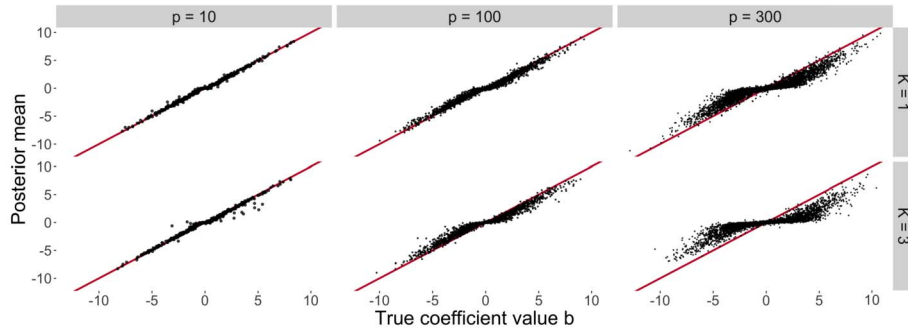


FIG 10. Posterior shrinkage of the overall coefficients pooled over the different hyperparameter setups and simulations. We show the true value versus the posterior mean. The red diagonal line is the identity function. Shrinkage is higher for weaker (in magnitude) signals resulting in “chair-like” plots. As the complexity increases, the prior increases the threshold above which (in absolute terms) shrinkage is noticeably weaker.

the elpd can be estimated by

$$\widehat{\text{elpd}} = \sum_{i=1}^N \log \left(\frac{1}{S} \sum_{s=1}^S p(y_i | \theta^{(s)}) \right).$$

We proceed to discuss how out-of-sample predictive performance (as measured via elpd) is intrinsically related to the amount of shrinkage done by the prior. We also discuss how prediction can be improved by imposing a higher amount of shrinkage via the hyperparameters. As a reference, the elpd values on the training set are usually close to -400 ± 50 for all cases, even as the number of covariates increases. Considering the test data data, Table 2 shows that in the comparably simple case of $p = 10$, all hyperparameter setups are providing similar point estimates for the elpd on test data and that it is similar to the elpd on training data. Figure 11 shows the distributions of the point estimator of elpd pooled across the different $T = 40$ trials arranged by setup and number of covariates along with 95% confidence intervals and the median. When $p = 10$, the distributions have similar behavior, indicating that under this relatively simple case (few predictor terms), we are able to properly generalize under the different hyperparameter setups.

When $p = 100$, Table 2 shows that there is a noticeable variability between the results of the point estimates of the elpd in the test set. Figure 11 shows that the distributions have different concentration around their median and in the width of their 95% confidence intervals. The fact that elpd decreases (while holding the number of test observations constant) is an indicator that the problem becomes more complex, to which the prior responds by increasing the amount of shrinkage it exerts.

Figure 12 shows the distributions of the posterior medians of m_{eff} across the different T trials arranged by hyperparameter setup and number of covariates, along with 95% confidence intervals and the median. Indeed, setup (5) pro-

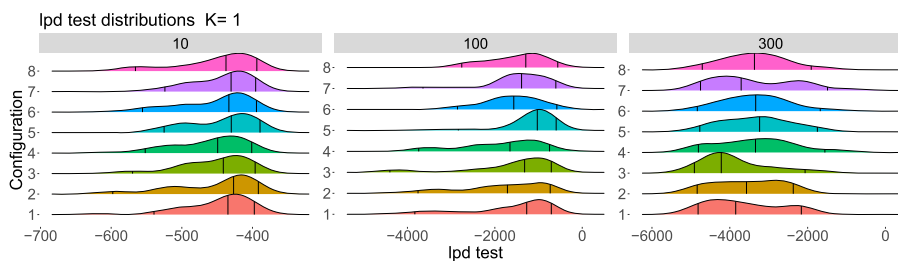


FIG 11. Estimated densities of the elpd estimators on the test set as p increases and arranged by hyperparameter setups when $K = 1$. The variation depicted is across $T = 40$ simulation trials. The vertical lines inside each density represent the 5%, 50%, 95% quantiles.

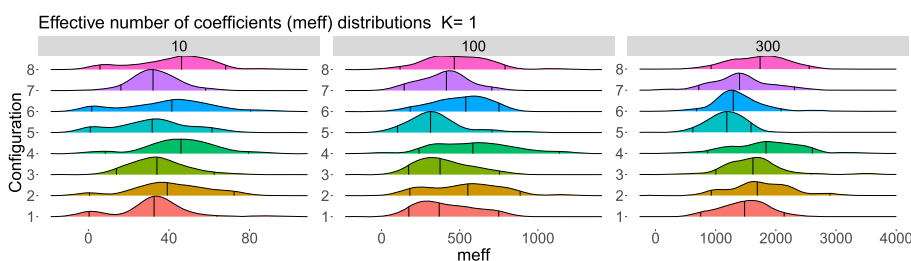


FIG 12. Estimated densities of the posterior median of the effective number of coefficients per hyperparameter setup and number of covariates when $K = 1$. The variation depicted is across $T = 40$ simulation trials. The vertical lines inside each density represent the 5%, 50%, 95% quantiles.

duces the most shrinkage, since it sacrifices detection of large signals by having a light tail and shifting mass near the origin, thus shrinking more. It is also expected that this hyperparameter setup performs well with respect to elpd since the current simulation setup has high levels of sparsity present (around 95% of the coefficients are truly zero). A somewhat better balanced condition, which presents a trade-off between predictive performance, detection of large signals and shrinkage, would be to consider conditions similar to setup (1) where $(\mu_{R^2}, \varphi_{R^2}) = (0.1, 0.5)$, which still allows for heavy tails. Figure 13 shows the relationship between m_{eff} and elpd when estimating it over the different setups for a fixed value of p across the different T trials. Similar results are obtained when averaging over different hyperparameter setups but are not shown here. Importantly, this behavior indicates that over-shrinking has not yet occurred in any of the different configurations, as we would see an increasing and then decreasing behavior of elpd as m_{eff} increases. Overall, a general trend emerges in these results: As m_{eff} decreases, elpd increases and out-of-sample predictive performance thus improves. Therefore, to improve the latter, the user can impose hyperparameters $(\mu_{R^2}, \varphi_{R^2})$ that serve the purpose of regularizing instead of necessarily reflecting prior knowledge directly (as the latter might suggest a higher R^2 than is sensible for regularization).




FIG 13. Relationship between $elpd$ and shrinkage by hyperparameter setup when $K = 1$. The non-linear relationships described by the colored lines in general shows a decreasing behavior as m_{eff} increases. Curves were estimated via thin-plate splines as implemented in the R package *mgcv* (Wood, 2011).

TABLE 3
Error Table when $K = 1$ and $\alpha = 0.05$.

		Scenario $K = 1, \rho = 0, p_i = 0.95, R_0^2 = 0.75, N = 200, \alpha = 0.05$							
p	Metric	1	2	3	4	5	6	7	8
10	Type I error	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)
	Type II error	0.24 (0.20)	0.24 (0.17)	0.15 (0.09)	0.17 (0.12)	0.23 (0.18)	0.28 (0.23)	0.09 (0.08)	0.25 (0.17)
	FDR	0.04 (0.04)	0.01 (0.01)	0.03 (0.03)	0.01 (0.02)	0.00 (0.01)	0.02 (0.05)	0.00 (0.00)	0.09 (0.07)
100	Type I error	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)
	Type II error	0.43 (0.26)	0.40 (0.29)	0.35 (0.22)	0.45 (0.27)	0.44 (0.24)	0.47 (0.27)	0.41 (0.22)	0.36 (0.23)
	FDR	0.01 (0.01)	0.02 (0.02)	0.01 (0.01)	0.08 (0.08)	0.01 (0.01)	0.05 (0.05)	0.04 (0.04)	0.01 (0.01)
300	Type I error	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)
	Type II error	0.72 (0.58)	0.68 (0.55)	0.73 (0.57)	0.73 (0.59)	0.72 (0.59)	0.75 (0.64)	0.69 (0.58)	0.76 (0.61)
	FDR	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)

We show the results for all the coefficients as well as results for the overall coefficients inside the parenthesis.

Credible interval coverage We evaluated the coverage properties of the R2D2M2 model by inspecting marginal credibility intervals of the regression coefficients of different size $1 - \alpha$. Table 3 reports Type I errors, Type II errors and False Discovery Rates (FDR) on the set of all coefficients regardless of the type and on the set of overall coefficients in parenthesis. To carry out the hypothesis $H_0 : \theta = \theta_0$ vs $H_1 : \theta \neq \theta_0$ at level α , we form the credibility interval of size $1 - \alpha$ denoted by I_α . We reject H_0 at level α if $\theta_0 \notin I_\alpha$. To estimate the Type I and Type II error we consider those coefficients that are truly zero and truly non-zero, respectively. Denote by V and R the total number of hypothesis that are rejected but are true (also known as “False Positives”) and the total number of hypothesis respectively (namely, “discoveries”) and let $Q = V/R$, then the FDR is defined as $FDR = \mathbb{E}(Q)$ (Benjamini and Hochberg, 1995). The FDR attempts to conceptualize Type I error in the multiple comparison setting and to identify as many discoveries, while incurring in a relatively low proportion of false positives. To estimate the FDR, we consider the truly zero coefficients and make use of the same test as before, while recording Q for each simulation. When interpreting these results, keep in mind that we should not necessarily expect correct frequentist coverage from our models, since they were not built to satisfy that goal.

For a pre-specified level of $\alpha = 0.05$, Table 3 shows that Type I errors are

controlled for (at level α) in both the set of all coefficients and the set of overall coefficients (presented in parenthesis). The FDR is also controlled at this level for both types of terms. On the other hand, Type II errors increase monotonically across the hyperparameter combinations with increasing number of overall coefficients in the model. This is expected since credibility intervals contain zero even for the majority of true non-zero coefficients due to the shrinkage the prior imposes when responding to the complexity of the problem at hand. What is more, Type II errors are worse for the set of all coefficients than for the set of overall, which is the result of marginal posterior credibility intervals of varying coefficients being wider than their overall counterparts. This is bound to happen, since we have less information about the varying coefficients than we do about the overall coefficients and, as mentioned before in Section 3.2, the former are thus subject to higher shrinkage.

The size of α determines the width of the credibility intervals and therefore has a direct influence on the error rates. Since the R2D2M2 prior produces Type I errors that are sufficiently low (and much lower than α itself), we can adjust α to decrease the Type II error, sacrificing Type I error control in exchange for statistical power. Type I errors are usually known as False Positive Rates (FPR). In this context, it is usual to study the complement of the Type II error known as the True Positive Rate (TPR). Likewise, modifying α will lead to a trade-off between the FPR and TPR and allows us to study ROC curves, which are usually presented when evaluating the efficiency of a classification algorithms (Fawcett, 2006).

