

Fast and Accurate Estimation of Non-Nested Binomial Hierarchical Models Using Variational Inference

Max Goplerud

Abstract.

Non-linear hierarchical models are commonly used in many disciplines. However, inference in the presence of non-nested effects and on large datasets is challenging and computationally burdensome. This paper provides two contributions to scalable and accurate inference. First, I derive a new mean-field variational algorithm for estimating binomial logistic hierarchical models with an arbitrary number of non-nested random effects. Second, I propose “marginally augmented variational Bayes” (MAVB) that further improves the initial approximation through a step of Bayesian post-processing. I prove that MAVB provides a guaranteed improvement in the approximation quality at low computational cost and induces dependencies that were assumed away by the initial factorization assumptions.

I apply these techniques to a study of voter behavior using a high-dimensional application of the popular approach of multilevel regression and post-stratification (MRP). Existing estimation took hours whereas the algorithms proposed run in minutes. The posterior means are well-recovered even under strong factorization assumptions. Applying MAVB further improves the approximation by partially correcting the under-estimated variance. The proposed methodology is implemented in an open source software package.

Keywords: hierarchical models, variational Bayes, marginal augmentation, scalable statistical methodology.

1 Introduction and Motivating Example

Hierarchical models, often known as multilevel, mixed, or random effects models, are ubiquitous in the social sciences (Gelman and Hill 2006; Rabe-Hesketh and Skrondal 2008). In political science alone, these models are used for addressing unobserved heterogeneity, explicitly modeling dependence between observations, allowing effects to vary across space or time, and many other applications (e.g. Clark and Linzer 2015; Bell and

*Open-source software to implement the models in this paper in R can be downloaded from github.com/mgoplerud/vglmer. I thank the anonymous reviewers, Michael Auslen, Naoki Egami, Shusei Eshima, Justin Grimmer, June Hwang, Kosuke Imai, Pierre Jacob, Gary King, Shiro Kuriwaki, Marc Ratkovic, Sun Young Park, Casey Petroff, Marc Ratkovic, Tyler Simko, Diana Stanescu, Dustin Tingley, Soichiro Yamauchi, and participants at PolMeth 2020 and the University of Pittsburgh’s Statistics Seminar for helpful comments on earlier versions of this paper. All remaining errors are my own.

[†]Assistant Professor, Department of Political Science, University of Pittsburgh; mgoplerud@pitt.edu.

Jones 2015; Steenbergen and Jones 2002; Stegmueller 2013). They are also popular in other fields such as educational research and psychology.

The benefits and challenges of these models can be illustrated by an increasingly popular application for survey research in social science: Multilevel Regression and Post-Stratification (MRP; Gelman and Little 1997; Park et al. 2004; Gao et al. 2020). Described in more detail in Section 5, the core purpose of this method is to extrapolate outcomes from nationally representative surveys to small geographic areas with limited data (e.g. city, state, or legislative district) using (i) a rich hierarchical model fit on the national survey and the (usually) binary or binomial outcome and (ii) post-stratification of predicted values based on the underlying population. This method has been widely applied to a variety of questions such as measuring public opinion on a wide variety of policies, examining ideology at the city level, and exploring determinants of vote choice and turnout decisions (e.g. Ghitza and Gelman 2013; Lax and Phillips 2009a, 2012; Butticie and Highton 2013; Tausanovitch and Warshaw 2014).

Early applications of these models usually additively included reasonable number of non-nested effects (e.g. four), but subsequent work noted the inability of such models to capture the rich complexity of the data (Ghitza and Gelman 2013). That paper increased the complexity of the model substantially by using *eighteen* mostly non-nested random effects and thus specifying a model with thousands of parameters. More broadly, the idea of using a more complex model has led to a variety of papers implementing more complex hierarchical models (Gelman et al. 2016; Gao et al. 2020) or relying on machine learning methods (Bisbee 2019; Ornstein 2020; Goplerud et al. 2018). Regardless of whether one relies on a “traditional” MRP or a recent extension, it is clear that comparing multiple specifications in a principled way is fundamental to performing reliable inference. Given the long history and popularity of using traditional hierarchical models when performing MRP, it is essential that there is a method to fit those models reliably and quickly given computational constraints for many practitioners.

Unfortunately, inference for non-linear hierarchical models—especially at the complexity needed to be competitive with machine learning alternatives—can be challenging as the likelihood function contains an intractable, high-dimensional, integral. There are two popular methods for applied researchers (Stegmueller 2013): First, one can approximate the integral numerically (e.g. Bates et al. 2015; Rabe-Hesketh et al. 2004). Second, one can use a fully Bayesian approach and sample from the joint distribution of all of the parameters of the model (e.g. Carpenter et al. 2017). The key downside of these methods is that they can be slow even on modestly sized problems, and thus it is challenging to get estimates of reasonable quality in a modest period of time. This is a problem of “scalability” to the large and complex models required for many empirical applications. A key downside of non-scalable models is that common techniques such as K -fold cross-validation or bootstrapping are prohibitively expensive.

This paper makes two contributions to tackling this problem. First, I outline a series of new variational algorithms based on Polya-Gamma augmentation that allow coordinate ascent variational inference to be implemented for binomial logistic regression for an arbitrary number of (non-nested) random effects while imposing only a mean-field factorization assumption. This extends existing work on variational methods for

this class of model, as there does not appear to be a tailored algorithm to estimate models with more than two non-nested random effects.¹ Further, the algorithm can be implemented without assuming independence between the “fixed” (i.e. fully pooled) and random effects.

Second, I outline a generic procedure for improving an initial variational approximation when a parameter expansion of the underlying Bayesian model exists. I do this by drawing a connection to “marginal augmentation” from the Markov Chain Monte Carlo literature (e.g. [Liu and Wu 1999](#); [Van Dyk and Meng 2001](#)) and showing that this parameter expansion often permits a nearly costless improvement of the initial approximation. The method (“marginally augmented variational Bayes”—MAVB) transforms the initial approximation by sampling the expansion parameter and re-transforming the original samples while maintaining the stationarity of the target posterior. This induces dependencies between the parameters that were assumed away in estimating the initial procedure and provides a provable *guaranteed* improvement upon the original approximation.

Methodologically, this pushes forward the literature on variational inference for hierarchical models by extending work in the case of a single random effect ([Hall et al. 2011](#); [Ormerod and Wand 2012](#); [Tan and Nott 2013](#); [Hall et al. 2019](#)) or two non-nested random effects ([Jeon et al. 2017](#); [Menictas et al. 2019](#)) to the general case. The proposed method requires no integration, unlike many existing methods for binary outcomes ([Ormerod and Wand 2012](#); [Tan and Nott 2013](#); [Jeon et al. 2017](#)). It further provides a link to existing work that seeks to combine Markov Chain Monte Carlo and variational inference by stochastic optimization (e.g. [Salimans et al. 2015](#); [Ruiz and Titsias 2019](#); [Yin and Zhou 2018](#)). Instead of optimizing the transformed density, MAVB transforms the samples from the initial approximation with a partial step of MCMC using marginal augmentation that, in practice, appears as performing a stochastic location/scale transformation of the sampled parameters. This leverages a sampler that is known to mix well in the case of fully Bayesian MCMC and lacks internal tuning parameters as its primary goal is to find a computationally inexpensive way to improve an initial approximation. While it bears some similarities to work on re-parameterization in hierarchical models for variational algorithms (e.g. [Tan and Nott 2013](#); [Tan 2021](#)), it does not fix the re-parameterization in advance of estimation. It differs from other approaches that seek to improve an initial approximation (e.g. linear response variational Bayes; [Giordano et al. 2015](#)) in that it has a guarantee on improving the approximation quality. Future work could examine how such methods work alongside Polya-Gamma data augmentation.