Figure 14 shows how we can improve the TPR while increasing the FPR as we vary $\alpha \in \{0.02, 0.05, 0.10, 0.20, 0.50, 0.66\}$ when forming the credibility interval I_α and considering all of the coefficients. Likewise, the same behavior is observed when taking into account only overall coefficients in Figure 15. To illustrate how we can make better trade-offs between TPR and FPR, consider the most complex case when $p = 300$ and a value of $\alpha = 0.2$. This results in FPR ≈ 0.2 and TPR ≈ 0.8 for the set of overall coefficients. Figures 14 and 15 also indicate that our procedure is substantially better than random classification even in the most complex cases. This occurs even without performing correction for multiple comparisons and although the prior was not built for variable selection per se. Nonetheless, we recommend to separate inference and selection procedures as argued by Piironen et al. (2020). After performing inference via the R2D2M2 prior, projection prediction methods could be applied in a second step Piironen et al. (2020); Catalina et al. (2020), but studying this is out of scope of the present paper.

Table 4 reports coverage properties of 95% credibility intervals for the set of all coefficients and for the set of all coefficients. The properties related to the subset of non-zero coefficients are reported in parenthesis. Two main points stand out: 1) Under a sparse regime, coverage of non-zero coefficients is in general a harder problem and the proportion of coverage is less for non-zero coefficients than when considering the set of all coefficients at once. Due to the high number of truly zero coefficients (95% in this scenario) the prior attempts to adapt to the sparsity at the expense of the non-zero coefficients. 2) On average, credibil-

TABLE 4
Coverage table when $K = 1$.

		Scenario $K = 1, \rho = 0, p_i = 0.95, R_0^2 = 0.75, N = 200, \alpha = 0.05$								
p	Description	Metric	1	2	3	4	5	6	7	8
10	All	Coverage	0.99 (0.91)	0.99 (0.89)	0.99 (0.90)	0.99 (0.89)	0.99 (0.89)	0.99 (0.81)	0.99 (0.91)	0.99 (0.92)
		Width	0.86 (2.65)	1.02 (2.56)	0.96 (2.35)	1.03 (2.21)	0.85 (2.42)	0.95 (2.27)	0.91 (1.63)	0.98 (2.30)
	Overall	Coverage	0.97 (0.83)	0.98 (0.86)	0.98 (0.92)	0.98 (0.89)	0.98 (0.91)	0.96 (0.78)	0.98 (0.90)	0.96 (0.83)
		Width	0.44 (0.65)	0.52 (0.74)	0.50 (0.73)	0.52 (0.69)	0.44 (0.66)	0.50 (0.70)	0.47 (0.67)	0.50 (0.66)
100	All	Coverage	0.99 (0.71)	0.99 (0.68)	0.99 (0.66)	0.99 (0.62)	0.99 (0.68)	0.99 (0.65)	0.99 (0.70)	0.99 (0.67)
		Width	1.40 (2.32)	1.43 (1.98)	1.33 (2.25)	1.39 (1.94)	1.33 (2.35)	1.54 (2.18)	1.3 (2.22)	1.34 (1.99)
	Overall	Coverage	0.98 (0.81)	0.98 (0.80)	0.98 (0.76)	0.97 (0.77)	0.98 (0.82)	0.98 (0.8)	0.99 (0.81)	0.98 (0.72)
		Width	0.88 (1.35)	0.94 (1.36)	0.84 (1.35)	0.90 (1.33)	0.85 (1.40)	1.01 (1.45)	0.8 (1.26)	0.88 (1.33)
300	All	Coverage	0.99 (0.53)	0.99 (0.45)	0.99 (0.54)	0.99 (0.44)	0.99 (0.51)	0.99 (0.42)	0.99 (0.54)	0.99 (0.44)
		Width	1.75 (2.16)	1.66 (1.96)	1.79 (2.20)	1.68 (2.01)	1.73 (2.14)	1.78 (2.07)	1.67 (2.05)	1.76 (2.06)
	Overall	Coverage	0.98 (0.67)	0.97 (0.52)	0.98 (0.69)	0.97 (0.53)	0.98 (0.66)	0.96 (0.47)	0.98 (0.64)	0.97 (0.49)
		Width	1.41 (2.21)	1.38 (2.08)	1.44 (2.28)	1.39 (2.10)	1.4 (2.24)	1.49 (2.17)	1.34 (2.12)	1.47 (2.17)

Credibility intervals were formed at a $\alpha = 0.05$ level. We show results for the set of all coefficients mentioned in the description and for the non-zero corresponding cases inside parenthesis.

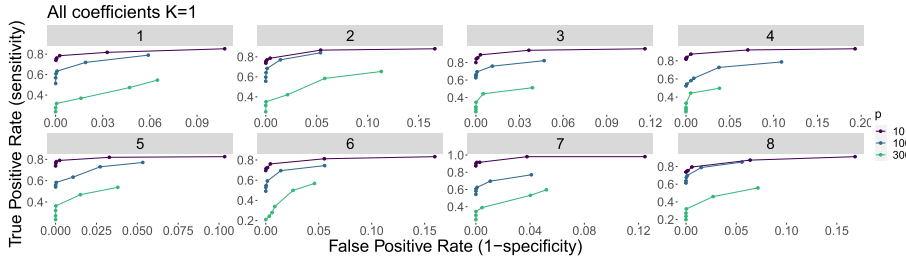


FIG 14. ROC curves for all coefficients arranged by hyperparameter setups when $K = 1$. Points are calculated by changing α when computing credibility intervals. All ROC curves shown are above the identity line.

ity intervals are wider for the non-zero coefficients. Since the prior adapts itself to the sparsity, it behaves conservatively when it detects a signal, performing shrinkage if the true value of the coefficient is not sufficiently high enough. Additionally, on average, the width of the credibility intervals increases with p (or with q for that matter) since there remains more epistemic uncertainty as model complexity increases given constant data set size.

4.3. Application to riboflavin production data

To illustrate the performance and usefulness of our developed models on real life challenges, we consider the `riboflavinGrouped` dataset provided by DSM (Kaiseraugst, Switzerland) and made publicly available by Buehlmann et al. (2014). This dataset has been used to study the associations between gene expressions and riboflavin (vitamin B_2) (Buehlmann et al., 2014; Lin et al., 2020). The response variable y to be predicted is the logarithm of the riboflavin production rate. The data contains the expression levels of $p = 4088$ candidate genes that might influence the riboflavin production. There are $N = 111$ observations arranged into $L = 28$ groups. Each group represents a strain (specimen)

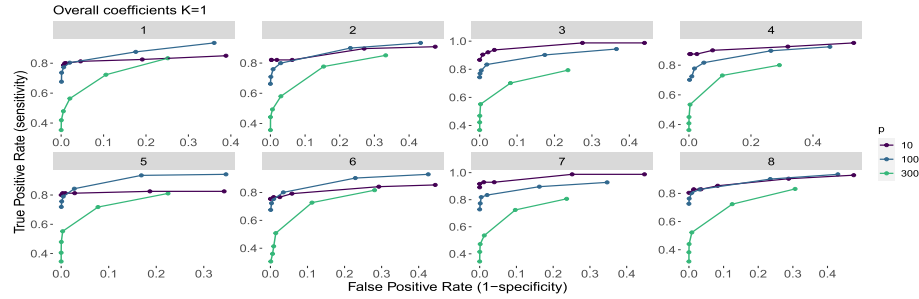


FIG 15. ROC curves for the overall coefficients when $K = 1$. Points are calculated by changing α when computing credibility intervals. All ROC curves shown are above the identity line.

of a genetically engineered *Bacillus subtilis* from which observations were taken. Different groups correspond to different strains of *Bacillus subtilis*, each having from 2 to 6 observations at different time points where the logarithm of the riboflavin production rate and the expression levels of the $p = 4088$ candidate genes were recorded.

The size of the data $N = 111$ is very small when compared to the number of candidate genes. Buehlmann et al. (2014) implemented a model with $p = 4088$ overall coefficients of which $q = 2$ were considered as varying among the different groups, however it is not specified in their work which are the ones they have considered to vary and why they have considered them. Lin et al. (2020) proposed modeling the association between the gene expressions and the production of riboflavin with a varying intercept model, i.e., $p = 4088$ overall terms and $q = 1$ varying term, however they do not take into account varying slopes nor models of different sizes. As an important result, both author groups reported the expression level of a specific gene, *YXLE-at*, as significantly associated with riboflavin production.

Our main goal in this case study is to test under which circumstances the R2D2M2 model is able to detect the gene *YXLE-at* as influential on the riboflavin production rate as well, even when considering much more complicated models than in earlier analysis of this dataset. To do so we built models of increasing size in terms of number of regression coefficients that will always include this main covariate of interest. A model with p overall coefficients is formed in the following way: we first include the expression level of gene *YXLE-at* as a covariate and then include the $p - 1$ covariates most correlated with *YXLE-at*; excluding those 5 covariates that are correlated with *YXLE-at* above $|r| = 0.98$, which makes them practically indistinguishable from the latter. Thus the maximal model in terms of overall coefficients is of size $p = 4083$. Once a covariate is included in the model, its varying coefficients counterpart is also included along with a varying intercept term. Including overall and varying coefficients the size of the model in terms of regression coefficients ranges from $P = 57$ when $p = 1$ to $P = 118406$ when $p = 4083$. The hyperparameters of the R2D2M2 prior are set to $(\mu_{R^2}, \varphi_{R^2}, a_\pi) = (0.1, 1, 0.25)$ to encode that we expect a fair amount of

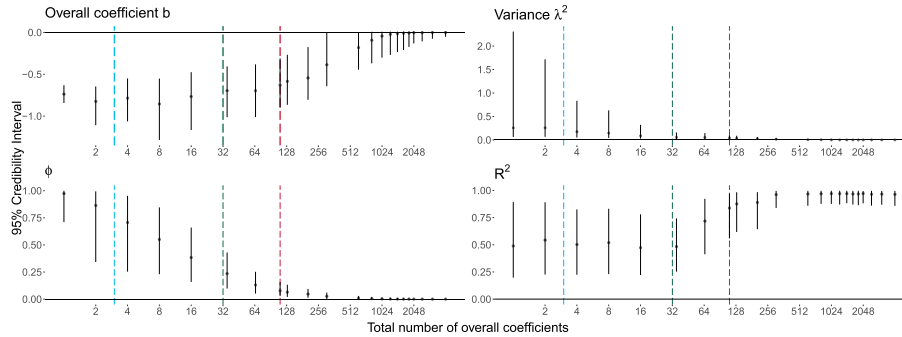


FIG 16. Plots of posterior means and 95% credibility intervals related to model size in terms of the total number of overall coefficients in the model. Results for the overall coefficient and corresponding variance component of the overall coefficient of the label *YXLE-at* are shown in the top left and top right plots respectively. The bottom left and bottom right plot show results for the proportion of allocated variance of *YXLE-at* and R^2 respectively. The blue, green and red vertical lines indicate the point where the total number of coefficients, total number of terms and total number of overall coefficient exceed the number of observations $N = 111$.

noise and to allow it to shrink most of the covariates to zero. This makes sense here, since only very few genes are expected to influence riboflavin production rate.

For the gene expression level of *YXLE-at*, Figure 16 shows plots of posterior means and 95% credibility intervals of the corresponding overall coefficient b , variance λ^2 of the overall coefficient, and allocation component ϕ (we drop the indices since we will only discuss the case of this gene) as well as R^2 related to the size of the model in terms of overall coefficients. The point in which the amount of observations is surpassed by the total number of coefficients $p + (q + 1)L$, number of terms $p + q + 1$ (i.e., number of components in ϕ) and number of covariates p is indicated by blue, green and red vertical lines, respectively. The gene *YXLE-at* is deemed as influential up to a certain point of p , when $p = 613$ to be precise, as the credibility intervals do not include zero; however, detection becomes more complicated as the size of the model increases, since the model responds to the complexity of the problem by shrinking coefficients stronger towards zero until this point, even the coefficients of a gene we now know to be truly relevant. This is in line with what we have shown in our simulations in Subsection 4.2 and demonstrates that the theoretically expected behavior of the prior can also be found when analysing real data. We reiterate here that shrinkage priors are not made to directly select variables but rather to provided sensible inference. If selection was really the goal, two step procedures would be recommended, of which fitting an R2D2M2 prior model would only be the first step (Piiroinen et al., 2020; Catalina et al., 2020).

5. Discussion

In this work, we have studied how joint regularization of all regression coefficients in Bayesian linear multilevel models can be carried out. To do so, we have proposed the R2D2M2 prior that imposes a prior on the global coefficient of determination R^2 and propagates the resulting regularization to the individual regression coefficients via a Dirichlet Decomposition. The chosen parameterization of the R2D2M2 prior offers intuitive and interpretable hyperparameters that are easy to understand in practice and offers analysts the ability to readily incorporate prior beliefs into their multilevel model. Together with the fact that only very few hyperparameters have to be specified, this greatly simplifies prior specification for these models.

We have also derived shrinkage factors and demonstrated how to compute the effective number of non-zero coefficients of the model. This allows the analyst to have a direct understanding of the implications that the chosen hyperparameters have on the amount of shrinkage imposed by the prior. We have shown, theoretically and via simulations, that the R2D2M2 prior enables both local and global shrinkage of the regression coefficients and that it possesses vital properties that are required when working on high-dimensional regression problems, such as sufficient concentration of mass near the origin as well as heavy tails. Finally we have demonstrated, via intensive simulations and by analysing real life data, that our prior implementation is not only well calibrated, but also offers reliable estimation even in the most complex cases that we investigated.

Previously, established priors for regression coefficients in the context of multilevel modelling have only been able to either jointly regularize the overall coefficients or the varying coefficients but never both at the same time. In the case of the latter, it has been done either in the standard way of using multivariate normal distribution with fixed hyperparameters (Gelman and Hill, 2006) or by constructing a joint prior over their variance parameters (Fuglstad et al., 2019). Additionally, dealing with high dimensionality in the multilevel context has been fairly restrictive, since it has been concerned with the scenario of a larger number of overall coefficients, i.e., $p > N$, while including only very few varying terms, i.e $q \ll N$. The R2D2M2 prior allows for a joint regularization of both the overall and varying coefficients and provides the opportunity to study high dimensional scenarios in which both $p > N$ and $q > N$.

Even though it is natural to wonder how the R2D2M2 prior model compares to frequentist methods to analyze multilevel models, we stress that we are focused on very high dimensional scenarios ($p, q > N$), where we need to regularize strongly due to the complexity of the problem. Existing frequentist methods and corresponding implementations for multilevel models that we are aware of are however not able to handle a high number of varying terms ($q > N$) and also struggle with a high number of overall terms ($p > N$) due to the resulting rank-deficient design matrices. Accordingly, comparing our proposed Bayesian model with frequentist methods would not be sensible at this point.

The present work can be extended in several ways. For instance, we could consider studying regularized versions of the R2D2M2 prior as Piironen and Ve-

htari (2017) did for the original Horseshoe prior (Carvalho et al., 2010). We have shown in simulations that as the complexity of the problem increases in terms of number of coefficients included the model, the prior reacts by increasing the amount of exerted shrinkage. However, coefficients that have a sufficiently high magnitudes will basically remain unaltered. Even though this is considered a key strength of shrinkage priors by some accounts (Van Erp et al., 2019), it may not always be desirable and can be harmful when the parameters are only weakly identified by the data (Piiironen and Vehtari, 2017). Alternatively, deriving joint priors for non-normal likelihoods is an important area for future research. For instance Gelman et al. (2020b) demonstrate for logistic regression that, as the number of covariates increases, quasi-complete separability becomes more likely, thus reducing the stability of the model if no substantial regularization is imposed.

Prior specification is a vital step in the Bayesian Workflow (Gelman et al., 2020b) and directly influences the performance of the models under consideration. Specifying joint priors allows us to account for the increasing complexity implied by increasing the amount of terms and parameters in the model, for example, in the shape of multilevel structure. What is more, they allow us to avoid some of the undesirable consequences of using independent, weakly informative priors in high dimensional setting. Ideally, joint priors should be able to express the user’s prior knowledge while at the same time being able to perform efficient regularization, both globally and locally. These ideas together with the new developments presented here lay out a new avenue of research in the field of prior specification and elicitation.

Appendix A: Proofs

Proof of Proposition 3.1

We provide a proof for the overall coefficients $b_i, i = 1, \dots, p$. The case of the varying coefficients u_{ig_j} is handled similarly. We begin by providing an alternative parametrization of the R2D2M2 prior which is used in the proofs of propositions 3.1, 3.2, and 3.3. The case of the varying coefficients is done in a similar way.

The prior for b_i is given by

$$\tau^2 \sim \text{BetaPrime}(a_1, a_2), \phi \sim \text{Dirichlet}(\alpha), b_i | \phi, \tau^2, \sigma^2 \sim \text{Normal}\left(0, \frac{\sigma^2}{\sigma_{x_i}^2} \phi_i \tau^2\right), \quad (23)$$

where $a_1 = \mu_{R^2} \varphi_{R^2}$, $a_2 = (1 - \mu_{R^2}) \varphi_{R^2}$. An observation from $\tau^2 \sim \text{BetaPrime}(a_1, a_2)$ can be simulated by the chain (Zhang et al. (2020))

$$\tau^2 | \xi \sim \text{Gamma}(a_1, \xi), \quad \xi \sim \text{Gamma}(a_2, 1),$$

where $\text{Gamma}(a, b)$ denotes a Gamma distribution with shape a and rate b . This provides an alternative representation of the R2D2M2 prior given by

$$\begin{aligned} \xi &\sim \text{Gamma}(a_2, 1), \tau^2 | \xi \sim \text{Gamma}(a_1, \xi), \phi \sim \text{Dirichlet}(\alpha) \\ b_i | \phi, \tau^2, \sigma^2 &\sim \text{Normal} \left(0, \frac{\sigma^2}{\sigma_{x_i}^2} \phi_i \tau^2 \right) \end{aligned} \quad (24)$$

Let α consist of a single repeating element a_π , where $a_\pi = \frac{a_1}{\dim(\alpha)} = \frac{\mu_{R^2} \varphi_{R^2}}{\dim(\alpha)}$. This is an automatic way of specifying a_π given that the user has provided a prior mean μ_{R^2} and a prior variance φ_{R^2} for R^2 . Then it follows that $\phi_i \tau^2 | \xi \sim \text{Gamma}(a_\pi, \xi)$. By definition, $\lambda_i^2 = \phi_i \tau^2$ and we can write the prior for b_i as

$$\xi \sim \text{Gamma}(a_2, 1), \lambda_i^2 | \xi \sim \text{Gamma}(a_\pi, \xi), b_i | \sigma^2, \lambda_i^2 \sim \text{Normal} \left(0, \frac{\sigma^2}{\sigma_{x_i}^2} \lambda_i^2 \right), \quad (25)$$

or as

$$\lambda_i^2 \sim \text{BetaPrime}(a_\pi, a_2), b_i | \sigma^2, \lambda_i^2 \sim \text{Normal} \left(0, \frac{\sigma^2}{\sigma_{x_i}^2} \lambda_i^2 \right).$$

Let $q_i^2 = \frac{\sigma^2}{\sigma_{x_i}^2}$, the prior marginal density of b_i is given by

$$\begin{aligned} p(b_i | \sigma) &= \int_0^\infty \frac{1}{\sqrt{2\pi q_i^2 \lambda_i^2}} \exp \left\{ -\frac{b_i^2}{2q_i^2 \lambda_i^2} \right\} \frac{1}{\text{B}(a_\pi, a_2)} (\lambda_i^2)^{a_\pi-1} (\lambda_i^2 + 1)^{-a_\pi-a_2} d\lambda_i^2 \\ &= \frac{1}{\sqrt{2\pi q_i^2} \text{B}(a_\pi, a_2)} \int_0^\infty \exp \left\{ -\frac{b_i^2}{2q_i^2} t_i \right\} t_i^{\eta-1} (t_i + 1)^{\nu-\eta-1} dt_i \\ &= \frac{1}{\sqrt{2\pi q_i^2} \text{B}(a_\pi, a_2)} \Gamma(\eta) U(\eta, \nu, z_i), \end{aligned}$$

where $t_i = \frac{1}{\lambda_i^2}$, $\eta = a_2 + 1/2$ and $\nu = 3/2 - a_\pi$. $U(\eta, \nu, z_i)$ represents the confluent hypergeometric function of the second kind (see Jeffrey et al. (2007)) and $z_i = \frac{|b_i|^2}{2q_i^2}$. $U(\eta, \nu, z_i)$ is defined as long as the real part of η and z_i are positive, which is always the case since $a_2 > 0$, $q_i^2 > 0$ and $|b_i| > 0$.

Proof of Proposition 3.2

The confluent hypergeometric function of the second kind $U(\eta, \nu, z)$ satisfies the following equation (see equation 13.2.40 Olver et al. (2010))

$$U(\eta, \nu, z) = z^{1-\nu} U(\eta - \nu + 1, 2 - \nu, z). \quad (26)$$

We drop the index in z since all cases are handled equally. Substituting $\eta = a_2 + 1/2$, $\nu = 3/2 - a_\pi$ and $z = \frac{t^2}{2q^2}$ results in

$$U \left(a_2 + 1/2, 3/2 - a_\pi, \frac{t^2}{2q^2} \right) = (2q^2)^{1/2-a_\pi} t^{2a_\pi-1} U \left(a_2 + a_\pi, a_\pi + 1/2, \frac{t^2}{2q^2} \right).$$

The right hand side has a singularity at $t = 0$ when $a_\pi < 1/2$.