The remainder of the paper proceeds as follows. Section 2 states multiple factorization assumptions under which Polya-Gamma augmentation can be used to estimate a variational approximation for a binomial logistic hierarchical model. Section 3 links parameter expansion to variational Bayes and explains MAVB formally.

Sections 4 and 5 conduct simulations and examine performance on the empirical example ([Ghitza and Gelman 2013](#)). The latter shows dramatic gains in speed: Even

¹Generic methods for variational inference, e.g. stochastic variational inference or automatic differentiation variational inference (ADVI; [Kucukelbir et al. 2017](#)), can be applied to most models, including hierarchical ones. I compare ADVI against my “tailored” algorithms and show it performs worse.

after applying MAVB and drawing 4,000 samples, the fastest variational algorithm is nearly 60 times faster than Laplace approximation and nearly 350 times faster than Hamiltonian Monte Carlo for the most complex models. This reduces the run time from hours to minutes. All variational methods well-recover the posterior means. While the strongest factorization assumptions have poor performance in terms of estimating the posterior variance, applying MAVB corrects a large amount of the problem.

Section 5 then uses this algorithm to engage in model comparison that was computationally infeasible in Ghitza and Gelman (2013). I perform 10-fold cross-validation across nine models ranging from having four to 18 random effects and thousands of parameters. The process takes around 30 minutes compared to the hours needed to fit even a single model once using existing approaches. The results provide some evidence of over-fitting in the original specification suggesting that the most complex model does not outperform models of intermediate complexity. I use this to draw out some guidance for practitioners of MRP in other substantive domains.

2 Mean-Field Variational Inference for Binomial Hierarchical Models

I focus on the following generative model that is broader than MRP but also captures the majority of applications. For each observation $i \in \{1, \dots, N\}$, the researcher observes y_i “successes” out of n_i trials (e.g. how many individuals in a population of size n_i turn out to vote). I model this using a binomial distribution with probability of success p_i defined via a linear predictor (ψ_i) put through a logistic link. Equation 1 expresses this model using a “general design” notation (Zhao et al. 2006). Appendix A shows the model using Gelman and Hill (2006)’s notation and a plate diagram.

$$y_i | \beta, \alpha \sim \text{Binom}(n_i, p_i), \quad p_i = \frac{\exp(\psi_i)}{1 + \exp(\psi_i)}, \quad \psi_i = \mathbf{x}_i^T \beta + \mathbf{z}_i^T \alpha \quad (1a)$$

$$\alpha_j | \Sigma_j \sim N(\mathbf{0}, \mathbf{I}_{g_j} \otimes \Sigma_j), \quad \Sigma_j \sim \text{IW}(\nu_j, \Phi_j), \quad p(\beta) \propto 1 \quad (1b)$$

$$\mathbf{z}_{i,j} = \mathbf{m}_{i,j} \otimes \mathbf{z}_{i,j}^b, \quad \alpha^T = [\alpha_1^T, \dots, \alpha_J^T], \quad \mathbf{z}_i^T = [\mathbf{z}_{i,1}^T, \dots, \mathbf{z}_{i,J}^T] \quad (1c)$$

As is standard in hierarchical models, the linear predictor consists of p “fixed” effects: $\mathbf{x}_i \in \mathbb{R}^p$. The hierarchical component contains J random effects indexed from $j \in \{1, \dots, J\}$. For each random effect j , there is a d_j dimensional covariate vector indexed by $\mathbf{z}_{i,j}^b$ where $\mathbf{z}_{i,j}^b = 1$ represents the ubiquitous “random intercept.” Each random effect has g_j groups and each observation i is assigned to exactly one group for each random effect; define its membership for random effect j as a one-hot vector $\mathbf{m}_{i,j} \in \{0, 1\}^{g_j}$.

The notation in Equation 1 stacks together the hierarchical components as follows; first, for each random effect j , $\mathbf{z}_{i,j}$ represents a $d_j \times g_j$ length vector (mostly sparse) by the Kronecker product (\otimes) of the group membership vector $\mathbf{m}_{i,j}$ and the base covariate. This repeats $\mathbf{z}_{i,j}^b$ once in the position corresponding to the group of which i is a member for random effect j . This allows us to model the distribution of the entire

parameter vector for random effect j ($\alpha_j \in \mathbb{R}^{g_j \cdot d_j}$) as a multivariate normal with a block diagonal matrix where each block is given an identical Inverse Wishart prior as noted in Equation 1b ($\Sigma_j \in \mathbb{R}^{d_j \times d_j}$; $\Phi_j \in \mathbb{R}^{d_j \times d_j}$). Using such priors is standard in the literature on variational inference for hierarchical models (e.g. Tan and Nott 2013), although extensions to more weakly informative priors are possible (e.g. Huang and Wand 2013). The compact notation in Equation 1a stacks together all random effects j into a single vector $\mathbf{z}_i \in \mathbb{R}^{\sum_{j=1}^J g_j \cdot d_j}$ that is highly sparse. It thus accommodates designs with arbitrary patterns of crossing (non-nesting) amongst the J random effects.

A key distinguishing feature of this model as applied to MRP is that J can be large (e.g. greater than ten) and g_j ranges widely from a handful up to over a thousand (e.g. $g_j = 4$ for ethnicity and $g_j = 1,020$ for state-ethnicity-age combinations in Ghitza and Gelman 2013). In most applications for MRP, $d_j = 1$ and $\mathbf{z}_{i,j}^b = 1$ (random intercept) but sometimes $d_j = 2$ in the case of a random slope and intercept (Gelman and Hill 2006). Regarding the other parameters, for most applications of MRP, N is often relatively modest given post-stratification requirements (see Section 5) and that surveys can be collapsed into units with identical state-demographic covariates by allowing varying n_i . Thus, in many studies, N can be made smaller than 10,000 (e.g. below 5,000 in Park et al. 2004; Ghitza and Gelman 2013). The size of β (p) is also usually modest and below ten.

By using Polya-Gamma augmentation, the model in Equation 1 can be rendered conditionally conjugate, enabling the straightforward application of numerous standard algorithms for Bayesian inference (Polson et al. 2013). Specifically, Equation 2 from Polson et al. (2013) states that for any $a, b > 0$ the following identity holds, where $f_{PG}(\omega|b, c)$ denotes the Polya-Gamma density with parameters b and c . The definition of a Polya-Gamma variable as a weighted infinite convolution of Gamma random variables is also shown.

$$\frac{\exp(\psi)^a}{[1 + \exp(\psi)]^b} = 2^{-b} \int \exp(s\psi - \psi^2/2\omega) f_{PG}(\omega|b, 0) d\omega, \quad s = a - b/2 \quad (2a)$$

$$\omega \sim PG(b, c) := \omega = \frac{1}{2\pi^2} \sum_{k=1}^{\infty} \frac{Z_k}{(k - 1/2)^2 + c^2/(4\pi^2)}, \quad Z_k \stackrel{i.i.d.}{\sim} \text{Gamma}(b, 1) \quad (2b)$$

Thus, the complete data likelihood can be expressed as follows where Ω denotes the $N \times N$ diagonal matrix of the corresponding ω_i , \mathbf{X} , \mathbf{Z} stack the data for each observation into a $N \times p$ and $N \times \sum_{j=1}^J d_j g_j$ design matrices, and \mathbf{s} is a $N \times 1$ vector with $[\mathbf{s}]_i = y_i - n_i/2$.

$$p(\mathbf{y}, \Omega | \alpha, \beta) \propto \exp \left(\mathbf{s}^T [\mathbf{X}\beta + \mathbf{Z}\alpha] - \frac{1}{2} [\mathbf{X}\beta + \mathbf{Z}\alpha]^T \Omega [\mathbf{X}\beta + \mathbf{Z}\alpha] \right) \prod_{i=1}^N f_{PG}(\omega_i | n_i, 0) \quad (3)$$

Noting the result from Polson et al. (2013) that the full conditional of $\omega_i | \mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \{\boldsymbol{\Sigma}_j\}_{j=1}^J$ has a Polya-Gamma distribution $PG(n_i, \mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{z}_i^T \boldsymbol{\alpha})$, it immediately follows that a Gibbs Sampler exists to sample all of the parameters in the model where the full conditionals on $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$ are normal and $\boldsymbol{\Sigma}_j$ is Inverse Wishart.

2.1 Variational Inference

The first contribution of this paper is to use the Polya-Gamma representation above to find a tractable variational algorithm to approximate the joint posterior of $p(\boldsymbol{\beta}, \boldsymbol{\alpha}, \{\boldsymbol{\Sigma}_j\}, \boldsymbol{\Omega} | \mathbf{y})$ and thus the joint posterior on the parameters excluding $\boldsymbol{\Omega}$. Blei et al. (2017) provides a recent review of these methods. Equation 4 formulates the problem where \mathcal{X} denotes some (restricted) set of distributions to optimize over. It can be equivalently expressed as finding the closest distribution in \mathcal{X} to the true posterior in terms of KL-divergence. For notational simplicity, denote $\boldsymbol{\theta} = \{\boldsymbol{\beta}, \boldsymbol{\alpha}, \{\boldsymbol{\Sigma}_j\}_{j=1}^J, \boldsymbol{\Omega}\}$.

$$q^*(\boldsymbol{\theta}) = \arg \max_{q(\boldsymbol{\theta}) \in \mathcal{X}} \text{ELBO}_{q(\boldsymbol{\theta})} \quad \text{where} \quad \text{ELBO}_{q(\boldsymbol{\theta})} = E_{q(\boldsymbol{\theta})} [\ln p(\mathbf{y}, \boldsymbol{\theta})] - E_{q(\boldsymbol{\theta})} [\ln q(\boldsymbol{\theta})] \quad (4)$$

A common method for solving this problem is known as “coordinate ascent variational inference” (CAVI; Blei et al. 2017). It maximizes or increases the target ELBO with respect to some sub-block of $\boldsymbol{\theta}$. By cycling through $\boldsymbol{\theta}$ repeatedly, a local optimum can be obtained. The choice of restriction \mathcal{X} is crucial to the accuracy of the approximation method; an extremely popular choice is a “mean-field” factorization assumption where blocks of parameters are assumed to be independent.

Leveraging the existence of a Gibbs Sampler, Result 1 states that the augmented posterior on $q(\boldsymbol{\theta})$ can be approximated using a number of mean-field assumptions with no further restrictions on distributional form, all updates having closed analytical forms, and for arbitrary J, d_j, g_j . Appendix A provides the full derivations as well as noting how to back out the corresponding Gibbs Sampler.

Result 1 (Existence of CAVI). *Consider the three factorization assumptions:*

$$\begin{aligned} \text{Scheme I:} \quad & \text{“Strong Factorization”} & \quad & \mathcal{X}_1 = q(\boldsymbol{\beta}) \prod_{j=1}^J q(\boldsymbol{\alpha}_j) q(\boldsymbol{\Sigma}) q(\boldsymbol{\Omega}) \\ \text{Scheme II:} \quad & \text{“Partial Factorization”} & \quad & \mathcal{X}_2 = q(\boldsymbol{\beta}) q(\boldsymbol{\alpha}) q(\boldsymbol{\Sigma}) q(\boldsymbol{\Omega}) \\ \text{Scheme III:} \quad & \text{“Limited Factorization”} & \quad & \mathcal{X}_3 = q(\boldsymbol{\beta}, \boldsymbol{\alpha}) q(\boldsymbol{\Sigma}) q(\boldsymbol{\Omega}) \end{aligned}$$

For the model in Equation 1 and for each choice of \mathcal{X}_k above, each step of the CAVI algorithm can be implemented exactly in closed form, with no additional assumptions. For each \mathcal{X}_k , the optimal approximation for $q(\boldsymbol{\beta}, \boldsymbol{\alpha})$ is multivariate normal, $q(\boldsymbol{\Sigma})$ is the product of J independent Inverse Wishart densities, and $q(\boldsymbol{\Omega})$ is the product of N independent Polya-Gammas.

Algorithm 1 explicitly outlines the updates for Scheme I. Experiments showed that convergence could be improved at little computational cost by jointly updating the mean parameters of $q(\boldsymbol{\beta})$ and $q(\boldsymbol{\alpha})$; see Appendix C for discussion. All models estimated in the paper use this acceleration technique.

Algorithm 1 CAVI for Scheme I

Set Priors of Inverse Wishart: $\{\nu_j, \Phi_j\}_{j=1}^J$; **Set Number of Iterations:** T
Initialize Variational Parameters: $\{\tilde{b}_i, \tilde{c}_i\}_{i=1}^N$ (for Poly-Gamma); $\tilde{\mu}_\beta, \tilde{\Lambda}_\beta, \tilde{\mu}_\alpha, \tilde{\Lambda}_\alpha$
(for β, α); $\{\tilde{\nu}_j, \tilde{\Phi}_j\}_{j=1}^J$ (for Σ_j)
For t in $1, \dots, T$

1. Update Poly-Gammas - $q(\{\omega_i\}_{i=1}^N)$: $\tilde{b}_i = n_i, \quad \tilde{c}_i = \sqrt{E_{q(\alpha, \beta)}[(\mathbf{x}_i^T \beta + \mathbf{z}_i^T \alpha)^2]}$
2. Update $q(\beta) \sim N(\tilde{\mu}_\beta, \tilde{\Lambda}_\beta)$:

$$\tilde{\Lambda}_\beta = \left(\sum_{i=1}^N E_{q(\omega_i)}[\omega_i] \mathbf{x}_i \mathbf{x}_i^T \right)^{-1}, \quad \tilde{\mu}_\beta = \tilde{\Lambda}_\beta \mathbf{X}^T \left(\sum_{i=1}^N \left(y_i - \frac{n_i}{2} \right) - E_{q(\omega_i)}[\omega_i] \cdot \mathbf{z}_i^T E_{q(\alpha)}[\alpha] \right)$$

3. Update $q(\alpha_j) \sim N(\tilde{\mu}_{\alpha,j}, \tilde{\Lambda}_{j,\alpha})$, where \mathbf{T}_j stacks the block diagonal expectation of the precision on the random effects (Σ_j^{-1}):

$$\tilde{\Lambda}_{\alpha,j} = \left(\mathbf{T}_j + \sum_{i=1}^N E_{q(\omega_i)}[\omega_i] \mathbf{z}_{i,j} \mathbf{z}_{i,j}^T \right)^{-1}, \quad \mathbf{T}_j = E_{q(\Sigma_j)}[\mathbf{I}_{g_j} \otimes \Sigma_j^{-1}]$$

$$\tilde{\mu}_{\alpha,j} = \tilde{\Lambda}_{\alpha,j} \mathbf{Z}_j^T \left[\sum_{i=1}^N \left(y_i - \frac{n_i}{2} \right) - E_{q(\omega_i)}[\omega_i] \cdot \left(\mathbf{x}_i^T E_{q(\beta)}[\beta] + \sum_{j': \{1, \dots, J\} \setminus j} \mathbf{z}_{i,j'}^T E_{q(\alpha_{j'})}[\alpha_{j'}] \right) \right]$$

4. Update $q(\{\Sigma_j\}_{j=1}^J)$: $\tilde{\nu}_j = \nu_j + g_j, \quad \tilde{\Phi}_j = \Phi_j + \sum_{g=1}^{g_j} E_{q(\alpha_{j,g})}[\alpha_{j,g} \alpha_{j,g}^T]$
5. Check for convergence, evaluate ELBO (see Appendix A for derivation).