The limiting form U has as $|z| \rightarrow 0$ follows equation 13.2.18 in Olver et al. (2010) given by

$$U(\eta, \nu, z) = \frac{\Gamma(\nu - 1)}{\Gamma(\eta)} z^{1-\nu} + \frac{\Gamma(1 - \nu)}{\Gamma(\eta - \nu + 1)} + \mathcal{O}(z^{2-\nu}), \quad 1 < \nu < 2.$$

Substituting $\eta = a_2 + 1/2, \nu = 3/2 - a_\pi$ and $z = \frac{t^2}{2q^2}$ results in

$$U\left(a_2 + 1/2, 3/2 - a_\pi, \frac{t^2}{2q^2}\right) = Ct^{2a_\pi - 1} + \frac{\Gamma(a_\pi + 1/2)}{\Gamma(a_2 + 1/2)} + \mathcal{O}(|t|^{1+2a_\pi}), \quad 0 < a_\pi < 1/2,$$

where $C = (2q^2)^{1/2 - a_\pi} \frac{\Gamma(1/2 - a_\pi)}{\Gamma(a_2 + 1/2)}$. The last expression is unbounded at $t = 0$ when $a_\pi < 1/2$. The case when $a_\pi = 1/2$ follows equation 13.2.19 given by

$$U(\eta, \nu, z) = -\frac{1}{\Gamma(\nu)} (\ln(z) + \psi(\eta) + 2\gamma) + \mathcal{O}(z \ln(z)), \quad 1 < \nu < 2,$$

where $\psi(\cdot), \gamma$ represent the digamma function and the Euler-Mascheroni constant respectively. Substituting $\eta = a_2 + 1/2, \nu = 3/2 - a_\pi$ and $z = \frac{t^2}{2q^2}$ gives

$$U\left(a_2 + 1/2, 3/2 - a_\pi, \frac{t^2}{2q^2}\right) = -\frac{1}{\Gamma(a_2 + 1/2)} \left(\ln\left(\frac{t^2}{2q^2}\right) + \psi(a_2 + 1/2) + 2\gamma \right) + \mathcal{O}\left(t^2 \ln\left(\frac{t^2}{2q^2}\right)\right).$$

The last expression is not defined at $t = 0$ and as $|t| \rightarrow 0$ it goes to infinity. To see the marginal priors are bounded when $a_\pi > 1/2$ it is sufficient to consider the cases given by equations 13.2.20-13.2.22 in Olver et al. (2010) and making the proper substitutions of $\eta = a_2 + 1/2, \nu = 3/2 - a_\pi$ and $z = \frac{t^2}{2q^2}$. In terms of the values a_π takes we have

$$U(\eta, \nu, t) \sim \begin{cases} \mathcal{O}(t^{2a_\pi - 1}), & 1/2 < a_\pi < 3/2 \\ \mathcal{O}(t^2 \ln(t^2/2)), & a_\pi = 3/2 \\ \mathcal{O}(t^2), & a_\pi > 3/2 \end{cases}$$

The derivative of $U(\eta, \nu, z)$ is given by equation 13.3.22 in Olver et al. (2010) as

$$\frac{d}{dz} U(\eta, \nu, z) = -\eta U(\eta + 1, \nu + 1, z).$$

Combining this with Equation (26) we have

$$\frac{d}{dz} U(\eta, \nu, z) = -\eta z^{-\nu} U(\eta - \nu + 1, 1 - \nu, z).$$

Substituting $\eta = a_2 + 1/2, \nu = 3/2 - a_\pi$ and $z = \frac{t^2}{2q^2}$ results in

$$\frac{d}{dt} U\left(a_2 + 1/2, 3/2 - a_\pi, \frac{t^2}{2q^2}\right) = Ct^{2a_\pi - 2} U\left(a_2 + a_\pi, a_\pi - 1/2, \frac{t^2}{2q^2}\right), \quad (27)$$

where $C = -(a_2 + 1/2)2^{3/2-a_\pi}(q^2)^{1/2-a_\pi}$. Equation (27) is undefined at $t = 0$ when $a_\pi < 1$.

Making $t = b_i$ or $t = u_{ig_j}$ proves the proposition.

Proof of Proposition 3.3

To prove the proposition we make use of Watson's lemma as found in Miller (2006).

Lemma A.1 (Watson's lemma). *Let $0 \leq T \leq \infty$ be fixed. Assume $f(t) = t^\lambda g(t)$, where $g(t)$ has an infinite number of derivatives in the neighborhood of $t = 0$, with $g(0) \neq 0$, and $\lambda > -1$. Suppose, in addition, either that $|f(t)| < Ke^{ct}$ for any $t > 0$, where K and c are independent of t . Then it is true that for all positive x that*

$$\left| \int_0^T e^{-xt} f(t) dt \right| < \infty$$

and that the following asymptotic equivalence holds:

$$\int_0^T e^{-xt} f(t) dt \sim \sum_{n=0}^{\infty} \frac{g^{(n)}(0)\Gamma(\lambda + n + 1)}{n!x^{\lambda+n+1}},$$

for $x > 0$ as $x \rightarrow \infty$.

The marginal distribution of an overall coefficient b is given by

$$\begin{aligned} p(b|\sigma) &= \frac{1}{\sqrt{2\pi q^2}B(a_\pi, a_2)} \int_0^\infty \exp\left\{-\frac{|b|^2}{2q^2}t\right\} t^{\eta-1}(t+1)^{\nu-\eta-1} dt \\ &= \int_0^\infty \exp\{-zt\} f(t) dt, \end{aligned}$$

where $z = \frac{|b|^2}{2q^2}$, $f(t) = C t^{\eta-1}(t+1)^{\nu-\eta-1} = t^{\eta-1}g(t)$, $C = \left(\sqrt{2\pi q^2}B(a_\pi, a_2)\right)^{-1}$, and $g(t) = C(t+1)^{\nu-\eta-1}$. If we make $\lambda = \eta - 1$, the hypothesis $\lambda > -1$ is satisfied since $a_2 - 1/2 > -1$ for $a_2 > 0$. $g(t)$ is infinitely differentiable around $t = 0$ and $g(0) = 0$. By Watson's Lemma, since $|f(t)| < Ke^{ct}$ for all $t > 0$ where K and c are independent of t , then as $|b| \rightarrow \infty$,

$$p(z|\sigma) = \sum_{n=0}^{\infty} \frac{g^{(n)}(0)\Gamma(\lambda + n + 1)}{n!z^{\lambda+n+1}}.$$

Truncating the sum at $n = 2$ gives

$$\begin{aligned} p(z|\sigma) &= C \left\{ \frac{\Gamma(a_2+1/2)}{z^{a_2+1/2}} - \frac{(a_\pi+a_2)\Gamma(a_2+3/2)}{z^{a_2+3/2}} + \frac{(a_\pi+a_2)(a_\pi+a_2+1)\Gamma(a_2+5/2)}{z^{a_2+5/2}} \right\} \end{aligned}$$

$$\begin{aligned}
 & + \mathcal{O}\left(\frac{1}{z^{a_2+7/2}}\right) \\
 & \sim \mathcal{O}\left(\frac{1}{z^{a_2+1/2}}\right).
 \end{aligned}$$

Therefore $p(|b| | \sigma) \sim \mathcal{O}\left(\frac{1}{|b|^{2a_2+1}}\right)$. When $a_2 < 1/2$ as $|b| \rightarrow \infty$ and comparing with $\frac{1}{b^2}$ we have that

$$\frac{p(b|\sigma)}{\frac{1}{b^2}} \sim \mathcal{O}\left(\frac{1}{b^{2a_2-1}}\right) \rightarrow \infty.$$

Proof of Proposition 3.4

We will first show that the R2D2M2 prior can be represented as Horseshoe type prior (Carvalho et al., 2010). We then proceed to make use of Theorem 2 and Theorem 3 in Carvalho et al. (2010) to show that the R2D2M2 prior with normal base distributions for the coefficients is of bounded influence.

Consider representation (25) of the R2D2M2 prior given by

$$b_i | \sigma, \lambda_i^2 \sim \text{Normal}(0, \sigma^2 \lambda_i^2), \quad \lambda_i^2 | \xi \sim \text{Gamma}(a_\pi, \xi), \quad \xi \sim \text{Gamma}(a_2, 1).$$

It can be shown that this representation is equivalent to

$$b_i | \sigma^2, \lambda_i^2 \sim \text{Normal}(0, \sigma^2 \lambda_i^2), \quad \lambda_i^2 \sim \text{BetaPrime}(a_\pi, a_2).$$

When $a_\pi = 1/2$ and $a_2 = 1/2$, then $\lambda_i \sim \text{Cauchy}^+(0, 1)$ and we have

$$b_i | \sigma, \lambda_i^2 \sim \text{Normal}(0, \sigma^2 \lambda_i^2), \quad \lambda_i \sim \text{Cauchy}^+(0, 1).$$

The Horseshoe prior (Carvalho et al., 2009) has the following representation:

$$\begin{aligned}
 b_i | \lambda_i, w & \sim \text{Normal}(0, w^2 \lambda_i^2) \\
 \lambda_i & \sim \text{Cauchy}^+(0, 1),
 \end{aligned}$$

where λ_i are the local shrinkage parameters and w is the global shrinkage parameter. Therefore the R2D2M2 prior can be represented as a Horseshoe type prior where the global scale is fixed and $w = 1$. In the following consider $\sigma = 1$ and $y|b \sim \text{Normal}(b, 1)$, since by hypothesis there are no varying coefficients in the model. Theorem 2 in Carvalho et al. (2010) shows that conditioned on one sample y^* , we can write

$$\mathbb{E}(b|y^*) = y^* + \frac{d}{dy^*} \log m(y^*), \tag{28}$$

where $m(y^*)$ is the marginal density for y^* given by $m(y^*) = \int p(y^*|b)p(b)db$. Theorem 3 of Carvalho et al. (2010) shows that for the Horseshoe prior as $|y^*| \rightarrow \infty$ then

$$\lim_{|y^*| \rightarrow \infty} \frac{d}{dy^*} \log m(y^*) = 0,$$

implying that as $|y^*| \rightarrow \infty$ then $\mathbb{E}(b|y^*) \approx y^*$ and thus showing that the R2D2M2 prior is of bounded influence. The exact form of $m(y^*)$ for the R2D2 prior is given by

$$m(y^*) = \frac{1}{(2\pi^3)^{1/2}} \int_0^\infty \exp\left(-\frac{y^{*2}/2}{1+\lambda^2}\right) \frac{1}{(1+\lambda^2)^{3/2}} d\lambda.$$

Making $z = \frac{1}{\lambda^2+1}$ results in

$$\begin{aligned} m(y^*) &= \frac{1}{(2\pi^3)^{1/2}} \int_0^1 \exp(-1/2y^{*2}z) z^{-1/2} dz \\ &= \frac{1}{\pi} \frac{\operatorname{erf}(y^*/\sqrt{2})}{y^*}, \end{aligned}$$

where $\operatorname{erf}(\cdot)$ denotes the error function given by $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$. Hence, $\frac{d}{dy^*} \log m(y^*)$ is given by

$$\frac{d}{dy^*} \log m(y^*) = \frac{\sqrt{\frac{2}{\pi}} e^{-y^{*2}/2}}{\operatorname{erf}(y^*/\sqrt{2})} - \frac{1}{y^*},$$

and $\lim_{|y^*| \rightarrow \infty} \frac{d}{dy^*} \log m(y^*) = 0$. Using this and (28) we have that as $|y^*| \rightarrow \infty$ then $\mathbb{E}(b|y^*) \rightarrow y^*$.

Appendix B: Posterior computation

B.1. Stan

To obtain draws from the posterior distribution we have implemented our model in Stan (Carpenter et al., 2017; Stan Development Team, 2022), which uses a substantially extended implementation of the No-U-Turn Sampler (NUTS) from Hoffman and Gelman (2014). This sampler is an adaptive form of Hamiltonian Monte Carlo (Brooks et al., 2011). The following code is fully functional however we also maintain a current version in <https://osf.io/wgsth/>.

This implementation considers K grouping factors with multiple levels L each and D overall coefficients, including the overall intercept. All groups have the same number of levels L_g as well as the same number of varying coefficients D_g . D_g is the number of group level effects and is including the varying intercepts as well, i.e Group K level l has D_g varying coefficients. This implementation is able to work with data where D, K, L_g can vary with no need to manually modify the Stan code. However, the code can be easily modified to handle the case when groups have different amount of levels and terms per group. It is also possible to briefly modify this code in order to work with $q \leq p$ varying terms. The order of the varying terms should consider the relationship with the i th covariate since we scale by σ_{x_i} .

```

// Stan code for the R2D2M2 prior.
functions {

  vector R2D2(vector z, vector sds_X, vector phi, real tau2) {
    /* Efficient computation of the R2D2 prior
    * Args:
    *   z: standardized population-level coefficients
    *   phi: local weight parameters
    *   tau2: global scale parameter (sigma is inside tau2)
    * Returns:
    *   population-level coefficients following the R2D2 prior
    */
    return z .* sqrt(phi * tau2) ./ sds_X ;
  }
}

data {
  int<lower=1> N; // total number of observations
  vector[N] Y; // response variable
  int<lower=1> D; // number of population-level effects
  matrix[N, D] X; // population-level design matrix
                  // including column of 1s
  int<lower=0> K; // number of groups
  vector[D-1] sds_X; // column sd of X before centering.
                  //Estimate or real values.

  //---- data for group-level effects

  int<lower=1> Lg; // number of levels per group (constant)
  int<lower=1> Dg; // number of coefficients per level per group
                  //(D_g constant per group)
  int<lower=1> J[N,K]; // grouping indicator matrix
                  // per observation per group K

  //---- group-level predictor values
  matrix[Dg,N] Z[K];

  //---- data for the R2D2 prior
  vector<lower=0>[(D-1)+K+(Dg-1)*K] R2D2_alpha;
  real<lower=0> R2D2_mean_R2; // mean of the R2 prior
  real<lower=0> R2D2_prec_R2; // precision of the R2 prior
  int prior_only; // should the likelihood be ignored?
}

transformed data {
  int Dc = D - 1; //
  matrix[N, Dc] Xc; //centered version of X without an intercept
  vector[Dc] means_X; //column means of X before centering
  vector[Dc] var_X;

```

```

vector[N] Yc;
real Ymean;
for (i in 2:D) {
  means_X[i - 1] = mean(X[, i]);
  var_X[i-1]= sds_X[i-1]^2;
  Xc[, i - 1] = (X[, i] - means_X[i - 1]) ;
}

Ymean= mean(Y);
for (i in 1:N) {
  Yc[i]= Y[i]-mean(Y);
}
}

parameters {
  real Intercept; // temporary intercept for centered predictors
  vector[Dc] zb; // standardized population-level effects
  matrix[Dg,Lg] z[K]; // standardized group-level effects
  real<lower=0> sigma; // standard deviation of response

  // local parameters for the R2D2M2 prior
  simplex[Dc+K+(Dg-1)*K] R2D2_phi;
  // R2D2 shrinkage parameters
  /*Convention of indexing: First Dc for overall effects,
    group of K for varying intercepts,
    Batches of Dc for each grouping factor */

  real<lower=0,upper=1> R2D2_R2; // R2 parameter
}

transformed parameters {

  vector[Dc] b; // population-level effects

  matrix[Dg,Lg] r[K]; // actual group-level effects

  real R2D2_tau2; // global R2D2 scale parameter
  R2D2_tau2 = R2D2_R2 / (1 - R2D2_R2);

  // compute actual overall regression coefficients
  b = R2D2(zb, sds_X, R2D2_phi[1:Dc], (sigma^2) * R2D2_tau2);

  for(k in 1:K){
    // varying intercepts
    // Dc+k is the kth varying intercept
    r[k,1,] = (sigma * sqrt(R2D2_tau2 * R2D2_phi[Dc+k ]
      * (z[k,1,])));
    for(d in 2: Dg){

```

```

// group level effects
// (k-1)Dc indexes the beginning of the kth batch of scales
r[k,d,]= sigma /(sds_X[(d-1)]) *
      sqrt(R2D2_tau2 *
           R2D2_phi[Dc+K+ (k-1)*(Dg-1) +(d-1) ]) * (z[k,d,]);
}
}
}

model {
  // likelihood including constants

  if (!prior_only) {
    // initialize linear predictor term
    vector[N] mu = Intercept + rep_vector(0.0, N);
    for (n in 1:N) {
      // add more terms to the linear predictor

      for(k in 1:K){
        mu[n]+=dot_product(r[k,.,J[n,k]], Z[k,.,n]) ;
      }

    }
    // mu+ Xc*b
    target += normal_id_glm_lpdf(Yc | Xc, mu, b, sigma);
  }

  target += beta_lpdf(R2D2_R2 | R2D2_mean_R2 * R2D2_prec_R2,
                    (1 - R2D2_mean_R2) * R2D2_prec_R2); // R^2
  target += dirichlet_lpdf(R2D2_phi | R2D2_alpha); //phi

  target += normal_lpdf(Intercept | 0, 10); // Intercept
  target += std_normal_lpdf(zb); //zb: overall effects

  for(k in 1:K){
    for(d in 1: Dg){
      target += std_normal_lpdf(z[k,d,]); // z
    }
  }

  target += student_t_lpdf(sigma | 3, 0, sd(Yc)); //
}

generated quantities {
  //---actual population-level intercept
  real b_Intercept = Ymean+Intercept - dot_product(means_X, b);

  //---y_tilde quantities of interest

```

```

vector[N] log_lik;
real y_tilde[N];
vector[N] mu_tilde = rep_vector(0.0, N)+Ymean+Intercept +Xc*b;
vector<lower=0>[(D-1)+K+(Dg-1)*K] lambdas;

//---y_tilde calc

for (n in 1:N) {
  for(k in 1:K){
    mu_tilde[n]+=dot_product(r[k,,J[n,k]], Z[k,,n]) ;
  }
  log_lik[n] =normal_lpdf( Y[n] | mu_tilde[n], sigma);
  y_tilde[n]=normal_rng(mu_tilde[n], sigma);
}

//--- lambdas

//overall variances
lambdas[1:Dc]= sigma^2*R2D2_phi[1:Dc] ./ var_X *R2D2_tau2 ;
//varying int variances
lambdas[(Dc+1):(Dc+K)]= sigma^2*R2D2_phi[(Dc+1):(Dc+K)]
                        *R2D2_tau2;

for(k in 1:K){
  // group level variances
  // (k-1)(Dg-1) indexes the start of the kth batch of scales
  lambdas[(Dc+K+(k-1)*(Dg-1)+1):(Dc+K+(k-1)*(Dg-1)+Dg-1)]=
    sigma^2*R2D2_phi[(Dc+K+(k-1)*
                      (Dg-1)+1):(Dc+K+(k-1)*(Dg-1)+Dg-1)]
    ./ var_X*R2D2_tau2;
}
}

```

B.2. Gibbs sampling

To illustrate a Gibbs sampling approach to obtain posterior draws we consider, without loss of generality, $K = 1$ grouping factors and consider that the first q terms (i.e covariates x_1, \dots, x_q with $q \leq p$) are all varying over $L = l$ levels, where we will consider that $\sigma_{xi} = 1 \forall i = 1, \dots, p$. Therefore ϕ is a simplex of dimension $r = p + 1 + q$ (due to the inclusion of the varying intercept term). We consider a symmetric Dirichlet distribution for ϕ with concentration vector $\alpha = (a_\pi, \dots, a_\pi)'$. Our procedure can be generalized for arbitrary concentration vectors α . For ease of notation and to be able to develop a blocked Gibbs sampling we will write model (11) in matrix form. This will drastically improve performance over moving one regression term at a time.

Denote by $b = (b_1, \dots, b_p)'$, $u_j = (u_{0j}, u_{1j}, \dots, u_{qj})'$ the varying coefficients in level $j = 1, \dots, l$ and by $u = (u_1, \dots, u_l)' \in \mathbb{R}^{l(q+1)}$ the vector containing all vary-

ing coefficients. Notice that we have suppressed the notation that indicates the group since we only have one grouping factor K . Let $y = (y_1, \dots, y_N)'$ denote a vector containing the N observations $y_i, i = 1, \dots, N$, X denote the standardized design matrix of dimension $N \times p$ and Z denote the varying effects matrix (see Bates et al., 2015, for details on how to construct Z). Since we only have one grouping factor, Z is a block matrix, where each block contains the observations that vary in each level. In the following, we denote the linear predictor as $\mu = Xb + Zu$.

Let $\Sigma_b = \sigma^2 \Gamma_b$ denote the covariance matrix of the vector of overall coefficients b where $\Gamma_b = \text{diag} \{ \phi_1 \tau^2, \dots, \phi_p \tau^2 \}$. Let $\Sigma_u \in \mathbb{R}^{l(q+1) \times l(q+1)}$ denote the block diagonal covariance matrix of u which is given by $\Sigma_u = \sigma^2 \Gamma_u$ where $\Gamma_u = \text{diag} \{ \gamma_1, \dots, \gamma_l \}$, and γ_j is the covariance matrix associated to u_j for $j = 1, \dots, l$. In our context, $\gamma_j = \text{diag} \{ \phi_{p+1} \tau^2, \phi_{p+2} \tau^2, \dots, \phi_{p+1+q} \tau^2 \}$ for all $j = 1, \dots, l$. Finally, let I_N represent the identity matrix of order N . With the notation that has been introduced and using the representation (24), we can write the R2D2M2 model (11) in matrix form as

$$\begin{aligned}
 y | \mu, \sigma^2 &\sim \text{Normal}(\mu, \sigma^2 I_N), \\
 b | \sigma, \tau^2, \phi &\sim \text{Normal}(0, \Sigma_b), \quad u | \sigma, \tau^2, \phi \sim \text{Normal}(0, \Sigma_u) \\
 \phi &\sim \text{Dirichlet}(\alpha) \\
 \tau^2 | \xi &\sim \text{Gamma}(a_1, \xi), \quad \xi \sim \text{Gamma}(a_2, 1), \quad \sigma \sim p(\sigma).
 \end{aligned} \tag{29}$$