This improves upon existing mean-field schemes for logistic hierarchical models in a number of ways. First, for any of the factorization assumptions, no further distributional assumptions are required (cf. Ormerod and Wand 2012; Tan and Nott 2013 assuming normality). Second, most existing algorithms for binomial outcomes require the repeated evaluation of (low) dimensional integrals at each iteration whose number scales with g_j (cf. Ormerod and Wand 2012; Tan and Nott 2013; Jeon et al. 2017). Extending these algorithms to $J > 2$ would likely incur significant computational costs as the number of those integrals increases. None of the schemes in Result 1 require integration at any step as the Poly-Gamma augmentation turns inference into iteratively performing weighted ridge regression. In the models considered in this paper, the major bottleneck as one moves from Scheme I to Scheme III is in calculating the variance term of $q(\beta, \alpha)$; even relying on a (sparse) Cholesky decomposition, this involves inverting an increasingly dense lower triangular matrix as weaker independence assumptions are imposed. Appendix E disaggregates the run-time of Algorithm 1 by stage and scheme.

Most importantly, the ability to choose between Schemes I, II, and III allows the researcher to smoothly trade-off computational cost and accuracy as in Menictas et al.

(2019)’s work on $J = 2$ for linear mixed effects models. Scheme I with its strong implied factorization assumptions is immediately scalable to huge datasets with large J or g_j . However, the downside is that the strong factorization assumptions will likely degrade performance. Scheme III provides the ability to avoid these strong assumptions at a somewhat increased computational cost. The ability to avoid such factorization assumptions for arbitrary $J > 1$ and binomial outcomes appears to be a new result. The expectation is that it will have the best performance. Scheme II is a compromise between the two extremes, and other hybrid approaches are possible such as applying re-parameterizations to the augmented posterior (e.g. Tan and Nott 2013; Tan 2021).

3 Marginally Augmented Variational Bayes

The second major contribution of the paper is demonstrating that there is a computationally cheap way of improving the initial approximation resulting from Schemes I, II, or III. The key intuition, formalized below, is that once an initial approximation $q(\boldsymbol{\theta})$ is found, one can draw samples from this approximation, perform a single step of Markov Chain Monte Carlo through (some of) the parameters, and thereby “improve” the sample. Some existing work in computer science (e.g. Salimans et al. 2015; Ruiz and Titsias 2019) has leveraged this point to attempt to *optimize* over the intractable improved density which can be computationally expensive.

By contrast, this paper explores the idea that if one can find a transition kernel with good mixing, then simply doing a single partial step can provide considerable gains at limited computational cost. While many samplers can be employed for this purpose, initial experiments suggested that the key problem was the independence assumptions in $q(\boldsymbol{\beta}, \boldsymbol{\alpha})$ and thus I chose to focus on marginal augmentation and parameter expansion as it is inexpensive to use in fully Bayesian MCMC to improve a Gibbs Sampler, has demonstrated strong performance in hierarchical models, lacks internal tuning parameters, and was explicitly designed to link the fixed and random effects together (Liu and Wu 1999; Van Dyk and Meng 2001; Gelman et al. 2008). I focus on logistic hierarchical models although the procedure is itself much more general; Section 6 discusses some broader implications and Appendix B formulates the results in a more general fashion.

The key idea behind parameter expansion is to create an “over-parameterized” model where certain additional parameters ($\boldsymbol{\xi}$) are introduced such that they (i) maintain the observed data model but (ii) are not identifiable from the observed data itself. A careful choice of parameter expansion allows the construction of algorithms that have either faster mixing for MCMC (Liu and Wu 1999; Van Dyk and Meng 2001) or faster convergence for deterministic algorithms such as EM (Liu et al. 1998). The intuition behind its effectiveness is that it allows “moves” (either via sampling steps in MCMC or parameter updates in EM) in the un-identified space that can break or escape the strong associations between parameter blocks (e.g. $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$) that slow down mixing (Liu and Wu 1999) or lead to the algorithm getting “stuck” for many iterations near boundary conditions (e.g. a small sampled $\boldsymbol{\Sigma}_j$ shrinking $\boldsymbol{\alpha}_{j,g}$ leading to a small $\boldsymbol{\Sigma}_j$, etc.; Gelman et al. 2008). Liu et al. (1998) provide a useful explanation of parameter expansion in the context of EM as a “covariance adjustment” to the estimated parameters.

In the case of hierarchical models, the most popular parameter expansion appears as a location and/or scale transformation of the random effects (e.g. [Van Dyk and Meng 2001](#)). The location transformation, for example, allows the random effects to have a non-zero mean: $\alpha_{j,g} \sim N(\mu_j, \Sigma_j)$. Note that it is not possible to estimate μ_j from the observed data but that it could be estimated if $\alpha_{j,g}$ were known. Implementation is simple and appears as a location or scale transformation of the sampled parameters that leads to very large gains in performance (e.g. [Van Dyk and Meng 2001](#); [Gelman et al. 2008](#)).

Definition 1 generalizes this parameter expansion to the arbitrary J case, where the M_j notation is bookkeeping to note which element of \mathbf{x}_i (and β) corresponds to each element of $\mathbf{z}_{i,j}^b$ (and $\alpha_{j,g}$).

Definition 1 (Expansions for Hierarchical Models). *Define a set of expansion parameters ξ that consists, for each j , of a mean shift $\mu_j \in \mathbb{R}^{d_j}$ and a scale shift $R_j \in \mathbb{R}^{d_j \times d_j}$ such that R_j is invertible. I use superscript X to denote the “expanded” parameters.*

The mapping between θ^X and θ for a fixed ξ is denoted as $t_\xi(\theta^X)$ and listed below. M_j is a $p \times d_j$ matrix such that $[M_j]_{a,b} = 1$ if the covariate corresponding to $[z_{i,j}]_b$ is the same as the covariate for $[x_i]_a$. All other elements of M_j are zero. For simplicity, assume that each element of \mathbf{z}_i corresponds to some variable in \mathbf{x}_i , i.e. that each column of M_j has exactly one non-zero element.

$$[\beta, \alpha, \{\Sigma_j\}_{j=1}^J, \Omega] = t_\xi([\beta^X, \alpha^X, \{\Sigma_j^X\}_{j=1}^J, \Omega^X]) = \begin{cases} \beta = \beta^X + \sum_{j=1}^J M_j R_j \mu_j \\ \alpha_{j,g} = R_j (\alpha_{j,g}^X - \mu_j) \\ \Sigma_j = R_j \Sigma_j^X R_j^T \\ \Omega = \Omega^X \end{cases}$$

The augmented model is listed below for an important special case treated in detail (“Mean Expansion”) in the empirical analysis. The full expansion (“Translation Expansion”) is also listed.

- *Mean Expansion: Assume all $R_j = I_{d_j}$.*

$$\begin{aligned} \ln p(y_i | \omega_i, \beta^X, \alpha^X) &\propto \mathbf{s}^T [\mathbf{X} \beta^X + \mathbf{Z} \alpha^X] - 1/2 [\mathbf{X} \beta^X + \mathbf{Z} \alpha^X]^T \Omega [\mathbf{X} \beta^X + \mathbf{Z} \alpha^X] \\ p(\beta^X) &\propto 1, \quad \alpha_{j,g}^X | \Sigma_j^X \sim N(\mu_j, \Sigma_j^X), \quad p(\Sigma_j^X) \sim IW(\nu_j, \Phi_j) \end{aligned}$$

- *Translation Expansion:*

$$\begin{aligned} \ln p(y_i | \omega_i, \beta^X, \alpha^X) &\propto \mathbf{s}^T [\mathbf{X} \beta^X + \mathbf{Z} \mathbf{R} \alpha^X] - 1/2 [\mathbf{X} \beta^X + \mathbf{Z} \mathbf{R} \alpha^X]^T \Omega [\mathbf{X} \beta^X + \mathbf{Z} \mathbf{R} \alpha^X] \\ \mathbf{R} &= \text{blockdiag}(\{I_{g_j} \otimes R_j\}_{j=1}^J), \quad p(\beta^X) \propto 1, \quad \alpha_{j,g}^X | \Sigma_j^X \sim N(\mu_j, \Sigma_j^X) \\ p(\Sigma_j^X) &\sim IW(\nu_j, \mathbf{R}_j^{-1} \Phi_j \mathbf{R}_j^{-T}) \end{aligned}$$

Given such an expanded version of the hierarchical model, there are two ways to improve the algorithms in this paper. First, drawing on [Jaakkola and Qi \(2007\)](#), it is

possible to accelerate convergence of Algorithm 1 using “parameter expanded variational Bayes” (PX-VB). Appendix C derives a new application of PX-VB to the models in Result 1 and shows it can often improve the algorithm’s convergence by decreasing the number of iterations required at effectively no computational cost as, functionally, it involves centering the random effects to be mean zero and adjusting the mean of $q(\beta)$ correspondingly.

The main use of parameter expansions in this paper, however, is to improve the *quality* of the approximation by “improving” $q(\theta)$ by performing one step of marginal augmentation where the expansion parameters ξ are sampled and then the components of θ are re-sampled. Definition 2 outlines the procedure in a general case. The notation and procedure mirrors that in Liu and Wu (1999).

Definition 2 (Marginally Augmented Variational Bayes—MAVB). *Given an initial approximation $q(\theta)$, a proper prior on the expansion parameter $p_0(\xi)$, and a one-to-one and differentiable transformation such that $t_\xi(\theta^X) = \theta$, create a new approximation $\tilde{q}(\theta)$ using the following procedure:*

1. Sample $\theta \sim q(\theta)$ and $\xi_0 \sim p_0(\xi)$.
2. Create $\theta^X = t_{\xi_0}^{-1}(\theta)$.
3. Sample a new ξ_1 as follows where $J_\xi(\theta^X)$ is the Jacobian of t_ξ with respect to θ^X and $p(\theta|\mathbf{y})$ denotes the true posterior distribution.

$$\xi_1 \sim p(\xi|\theta^X, \mathbf{y}) \propto p(t_\xi(\theta^X)|\mathbf{y}) \cdot |J_\xi(\theta^X)| \cdot p_0(\xi)$$

4. Define $\tilde{\theta} = t_{\xi_1}(\theta^X) = t_{\xi_1}(t_{\xi_0}^{-1}(\theta))$

Theorem 1 states a key result for MAVB.

Theorem 1 (Guaranteed Improvement with MAVB). *For any (proper) choice of prior $p_0(\xi)$, the MAVB approximation $\tilde{q}(\theta)$ has a better ELBO than the initial approximation:*

$$\text{ELBO}_{\tilde{q}(\theta)} \geq \text{ELBO}_{q(\theta)}$$

The proof is in Appendix B and uses two lemmas from existing results. First, Theorem 1 in Liu and Wu (1999) demonstrates the transformation to generate MAVB maintains the stationarity of the posterior. Second, a data processing inequality noted by various authors (e.g. Ruiz and Titsias 2019) showing that this transformation which keeps the *true* posterior invariant results in a better approximating distribution.

It is known from the data augmentation literature that an increasingly diffuse prior on the expansion parameters (“working prior”) allows for the parameters themselves to “decide” the best expansion parameter ξ rather than being weighed down by the prior (e.g. Liu and Wu 1999, p. 1268), and I conjecture that a similar intuition applies for MAVB. Thus, in all applications, I use an improper prior (i.e. $p_0(\xi) \propto 1$); Appendix B

discusses the validity of this prior using existing theory (Liu and Wu 1999; Van Dyk and Meng 2001), provides the result for a proper prior on ξ , and notes Algorithm 2 can be found as the limit of a proper working prior $p_0(\mu_j) \sim N(0, \tau^2 \mathbf{I})$ as $\tau \rightarrow \infty$. Algorithm 2 shows how MAVB is implemented using the mean expansion noted in Definition 1.²

Algorithm 2 Applying MAVB to Non-Linear Hierarchical Models

Set the Number of Samples Desired: M

Estimate $q(\theta)$ **using** CAVI (e.g. Algorithm 1)

For m in $1, \dots, M$

1. Draw $\theta^{(m)} \sim q(\theta)$
2. Sample the expansion parameters μ_j for each j

$$\tilde{\mu}_j \sim N\left(\frac{1}{g_j} \sum_{g=1}^{g_j} \alpha_{j,g}^{(m)}, \frac{1}{g_j} \Sigma_j^{(m)}\right)$$

3. Adjust the initial draws to get the improved sample $\tilde{\theta}^{(m)}$

$$\tilde{\alpha}_{j,g}^{(m)} = \alpha_{j,g}^{(m)} - \tilde{\mu}_j, \quad \tilde{\beta}^{(m)} = \beta^{(m)} + \sum_{j=1}^J \mathbf{M}_j \tilde{\mu}_j$$

Thus, for this model and relying only on a location transformation, MAVB has a simple form that, as shown later, can result in considerable improvements in the performance of Scheme I. As noted in the earlier discussion, the presentation of MAVB in Algorithm 2 illustrates the close relationship to the location transformation noted earlier: It can be thought of a “stochastic” location transformation given that mean of the expansion parameter is the mean of the sampled $\alpha_{j,g}$.

Some additional remarks are in order: First, if MAVB is applied to an approximation resulting from Scheme I (i.e. with independence assumed between β and α), the resulting approximation *will not* imply such an assumption. Consider the correlation between $\alpha_{j,g}$ and β in Algorithm 2. Before applying MAVB, the two parameters are independent by assumption. After applying MAVB, they have a non-zero posterior correlation because of the shared dependence on $\tilde{\mu}_j$. While not sufficient to restore *all* missing dependencies (e.g. components of β that are not included in any random effect), this can at least address some of the shortcomings of Scheme I. MAVB can be applied to the outputs of Scheme II and III, although the expectation is that the improvement for these schemes should be less pronounced given that more of those dependencies are estimated directly.