In the following we denote by $Z \sim \text{GenInvGaussian}(\chi, \rho, \nu)$, the Generalized Inverse Gaussian distribution (Robert, 1991) with parameters $\chi > 0, \rho > 0$ and $\nu \in \mathbb{R}$ if

$$p(z) \propto z^{\nu-1} \exp \{ -(\rho z + \chi/z) / 2 \}.$$

We denote by $\text{InvGamma}(c, d)$ the Inverse Gamma distribution with shape c and scale d and consider that a priori $\sigma^2 \sim \text{InvGamma}(c, d)$. Consider representation (29) of the R2D2M2 prior, then the Gibbs sampling procedure is the following:

1. Set initial values for $b, u, \sigma, \phi, \tau^2, \xi$.
2. Sample $b | u, \phi, \tau, \sigma, y \sim \text{Normal}(\bar{b}, S_b)$, where

$$\begin{aligned}
 \bar{b} &= (X'X + \Gamma_b^{-1})^{-1} X'(y - Zu) \\
 S_b &= \sigma^2 (X'X + \Gamma_b^{-1})^{-1}.
 \end{aligned}$$

3. Sample $u | b, \phi, \tau, \sigma, y \sim \text{Normal}(\bar{u}, S_u)$, where

$$\begin{aligned}
 \bar{u} &= (Z'Z + \Gamma_u^{-1})^{-1} Z'(y - Xb) \\
 S_u &= \sigma^2 (Z'Z + \Gamma_u^{-1})^{-1}.
 \end{aligned}$$

4. Sample $\sigma^2 | b, u, \phi, \tau^2 \sim \text{InvGamma}(\dot{c}, \dot{d})$, where

$$\begin{aligned}
 \dot{c} &= c + \frac{1}{2} (N + p + l(q + 1)) \\
 \dot{d} &= d + \frac{1}{2} (\|y - Xb - Zu\|_2^2 + b' \Gamma_b^{-1} b + u' \Gamma_u^{-1} u).
 \end{aligned}$$

5. Sample $\tau^2 | b, u, \sigma, \xi \sim \text{GenInvGaussian}(\chi, \rho, \nu)$, where

$$\chi = \frac{1}{\sigma^2} (b' T_b^{-1} b + u' T_u^{-1} u), \quad \rho = 2\xi, \quad \nu = a_1 - 1/2(p + l(q + 1)),$$

where T_b and T_u are such that $\Gamma_b = T_b \tau^2$ and $\Gamma_u = T_u \tau^2$ respectively. Sampling from the Generalized Inverse Gaussian is not straightforward, but the reader can refer to Leydold and Hörmann (2011) for efficient methods.

6. Sample $\xi | \tau^2 \sim \text{Gamma}(a_1 + a_2, 1 + \tau^2)$.

7. Sample $\phi | b, u, \sigma, \xi$.

To sample $\phi | b, u, \sigma$ we make use of the following proposition

Proposition B.1. *The joint posterior of $\phi | b, u, \sigma, \xi$ has the same distribution as $(T_1/T, \dots, T_r/T)$ where T_j are independently drawn according to*

$$T_j \sim \text{GenInvGaussian} \left(\frac{b_j^2}{2\sigma^2}, 2\xi, a_\pi - 1/2 \right) \quad j = 1, \dots, p$$

$$T_j \sim \text{GenInvGaussian} \left(\frac{\sum_{j=1}^l u_{ij}^2}{2\sigma^2}, 2\xi, a_\pi - l/2 \right) \quad j = p + 1, \dots, q + p + 1,$$

where $r = p + 1 + q$, and $T = \sum_{i=1}^r T_i$.

The main idea is to integrate out τ^2 and after doing so set $\phi_i = T_i/T$. To see this works, consider the joint posterior of $\phi | b, u, \sigma, \xi$ that results from integrating τ^2 ,

$$p(\phi | b, u, \sigma, \xi) \propto \prod_{i=1}^p \phi_i^{a_\pi - 3/2} \prod_{i=p+1}^r \phi_i^{a_\pi - l/2 - 1} \times \int_0^\infty (\tau^2)^{\nu-1} \exp \left\{ -\frac{1}{2} [\rho \tau^2 + \chi/\tau^2] \right\} d\tau^2, \quad (30)$$

where $\nu = a_1 - 1/2(p + l(q + 1))$, $\rho = 2\xi$, $\chi = \frac{1}{\sigma^2} \left(\sum_{i=1}^p \frac{b_i^2}{\phi_i} + \sum_{i=p+1}^r \sum_{j=1}^l \frac{u_{ij}^2}{\phi_i} \right)$. Now consider the following proposition (see Kruijer et al., 2010, Annotation (36) for the proof).

Proposition B.2. *Suppose T_1, \dots, T_r are independent random variables, with T_i having a density f_i on $(0, \infty)$. Let $\phi_i = T_i/T$ with $T = \sum_i T_i$, then the joint density f of $(\phi_1, \dots, \phi_{r-1})$ supported on the simplex \mathcal{S}^{r-1} has the form*

$$p(\phi_1, \dots, \phi_{r-1}) = \int_0^\infty t^{r-1} \prod_{i=1}^n f_i(\phi_i t) dt, \quad (31)$$

where $\phi_r = 1 - \sum_{i=1}^{r-1} \phi_j$.

Setting $f_i(x)$ as

$$f_i(x) \propto x^{-\delta_1} \exp \left\{ \frac{b_i^2}{2\sigma^2} \frac{1}{x} \right\} \exp \{-\xi x\}, \quad i = 1, \dots, p$$

$$f_i(x) \propto x^{-\delta_2} \exp \left\{ -\frac{\sum_{j=1}^l u_{ij}^2}{2\sigma^2} \frac{1}{x} \right\} \exp \{-\xi x\}, \quad i = p + 1, \dots, p + q + 1,$$

substituting $f_i(\phi_i\tau^2)$ into (31) and simplifying we arrive to

$$p(\phi_1, \dots, \phi_{r-1}) = \prod_{j=1}^p \phi_j^{-\delta_1} \prod_{j=p+1}^{p+q+1} \phi_j^{-\delta_2} \int_0^\infty (\tau^2)^{r-\delta_1 p - \delta_2 (q+1) - 1} \times \exp \left\{ -\frac{1}{2} \left[\frac{1}{\sigma^2} \left(\sum_{i=1}^p b_i^2 / \phi_i + \sum_{i=p+1}^r \sum_{j=1}^l u_{ij}^2 / \phi_i \right) / \tau^2 + 2\xi\tau^2 \right] \right\} d\tau^2. \tag{32}$$

Comparing equations (30) and (32) we have $\delta_1 = 3/2 - a_\pi, \delta_2 = (l + 2)/2 - a_\pi$. The proof is completed by noticing that f_j are Generalized Inverse Gaussian distributions.

h) Repeat until convergence

Appendix C: Examples

C.1. Simulation based calibration

TABLE 5

Values assigned to the hyperparameters in the Simulation Based Calibration experiment 4.1.

Description	Hyperparameter	Values
Number of simulations	T	100
Number of iterations	S	3000
Groups	K	{0, 1}
Levels	L	20
Covariates	p	{10, 100, 300}
Prior mean of R^2	μ_{R^2}	{0.1, 0.5}
Prior precision of R^2	φ_{R^2}	{0.5, 1}
Concentration parameter	a_π	{0.5, 1}
Covariance matrix of x	Σ_x	{ $I_p, AR(\rho)$ } where $\rho \in \{0.5\}$

C.2. Simulations from sparse multilevel models

TABLE 6

Values assigned to the hyperparameters in the data generating process in experiment 4.2.

Description	Hyperparameter	Values
True value of R^2	R_0^2	{0.25, 0.75}
Groups	K	{1, 3}
Levels	L	20
Covariates	p	{10, 100, 300}
Prior mean of R^2	μ_{R^2}	{0.1, 0.5}
Prior precision of R^2	φ_{R^2}	{0.5, 1}
Concentration parameter	a_π	{0.5, 1}
Covariance matrix of x	Σ_x	{ $I_p, AR(\rho)$ } where $\rho \in \{0.5\}$
Level of induced sparsity	v, z	{0.5, 0.95}

C.2.1. Additional tables and figures

TABLE 7
Predictive Table results with $K = 3$.

Scenario $K = 3, \rho = 0, p_i = 0.95, R_0^2 = 0.75, N = 200$									
p	Metric	1	2	3	4	5	6	7	8
10	RMSE	0.05	0.09	0.08	0.09	0.04	0.06	0.09	0.03
	elpd	-477	-491	-491	-461	-470	-478	-497	-499
	m_{eff}	71	93	90	80	71	115	84	88
100	RMSE	0.01	0.01	0.01	0.02	0.01	0.02	0.02	0.02
	elpd	-2700	-3049	-3806	-3440	-2899	-2920	-2744	-3681
	m_{eff}	1109	1402	1388	1764	1101	1213	1260	1629
300	RMSE	0.02	0.02	0.01	0.02	0.01	0.01	0.02	0.01
	elpd	-6210	-4413	-7064	-6285	-4480	-3413	-6471	-5501
	m_{eff}	2532	2104	2461	2722	2025	1669	2642	2405

The table shows results for the Root Mean Squared Error (RMSE), expected logpointwise predictive density (elpd) on the test set and the effective number of nonzero coefficients m_{eff} .

TABLE 8
Error Table $K = 3$.

Scenario $K = 3, \rho = 0, p_i = 0.95, R_0^2 = 0.75, N = 200, \alpha = 0.05$											
p	Metric	1	2	3	4	5	6	7	8		
10	Type I error	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
	Type II error	0.34 (0.32)	0.32 (0.24)	0.23 (0.13)	0.31 (0.18)	0.23 (0.21)	0.21 (0.13)	0.17 (0.17)	0.22 (0.20)		
	FDR	0.12 (0.00)	0.11 (0.00)	0.11 (0.00)	0.11 (0.00)	0.03 (0.00)	0.08 (0.00)	0.15 (0.01)	0.02 (0.00)		
	TNDR	0.99 (0.95)	0.98 (0.96)	0.98 (0.98)	0.98 (0.97)	0.99 (0.96)	0.99 (0.98)	0.98 (0.97)	0.99 (0.97)		
100	Type I error	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)		
	Type II error	0.50 (0.37)	0.48 (0.39)	0.40 (0.33)	0.52 (0.35)	0.35 (0.27)	0.47 (0.35)	0.49 (0.37)	0.51 (0.30)		
	FDR	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.02 (0.02)		
	TNDR	0.99 (0.97)	0.99 (0.98)	0.99 (0.98)	0.99 (0.97)	0.99 (0.98)	0.99 (0.97)	0.99 (0.97)	0.99 (0.98)		
300	Type I error	0.01 (0.01)	0.01 (0.01)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)		
	Type II error	0.83 (0.72)	0.89 (0.80)	0.85 (0.76)	0.90 (0.78)	0.82 (0.73)	0.85 (0.81)	0.87 (0.75)	0.87 (0.80)		
	FDR	0.02 (0.02)	0.02 (0.02)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)		
	TNDR	0.99 (0.96)	0.99 (0.95)	0.99 (0.96)	0.99 (0.96)	0.99 (0.96)	0.99 (0.95)	0.99 (0.95)	0.99 (0.95)		

We show the results for all the coefficients as well as results for the overall coefficients inside the parenthesis. The numbers 1-8 indicate the hyperparameter setup used.

TABLE 9
Coverage Table $K = 3$.

Scenario $K = 3, \rho = 0, p_i = 0.95, R_0^2 = 0.75, N = 200, \alpha = 0.05$												
p	Description	Metric	1	2	3	4	5	6	7	8		
10	All	Coverage	0.99 (0.42)	0.99 (0.46)	0.98 (0.46)	0.98 (0.43)	0.99 (0.47)	0.99 (0.51)	0.98 (0.41)	0.99 (0.60)		
		Width	0.79 (1.50)	0.94 (2.16)	1.00 (2.06)	0.96 (2.04)	0.80 (2.02)	0.94 (2.14)	1.01 (1.87)	0.84 (1.84)		
	Overall	Coverage	0.97 (0.80)	0.97 (0.77)	0.98 (0.90)	0.96 (0.74)	0.97 (0.80)	0.97 (0.80)	0.98 (0.88)	0.98 (0.86)		
		Width	0.46 (0.75)	0.55 (0.82)	0.58 (0.89)	0.57 (0.80)	0.48 (0.73)	0.55 (0.81)	0.57 (0.84)	0.51 (0.72)		
	100	All	Coverage	0.99 (0.39)	0.99 (0.33)	0.99 (0.37)	0.99 (0.33)	0.99 (0.40)	0.99 (0.37)	0.99 (0.34)	0.99 (0.36)	
		Overall	Coverage	0.98 (0.73)	0.97 (0.55)	0.98 (0.68)	0.97 (0.53)	0.98 (0.71)	0.97 (0.54)	0.98 (0.69)	0.97 (0.59)	
300	All	Coverage	0.82 (1.42)	0.75 (1.43)	0.74 (1.34)	0.91 (1.50)	0.73 (1.33)	0.85 (1.42)	0.77 (1.37)	0.95 (1.57)		
		Width	0.98 (0.29)	0.98 (0.26)	0.99 (0.29)	0.99 (0.28)	0.99 (0.30)	0.99 (0.22)	0.99 (0.28)	0.99 (0.29)		
	Overall	Coverage	1.32 (1.55)	1.29 (1.51)	1.27 (1.55)	1.37 (1.53)	1.26 (1.58)	1.23 (1.37)	1.39 (1.61)	1.36 (1.62)		
		Width	0.95 (0.34)	0.94 (0.23)	0.96 (0.35)	0.96 (0.24)	0.96 (0.33)	0.96 (0.25)	0.96 (0.35)	0.96 (0.24)		
	300	All	Coverage	1.21 (2.03)	1.19 (1.79)	1.15 (1.95)	1.26 (1.91)	1.16 (2.01)	1.14 (1.73)	1.27 (2.16)	1.26 (1.92)	
		Overall	Width									

Credibility intervals were formed at a $\alpha = 0.05$ level. We show results for the set of all coefficients mentioned in the description and for the non-zero corresponding cases inside parenthesis. The numbers 1-8 indicate the hyperparameter setup used.




FIG 17. Densities of the elpd estimators on the test set as p increases and arranged by hyperparameter configurations when $K = 3$. The vertical lines inside each density represent the 5%, 50%, 95% quantiles from left to right respectively. This shows that, in average, we can expect similar results, however for single realizations differences can be observed due to the heavy tails present. Hence, proper hyperparameter specification should be done.




FIG 18. Densities of the posterior median of the effective number of coefficients per hyperparameter configuration and number of covariates when $K = 3$.

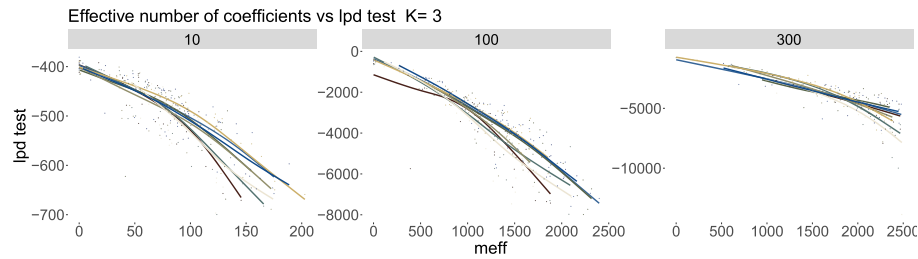


FIG 19. Relationship between elpd and shrinkage by hyperparameter configuration when $K = 3$. The non-linear relationships described by the colored lines in general shows a decreasing behavior as m_{eff} increases.

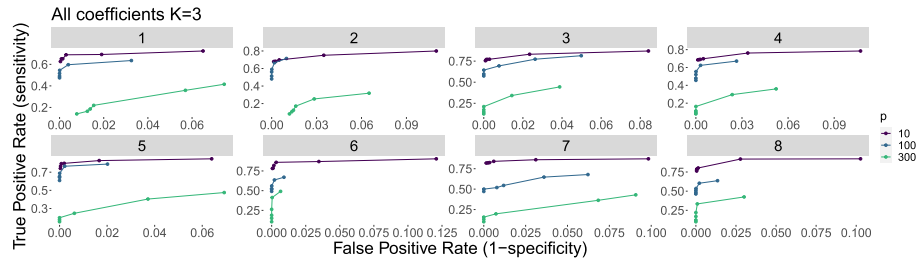


FIG 20. ROC curves for all the coefficients arranged by hyperparameter configurations when $K = 3$. Points are calculated by moving α when forming credibility intervals. Notice that the shape of the curve is due to the fact that even for high values of α , False Positive Rates don't increase. All ROC curves shown are above the identity line.

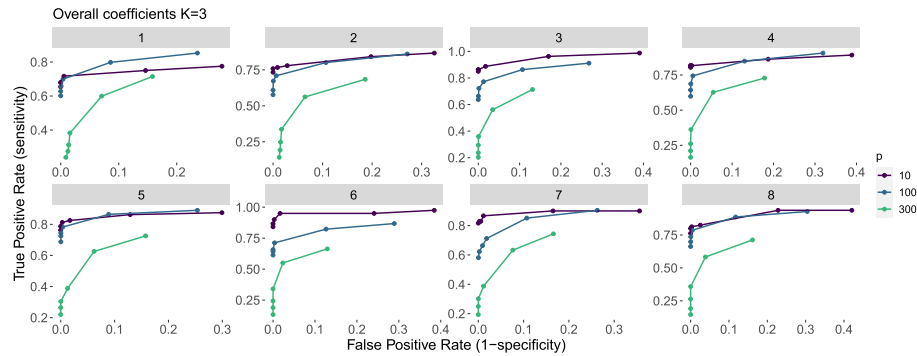


FIG 21. ROC curves for the overall coefficients when $K = 3$. Points are calculated by moving α when forming credibility intervals. Notice that the upper bound on the curve is due to the fact that even for high values of α , False Positive Rates don't increase. All ROC curves shown are above the identity line.

Acknowledgments

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2075-390740016 and DFG Project 500663361. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech).

References

- Aitchison, J. and Shen, S. M. (1980). "Logistic-Normal Distributions: Some Properties and Uses." *Biometrika*, 67(2): 261–272. URL <http://www.jstor.org/stable/2335470> MR0581723
- Bai, R. and Ghosh, M. (2019). "Large-scale multiple hypothesis testing with the normal-beta prime prior." *Statistics*, 53(6): 1210–1233. URL <https://doi.org/10.1080/02331888.2019.1662017> MR4034859

- Barr, D. J., Levy, R., Scheepers, C., and Tily, H. J. (2013). “Random effects structure for confirmatory hypothesis testing: Keep it maximal.” *Journal of memory and language*, 68 3.
- Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). “Fitting Linear Mixed-Effects Models Using lme4.” *Journal of Statistical Software*, 67(1): 1–48. URL <https://www.jstatsoft.org/index.php/jss/article/view/v067i01>
- Benjamini, Y. and Hochberg, Y. (1995). “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1): 289–300. URL <http://www.jstor.org/stable/2346101> MR1325392
- Bhadra, A., Datta, J., Polson, N. G., and Willard, B. (2016). “Default Bayesian analysis with global-local shrinkage priors.” *Biometrika*, 103(4): 955–969. URL <http://www.jstor.org/stable/26363497> MR3620450
- Bhattacharya, A., Pati, D., Pillai, N. S., and Dunson, D. B. (2015). “Dirichlet–Laplace Priors for Optimal Shrinkage.” *Journal of the American Statistical Association*, 110(512): 1479–1490. PMID: 27019543. URL <https://doi.org/10.1080/01621459.2014.960967> MR3449048
- Brooks, Gelman, S., and Jones, A. (2011). *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 1 edition.
- Browne, W. J. and Draper, D. (2006). “A comparison of Bayesian and likelihood-based methods for fitting multilevel models.” *Bayesian Analysis*, 1(3): 473–514. URL <https://doi.org/10.1214/06-BA117> MR2221283
- Buehlmann, P., Kalisch, M., and Meier, L. (2014). “High-Dimensional Statistics with a View Toward Applications in Biology.” *Annual Review of Statistics and Its Application*, 1(1): 255–278. URL <https://doi.org/10.1146/annurev-statistics-022513-115545>
- Bürkner, P.-C. (2017). “brms: An R Package for Bayesian Multilevel Models Using Stan.” *Journal of Statistical Software*, 80(1): 1–28. URL <https://www.jstatsoft.org/index.php/jss/article/view/v080i01>
- Bürkner, P.-C.(2021). “Bayesian Item Response Modeling in R with brms and Stan.” *Journal of Statistical Software*, 100(5): 1–54. URL <https://www.jstatsoft.org/index.php/jss/article/view/v100i05>
- Bürkner, P.-C. and Charpentier, E. (2020). “Modelling monotonic effects of ordinal predictors in Bayesian regression models.” *British Journal of Mathematical and Statistical Psychology*, 73(3): 420–451. URL <https://bpspsychub.onlinelibrary.wiley.com/doi/abs/10.1111/bmsp.12195>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). “Stan: A Probabilistic Programming Language.” *Journal of Statistical Software*, 76(1): 1–32. URL <https://www.jstatsoft.org/index.php/jss/article/view/v076i01>
- Carvalho, Polson, and Scott (2010). “The horseshoe estimator for sparse signals.” *Biometrika*, 97(2): 465–480. URL <http://www.jstor.org/stable/25734098> MR2650751
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2009). “Handling Sparsity via the Horseshoe.” In van Dyk, D. and Welling, M. (eds.), *Proceedings of*