Second, the cost of implementing MAVB is quite modest, unlike existing approaches that attempt to *optimize* over the improved density (e.g. Ruiz and Titsias 2019). After

²MAVB for “Translation Expansion” (i.e. \mathbf{R}_j is not fixed) is more delicate and thus not explored here, as it requires a specific choice of prior on Σ_j and a specific choice of improper working prior to be tractable; see Van Dyk and Meng (2001) for details. Examining whether this could be used with proper priors is an interesting area for future research.

drawing a sample from $q(\theta)$, all that is needed to perform MAVB is drawing $\sum_j d_j$ univariate Gaussians (22 in the largest model considered in this paper [Model 9]), some summation of the sampled random effects and then subtracting off the sampled expansion parameter $\tilde{\mu}_j$. Note that this MAVB procedures do not require sampling the Polya-Gammas as they are left un-transformed by the algorithm nor does the cost of MAVB depend on the size of the data (N) directly; even if g_j is large, MAVB will still be fast.

Third, while MAVB is guaranteed to increase performance, the quality of MAVB is difficult to ascertain analytically in most complex models. However, insights come from simpler cases: In a stylized hierarchical model, Liu and Wu (1999) show that marginal augmentation results in perfect sampling. In the more realistic case where Σ_j is not fixed and $J = 1$, studies show that certain forms of marginal augmentation result vastly improved mixing of MCMC samplers (Van Dyk and Meng 2001; Gelman et al. 2008). Thus, there is reason to be optimistic about the ability of MAVB to improve initial approximations as the scale/location transformations in fully Bayesian marginal augmentation seem to provide quite considerable benefits over simple Gibbs samplers.

Overall, while MAVB is likely to be helpful in improving the variational schemes in this paper, it is not a panacea. Its major benefit appears to be in “connecting” blocks of parameters that were assumed to be independent in a way that is guaranteed to improve the approximation quality at a very limited computational cost. The key limitation is that its speed and scalability depends on it *not* returning to the observed data (\mathbf{y}). Interestingly, this suggests a “stronger” version of MAVB that could be performed by implementing one full sweep of the Gibbs Sampler, i.e. sampling Polya-Gammas and cycling through all full conditionals, and then performing marginal augmentation. If this were to be performed many times, the samples would converge to the true posterior by standard properties of MCMC. While this might raise its own computational concerns, exploring this is an interesting area of future research.

4 Simulation Study

I perform a simulation study to assess the accuracy of the proposed methods. I compare my variational algorithms against two gold standards (Laplace approximation using `blme` - Bates et al. 2015; Chung et al. 2015; HMC in `STAN` using `brms`; Bürkner 2017) and Automatic Differentiation Variational Inference (ADVI; Kucukelbir et al. 2017).³ The latter is a useful comparison as it is easily implemented in `STAN` and is a generic approach to approximate complex models. I show results using its mean-field approximation (MF) and full rank (FR). To begin, I conducted a simulation where the linear predictor ψ_i was generated using the following scheme ($J = 2$).

³Using `blme` allows for an identical Inverse Wishart prior to be added to the Laplace approximation; models are fit using `optimx`’s `nlminb` algorithm (Nash and Varadhan 2011) that returned noticeably better performance. `brms` generates a model that can be manually adapted to place an Inverse Wishart prior as this is not permitted in the default options in pre-written `STAN` models at the time of writing (e.g. `rstanarm` or `brms`).

Draw the fixed effects $\beta \sim N(\mathbf{0}, [0.2]^2 \mathbf{I}_{10})$

For each group $g \in \{1, \dots, 10\}$, draw the random intercept $\alpha_{1,g} \sim N(0, 1)$

For each group $g' \in \{1, \dots, 10\}$, draw the random intercept $\alpha_{2,g'} \sim N(0, 1)$

For each observation $i \in \{1, \dots, 1000\}$, assign at random to groups g, g' . Draw its fixed effect $\mathbf{x}_i \sim N(\mathbf{0}, \Sigma)$ where $\Sigma_{j,j'} = 0.5^{|j-j'|}$; $j, j' \in \{1, \dots, 10\}$. Draw y_i such that:

$$y_i \sim \text{Bern}(p_i), \quad p_i = \frac{\exp(\mathbf{x}_i^T \beta + \alpha_{1,g[i]} + \alpha_{2,g'[i]})}{1 + \exp(\mathbf{x}_i^T \beta + \alpha_{1,g[i]} + \alpha_{2,g'[i]})}$$

All models are fit with a standard Inverse Wishart prior of $\text{IW}(d_j + 1; \mathbf{I}_{d_j})$ on Σ_j . I run each variational algorithm until the change in the ELBO is less than 10^{-8} or the largest parameter changes by less than 10^{-5} . For the HMC and MAVB methods, I draw 4,000 samples from the (approximate) posterior.

Table 1 reports four measures of performance; the first two measures compare the point estimates (posterior mean) against HMC. The third measure compares the full posterior using a measure of “accuracy” that modifies the integrated absolute error (e.g. [Faes et al. 2011](#)). Formally, this is defined as $1 - \frac{1}{2} \int_{-\infty}^{\infty} |q_k(\theta) - q_{\text{HMC}}(\theta)| d\theta$. I use kernel density estimation with a range over the shared support of the samples (`bkde`, `KernSmooth`; [Wand and Ripley 2020](#)) and then approximate the integral. Finally, to understand how the estimates of uncertainty fare against the unknown truth, I examine the “frequentist coverage”: Does an interval of ± 1.96 times the standard deviation of the parameter contain the truth? A value of around 0.95 would indicate correct coverage at the expected frequentist level.

The results are promising; looking at the bias and RMSE, the variational methods perform well; they have very small bias against the means estimated from HMC and an RMSE that is quite small, comparable to the Laplace approximation, and out performs both ADVI implementations.

Examining accuracy and frequentist coverage shows more separation across the methods. The accuracy and coverage of Scheme I are noticeably lower than the Laplace approximation. However, applying MAVB results in noticeable improvements in accuracy (around 6% for fixed effects; and nearly 25% for random effects) and increases coverage by similar amounts to near nominal levels. After this improvement, Scheme I is comparable to the best approximate method (Laplace approximation) having slightly lower accuracy for the fixed effects but noticeably better accuracy and coverage for the random effects. Scheme II has somewhat better initial performance but is also boosted considerably by applying MAVB.

Scheme III—the factorization that does not assume independence between $q(\alpha)$ and $q(\beta)$ —performs nearly as well as the Laplace approximation (and better in terms

Table 1: Results from Simulations

		Bias		RMSE		Accuracy		Coverage	
		FE	RE	FE	RE	FE	RE	FE	RE
	Laplace	-0.000	0.005	0.007	0.056	0.966	0.847	0.950	0.878
	HMC							0.949	0.960
	ADVI (MF)	-0.002	0.003	0.041	0.074	0.775	0.781	0.796	0.830
	ADVI (FR)	-0.000	0.001	0.048	0.108	0.921	0.881	0.969	0.948
	Scheme I	-0.000	0.004	0.007	0.034	0.870	0.706	0.862	0.736
	Scheme II	-0.000	0.004	0.007	0.030	0.870	0.844	0.862	0.871
	Scheme III	-0.000	0.004	0.007	0.026	0.936	0.948	0.923	0.934
MAVB +	Scheme I	-0.000	0.004	0.008	0.034	0.932	0.948	0.922	0.938
	Scheme II	-0.001	0.005	0.007	0.031	0.933	0.955	0.916	0.940
	Scheme III	-0.000	0.004	0.007	0.027	0.937	0.963	0.922	0.942

Note: This reports the bias (Bias), root mean squared error (RMSE) of the estimated posterior means against those estimated from HMC. The distance between the distributions (Accuracy) and frequentist coverage (Coverage) are reported; see the main text for an explanation of these measures. The statistics are disaggregated by fixed (FE) and random effects (RE). All results are created using all relevant parameters in each simulation and then averaged across one hundred simulations. ADVI (MF) uses the mean-field approximation; ADVI (FR) uses the full rank approximation in [Kucukelbir et al. \(2017\)](#).

of the random effects) before applying MAVB. Applying MAVB results in only slight improvements (e.g. a 1-2% boost in accuracy and coverage for the random effects).

Appendix D conducts additional simulations. First, I vary the magnitude of the true coefficients by changing the variance of the fixed and random effects. After applying MAVB, the coverage of the variational methods is near nominal (i.e. above 0.90) in all cases except when the variance of the true distribution of the fixed effects is larger where MAVB is insufficient to obtain nominal coverage on the fixed effects (0.80-0.85) although the coverage on the random effects remains good. While this is worthy of future exploration, I conjecture this occurs because of the large magnitudes of the linear predictors (with 5-95% interval of around -5.9 to 5.5 vs. -2.7 and 2.3 in the simulations in Table 1) and the highly bimodal distribution of p_i . It may be that a pass over the observed data and one full sweep of MCMC (discussed in Section 3 as a “stronger” MAVB) could result in more significant improvements in coverage.

Second, to examine simulations in a more realistic case, I fit a simple MRP model on the data from [Ghitza and Gelman \(2013\)](#) with random effects for age, income, ethnicity and state (Model 1 from Table 2, below) and take the parameter estimates from the Laplace approximation as “ground truth” to create simulated outcomes. It shows a similar pattern although with weaker performance across the board—Scheme I after applying MAVB outperforms ADVI (Mean Field) across all measures with noticeable improvements in accuracy (10%) and is comparable to ADVI (Full Rank). Scheme III performs the best of all approximate methods, including beating the Laplace and ADVI (Full Rank). The values of the linear predictor are relatively modest in this case (90%

of the HMC posterior means are between -0.67 and 1.99) and more comparable to those in the simulations in Table 1. This provides further evidence that for reasonably sized linear predictors, the variational approximations perform well and are improved by MAVB.

Finally, I examine the sensitivity of the algorithm to initial values to see if there is evidence of arriving at different local optima. I find little evidence of this for the models considered in this paper given reasonable random initializations, although researchers should check for this in their own applications.

5 Application: Estimation for Complex MRP

This section re-analyses the results in [Ghitza and Gelman \(2013\)](#) where I compare my results against Hamiltonian Monte Carlo (HMC). I then conduct 10-fold cross-validation using Scheme I to examine which model seems to be most appropriate to use for the final predictive task. I find that, contrary to the decision in [Ghitza and Gelman \(2013\)](#), a model with intermediate complexity is preferred.

5.1 Brief Explanation of MRP

Before proceeding, I provide a brief explanation of MRP (see, e.g. [Park et al. 2004](#); [Lax and Phillips 2009b](#); [Ghitza and Gelman 2013](#) for more detailed explanations). The key problem is that while it is easy to gather a representative survey at the national level, it is very expensive to gather a sufficiently large and representative survey at sub-national units (e.g. states) or sub-types of respondents (e.g. by race, education, income, their interactions, etc.). Further, the number of observations in any sub-group may be very small, rendering a direct analysis of their values unreliable ([Lax and Phillips 2009b](#); [Warshaw and Rodden 2012](#); [Buttice and Highton 2013](#)). However, the most substantively important questions exactly rely on drawing inferences about those sub-groups. MRP provides a model-based procedure to attempt to reliably estimate these sub-group effects by providing a principled way to extrapolate the nationally representative survey.

MRP is a two-step procedure. First, the researcher estimates a hierarchical model (“multilevel regression”) on the initial survey including covariates such as demographic characteristics and indicators for the relevant geographic unit (e.g. state) to get estimates for various “types” of respondents (e.g. age-income-ethnicity by state). The hierarchical model usually has a binomial or binary outcome. The second step calculates the expected response for each demographic-state profile. These can be examined directly or aggregated to get a measure of opinion at the desired geographic level (e.g. state). The aggregation or “post-stratification” occurs by taking a weighted average of those sub-group predictions from the known joint distribution in the population from some ground truth such as the Census. This paper has focused on the first step (“multilevel regression”).

[Ghitza and Gelman \(2013\)](#) apply this method to explore the decision to turn out to vote and party choice by age-race-income-state sub-groups in the 2004 and 2008

American presidential elections. They note that traditional MRP includes the random effects linearly and thus may be failing to capture important complexities or interactions between demography and geography. They thus fit a highly complex model with eighteen random effects and nearly 4,000 parameters on a dataset with around 4,000 observations. After doing so, they draw a variety of subtle and nuanced conclusions about the behavior of particular demographic sub-groups. For example, they qualify the conventional wisdom to show that turnout increases were concentrated amongst *non-white* younger voters instead of younger white voters (Ghitza and Gelman 2013, p. 771-772).

5.2 Estimating Complex Hierarchical Models

I begin by performing a direct comparison of Schemes I, II, and III against the gold standard approaches applied to Ghitza and Gelman (2013). To illustrate the many specifications available to the researcher, Table 2 shows nine possible specifications ranging from a simple MRP model with no interactions to the preferred model in Ghitza and Gelman (2013) (Model 9). I round y_i and n_i to the nearest integer to facilitate interpretation as a standard binomial regression. The intermediate models represent varying complexities that allow for some, but not all, interactions.

It clearly shows the scale of the difficulty for applied researchers: Fitting the published model (Model 9) takes hours using either specification on a machine similar to that available for many applied researchers (a Microsoft Azure instance; Ubuntu, 4 cores, 16 GB of RAM). Methods that require fitting the model repeatedly to facilitate common tasks as bootstrapping, model comparison via cross-validation, or ensemble analysis (Van der Laan et al. 2007, see Ornstein 2020 for an application to MRP) are clearly prohibitively expensive for all except the simplest models using the Laplace approximation.

Figure 1 illustrates the improvement after applying the variational algorithms to each model in Table 2 and performing MAVB. All reported times include estimation of the variational algorithm and drawing 4,000 samples using MAVB. Appendix E shows the time for estimation and MAVB separately; it takes around thirty seconds for Scheme I on the most complex model.

As shown on a linear scale, the time to estimate either the Laplace approximation or Hamiltonian Monte Carlo dwarfs that of any of the variational schemes. The right panel shows the results on a log-scale to allow for clearer comparisons; it shows that Scheme I remains remarkably fast estimating even Model 9 in around one minute versus hour(s) for either gold standard method. The performance of Schemes II and III degrade somewhat—taking around fifteen minutes to fit. This is still very reasonable, but may still be onerous if repeated fitting is required as in cross-validation.

The quality of the approximation is also crucial to assess. As the truth is unknown, I do this by comparing all methods against HMC as this seeks most directly to sample the posterior.⁴ Figure 2 begins by comparing the point estimates pooling across the

⁴This method is, of course, itself approximate as it may fail to accurately sample the posterior.