- the *Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, 73–80. Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA: PMLR. URL <https://proceedings.mlr.press/v5/carvalho09a.html>
- Catalina, A., Bürkner, P.-C., and Vehtari, A. (2020). “Projection Predictive Inference for Generalized Linear and Additive Multilevel Models.”
- Depaoli, S. and Clifton, J. P. (2015). “A Bayesian Approach to Multilevel Structural Equation Modeling With Continuous and Dichotomous Outcomes.” *Structural Equation Modeling: A Multidisciplinary Journal*, 22: 327–351. MR3357137
- Fawcett, T. (2006). “An introduction to ROC analysis.” *Pattern Recognition Letters*, 27(8): 861–874. ROC Analysis in Pattern Recognition. URL <https://www.sciencedirect.com/science/article/pii/S016786550500303X>
- Fuglstad, G.-A., Simpson, D., Lindgren, F., and Rue, H. (2019). “Constructing Priors that Penalize the Complexity of Gaussian Random Fields.” *Journal of the American Statistical Association*, 114(525): 445–452. URL <https://doi.org/10.1080/01621459.2017.1415907> MR3941267
- Gelman, A. (2006). “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper).” *Bayesian Analysis*, 1(3): 515–534. URL <https://doi.org/10.1214/06-BA117A> MR2221284
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., and Rubin, D. (2013). *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis. URL <https://books.google.de/books?id=ZXL6AQAAQBAJ>
- Gelman, A. and Hill, J. (2006). *Data Analysis Using Regression and Multi-level/Hierarchical Models*. Analytical Methods for Social Research. Cambridge University Press. MR2262986
- Gelman, A., Hill, J., and Vehtari, A. (2020a). *Regression and Other Stories*. Analytical Methods for Social Research. Cambridge University Press.
- Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., Kennedy, L., Gabry, J., Bürkner, P.-C., and Modrák, M. (2020b). “Bayesian Workflow.” URL <https://arxiv.org/abs/2011.01808>
- Good, I. J. (1962). “Theory of Probability Harold Jeffreys (Third edition, 447 + ix pp., Oxford Univ. Press, 84s.)” *Geophysical Journal International*, 6: 555–558. MR0187257
- Goodrich, B., Gabry, J., Ali, I., and Brilleman, S. (2020). “rstanarm: Bayesian applied regression modeling via Stan.” R package version 2.21.1. URL <https://mc-stan.org/rstanarm>
- Gorinova, M. I., Gordon, A. D., and Sutton, C. (2019). “Probabilistic Programming with Densities in SlicStan: Efficient, Flexible, and Deterministic.” *Proc. ACM Program. Lang.*, 3(POPL). URL <https://doi.org/10.1145/3290348>
- Harrell, F. (2013). *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. Springer Series in Statistics. Springer New York. URL <https://books.google.de/books?id=7D0mBQAAQBAJ>
- Hoerl, A. E. and Kennard, R. W. (1970). “Ridge Regression: Biased Estimation

- for Nonorthogonal Problems.” *Technometrics*, 12(1): 55–67. URL <http://www.jstor.org/stable/1267351>
- Hoffman, M. D. and Gelman, A. (2014). “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo.” *J. Mach. Learn. Res.*, 15(1): 1593–1623. [MR3214779](#)
- Jeffrey, A., Zwillinger, D., Gradshteyn, I., and Ryzhik, I. (2007). “8–9 - Special Functions.” In *Table of Integrals, Series, and Products (Seventh Edition)*, 859–1048. Boston: Academic Press, seventh edition edition. URL <https://www.sciencedirect.com/science/article/pii/B9780080471112500169> [MR2360010](#)
- Jones, M. (2009). “Kumaraswamy’s distribution: A beta-type distribution with some tractability advantages.” *Statistical Methodology*, 6(1): 70–81. URL <https://www.sciencedirect.com/science/article/pii/S1572312708000282> [MR2655540](#)
- Kim, S., Moon, H., Modrák, M., and Säilynoja, T. (2022). *SBC: Simulation Based Calibration for rstan/cmdstanr models*. <https://hyunjimoon.github.io/SBC/>, <https://github.com/hyunjimoon/SBC/>
- Kruijer, W., Rousseau, J., and Vaart, A. (2010). “Adaptive Bayesian Density Estimation with Location-Scale Mixtures.” *Electronic Journal of Statistics*, 4. [MR2735885](#)
- Kruschke, J. K. (2015). “Chapter 6 - Inferring a Binomial Probability via Exact Mathematical Analysis.” In Kruschke, J. K. (ed.), *Doing Bayesian Data Analysis (Second Edition)*, 123–141. Boston: Academic Press, second edition edition. URL <https://www.sciencedirect.com/science/article/pii/B9780124058880000064>
- Laura Dietz, B. O., Jaakko Luttinen (2010). “BayesNet.” URL <https://github.com/jluttine/tikz-bayesnet>
- Leydold, J. and Hörmann, W. (2011). “Generating generalized inverse Gaussian random variates by fast inversion.” *Computational Statistics and Data Analysis*, 55(1): 213–217. URL <https://www.sciencedirect.com/science/article/pii/S0167947310002847> [MR2736548](#)
- Lin, J. (2016). “On The Dirichlet Distribution by Jiayu Lin.”
- Lin, L., Drton, M., and Shojaie, A. (2020). “Statistical Significance in High-Dimensional Linear Mixed Models.” In *Proceedings of the 2020 ACM-IMS on Foundations of Data Science Conference, FODS’20*, 171–181. New York, NY, USA: Association for Computing Machinery. URL <https://doi.org/10.1145/3412815.3416883>
- Mikkola, P., Martin, O. A., Chandramouli, S., Hartmann, M., Pla, O. A., Thomas, O., Pesonen, H., Corander, J., Vehtari, A., Kaski, S., Bürkner, P.-C., and Klami, A. (2021). “Prior knowledge elicitation: The past, present, and future.”
- Miller, P. (2006). *Applied Asymptotic Analysis*. American Mathematical Society, 1 edition. [MR2238098](#)
- Morris, M., Wheeler-Martin, K., Simpson, D. P., Mooney, S. J., Gelman, A., and DiMaggio, C. J. (2019). “Bayesian hierarchical spatial models: Implementing the Besag York Mollié model in stan.” *Spatial and spatio-temporal*

- epidemiology*, 31: 100301.
- Nakagawa, S. and Schielzeth, H. (2013). “A general and simple method for obtaining R^2 from generalized linear mixed-effects models.” *Methods in Ecology and Evolution*, 4(2): 133–142. URL <https://besjournals.onlinelibrary.wiley.com/doi/abs/10.1111/j.2041-210x.2012.00261.x>
- Nishimura, A. and Suchard, M. A. (2022). “Shrinkage with Shrunk Shoulders: Gibbs Sampling Shrinkage Model Posteriors with Guaranteed Convergence Rates.” *Bayesian Analysis*, 1–24. URL <https://doi.org/10.1214/22-BA1308> MR4578057
- Olver, F., of Standards, N. I., (U.S.), T., Lozier, D., Boisvert, R., and Clark, C. (2010). *NIST Handbook of Mathematical Functions Hardback and CD-ROM*. Cambridge University Press. URL <https://books.google.de/books?id=3I15Ph1Qf38C> MR2723248
- Paananen, T., Catalina, A., Burkner, P.-C., and Vehtari, A. (2020). “Group Heterogeneity Assessment for Multilevel Models.” *arXiv: Methodology*.
- Park, T. and Casella, G. (2008). “The Bayesian Lasso.” *Journal of the American Statistical Association*, 103(482): 681–686. URL <https://doi.org/10.1198/016214508000000337> MR2524001
- Piironen, J., Paasiniemi, M., and Vehtari, A. (2020). “Projective inference in high-dimensional problems: Prediction and feature selection.” *Electronic Journal of Statistics*, 14(1): 2155–2197. URL <https://doi.org/10.1214/20-EJS1711> MR4097052
- Piironen, J. and Vehtari, A. (2017). “Sparsity information and regularization in the horseshoe and other shrinkage priors.” *Electronic Journal of Statistics*, 11(2): 5018–5051. URL <https://doi.org/10.1214/17-EJS1337SI> MR3738204
- Polson, N. G. and Scott, J. G. (2012). “On the Half-Cauchy Prior for a Global Scale Parameter.” *Bayesian Analysis*, 7(4): 887–902. URL <https://doi.org/10.1214/12-BA730> MR3000018
- Rights, J. D. and Sterba, S. K. (2019). “Quantifying explained variance in multilevel models: An integrative framework for defining R-squared measures.” *Psychological methods*, 24 3: 309–338.
- Robert, C. (1991). “Generalized inverse normal distributions.” *Statistics & Probability Letters*, 11(1): 37–41. URL <https://www.sciencedirect.com/science/article/pii/016771529190174P>
- Simpson, D., Rue, H., Riebler, A., Martins, T. G., and Sørbye, S. H. (2017). “Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors.” *Statistical Science*, 32(1): 1–28. URL <https://doi.org/10.1214/16-STS576> MR3634300
- Stan Development Team (2022). “Stan Modeling Language Users Guide and Reference Manual, Version 2.30.0.” URL <http://mc-stan.org/>
- Säilynoja, T., Bürkner, P.-C., and Vehtari, A. (2021). “Graphical Test for Discrete Uniformity and its Applications in Goodness of Fit Evaluation and Multiple Sample Comparison.” MR4402179
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., and Gelman, A. (2020). “Validating Bayesian Inference Algorithms with Simulation-Based Calibra-

- tion.”
- van der Pas, S. (2021). “Theoretical guarantees for the horseshoe and other global-local shrinkage priors.” In *Handbook of Bayesian Variable Selection*, 133–160. Chapman and Hall/CRC.
- Van Erp, S., Oberski, D., and Mulder, J. (2019). “Shrinkage priors for Bayesian penalized regression.” *Journal of Mathematical Psychology*, 89: 31–50. [MR3903921](#)
- Vehtari, A. and Ojanen, J. (2012). “A survey of Bayesian predictive methods for model assessment, selection and comparison.” *Statistics Surveys*, 6(none): 142 – 228. URL <https://doi.org/10.1214/12-SS102> [MR3011074](#)
- Wakefield, J. (2013). *Bayesian and Frequentist Regression Methods*. Springer Series in Statistics. Springer New York. URL <https://books.google.de/books?id=OUJEAAAQBAJ> [MR3025225](#)
- Wood, S. (2017). *Generalized Additive Models: An Introduction with R*. Chapman & Hall / CRC texts in statistical science. CRC Press/Taylor & Francis Group. URL <https://books.google.de/books?id=0itmjwEACAAJ>
- Wood, S. N. (2011). “Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models.” [MR2797734](#)
- Yanchenko, E., Bondell, H. D., and Reich, B. J. (2021). “The R2D2 Prior for Generalized Linear Mixed Models.” URL <https://arxiv.org/abs/2111.10718>
- Zanella, G. and Roberts, G. (2021). “Multilevel Linear Models, Gibbs Samplers and Multigrid Decompositions (with Discussion).” *Bayesian Analysis*, 16(4): 1309–2770. URL <https://doi.org/10.1214/20-BA1242> [MR4381136](#)
- Zhang, Y. D., Naughton, B. P., Bondell, H. D., and Reich, B. J. (2020). “Bayesian Regression Using a Prior on the Model Fit: The R2-D2 Shrinkage Prior.” *Journal of the American Statistical Association*, 0(0): 1–13. URL <https://doi.org/10.1080/01621459.2020.1825449> [MR4436318](#)