Table 2: Nine Possible Models for Predicting Turnout via MRP


	Model								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
State	◇	●	●	●	●	●	●	●	●
Age	◇	●	●	●	●	●	●	●	●
Eth	◇	●	●	●	●	●	●	●	●
Inc	◇	◇	◇	◇	◇	◇	◇	◇	◇
Region					●	●	●	●	●
State * Age				◇	◇	◇	◇	◇	◇
State * Eth				◇	◇	◇	◇	◇	◇
State * Inc				◇	◇	◇	◇	◇	◇
Eth * Age			◇	◇	◇	◇	◇	◇	◇
Eth * Inc			◇	◇	◇	◇	◇	◇	◇
Inc * Age			◇	◇	◇	◇	◇	◇	◇
Region * Age					◇	◇	◇	◇	◇
Region * Eth					◇	◇	◇	◇	◇
Region * Inc					◇	◇	◇	◇	◇
State * Eth * Age								◇	◇
State * Eth * Inc							◇	◇	◇
State * Inc * Age									◇
Eth * Inc * Age						◇	◇	◇	◇
<i>Number of Parameters</i>									
	74	139	198	864	945	1026	2047	2864	3885
<i>Run Time of Model in Minutes</i>									
Laplace - 2004	0.2	0.9	2.2	5.1	13.4	27.0	38.4	50.1	81.1
Laplace - 2008	0.2	0.8	2.2	5.5	12.1	23.9	37.0	52.5	84.9
HMC - 2004	113.8	131.5	166.9	199.9	326.6	288.6	375.4	430.0	469.0
HMC - 2008	101.5	132.4	174.1	196.6	353.8	327.3	307.9	402.5	463.8

Note: This table summarizes nine possible models to predict voter turnout. All models include six fixed effects: an intercept, (standardized) individual income, state-level income, state-level Republican vote share and the interaction between individual income and the latter two variables. [Ghitza and Gelman \(2013\)](#) use Model 9. The first panel indicates which random effects are included; a hollow diamond (◇) indicates that only a random intercept is used. A solid circle (●) indicates that a random intercept and a random slope allowing for the effect of (standardized) individual income to vary by group are included. The number of parameters is the number of fixed effects, random effects, and variance components for the random effects. The run times are for a Laplace approximation using `blme` ([Bates et al. 2015](#); [Chung et al. 2015](#)) and HMC in `STAN` (via `brms`; [Bürkner 2017](#)). All models were run on an instance with 16 GB of memory and 4 cores. HMC was estimated using four chains distributed in parallel.

18 models. As there are thousands of parameters to plot, I simplify the picture in the following way; I plot the absolute magnitude of the estimates averaged across j : $\bar{\alpha}_j = \frac{1}{g_j} \sum_{g=1}^{g_j} |\alpha_{j,g}|$ in solid circles and shade the background of the plot based on the density of the individual $|\alpha_{j,g}|$. This prevents the domination of the j with smaller numbers of groups (e.g. age, income, etc.) in the visualization. I also separately mark

Experiments suggested that setting “adapt delta” to 0.99 was required to eliminate all divergent transitions (except for one in Model 3 in 2004).

Figure 1: Speed of Estimation



Note: Each figure plots the run-time of each of the five methods (Laplace approximation, Hamiltonian Monte Carlo [HMC], Schemes I-III with drawing 4,000 samples using MAVB). The reported times are averaged across the 2004 and 2008 elections. The left figure shows the time in minutes on a linear scale; the right figure reports the same information on a log-scale. Model 1-9 are described in Table 2. All models are fit on a computer with 16 GB of RAM and 4 cores.

the fixed and random effects.


Consider first the Laplace approximation; it nearly exactly recovers the point estimates—its solid points and shading lie very near to the 45-degree line. For the variational methods, Scheme I is highly correlated with the posterior ($\rho = 0.996$ for $\bar{\alpha}_j$; $\rho = 0.964$ for the raw $|\alpha_{j,g}|$) although less so than the Laplace approximation. Schemes II and III show tight coupling with the estimates from HMC and are effectively equally accurate to the Laplace approximation. This matches the conventional wisdom that variational methods typically well-recover the posterior means.

Figure 3 presents an analogous figure for the posterior variability, plotting the standard deviation of each parameter. It smooths across random effects in the same way as Figure 2. In interpreting this figure, note that points in the upper left quadrant (above the 45-degree line) indicate worrying performance as the posterior variability is below that coming from the HMC estimates.

Again consider first the Laplace approximation; the standard deviation of its point estimates are often tightly clustered near the 45-degree line but there are a number of random effects that are noticeably smaller (above the 45-degree line).

The performance for the variational algorithms is rather mixed, by comparison. Looking at Scheme I, almost all points show a too small standard deviation—with many random effects being considerably too small. Scheme III, however, improves the

Figure 2: Comparing Posterior Means



Note: This figure plots the absolute value of estimated mean value from Schemes I-III and the Laplace approximation on the horizontal axis against the absolute value of the posterior mean from Hamiltonian Monte Carlo [HMC] on the vertical axis. Each parameter is plotted as a thin grey point; the average of the values inside each random effect are shown as larger points. The axes are on a square-root scale.


situation markedly. While slightly smaller—especially for points with large standard deviations—it tracks the 45-degree line closely and has better performance than the Laplace approximation. As expected, Scheme II is somewhat of an intermediate case; improving some parameters but still having significant problems.

Overall, therefore, Schemes I and II fall into the usual problem of understating posterior variance. By contrast, Scheme III appears to do rather well and lacks the obvious problems of lack of posterior variability versus a fully Bayesian baseline. This corroborates results from [Menictas et al. \(2019\)](#) that estimating $q(\beta, \alpha)$ jointly performs well for (linear) hierarchical models with $J > 1$.

Finally, I show how these estimates change when using MAVB. I focus on the effect on posterior variability as the means are not materially affected by MAVB; [Appendix E](#) shows the analogous figure. [Figure 4](#) presents the distribution of the gap between the variability between the HMC estimates and the other methods where negative values indicates a smaller standard deviation for the competitor methods. Any point below the dotted line indicates that that percentile of effects has a smaller standard deviation than the HMC estimates. To make results interpretable, I report the percentage gap, e.g. $(\text{sd}_k^{\text{Laplace}} - \text{sd}_k^{\text{HMC}}) / \text{sd}_k^{\text{HMC}} \cdot 100$ for all parameters k in (β, α) . To ensure that random effects with small g_j are counted, it presents the averaged statistic across g as in [Figures 2 and 3](#).

The results provide clear evidence for the important role of MAVB. Considering first

Figure 3: Comparing Posterior Variability




Note: This figure plots the estimated standard deviation from Schemes I-III and the Laplace approximation on the horizontal axis against the standard deviation of the posterior distribution from Hamiltonian Monte Carlo [HMC] on the vertical axis. Each parameter is plotted as a thin grey point; the average of the values inside each random effect are shown as larger points.

random effects in the left panel, it is first worth noting that the Laplace approximation—commonly used by researchers—has poor performance for a number of parameter blocks (e.g. the lower percentiles). Scheme I shows a clear lack of variability in the posterior estimates with all estimates being estimated at least 20% too precisely and around *half* of all estimates having less than 75% of the variability estimated in the fully Bayesian setting. After applying MAVB, the (red) solid line shows a considerable improvement although still markedly below the HMC estimates and performing worse than the Laplace approximation. Large improvements are seen for the fixed effects (β) where the estimates of the variability go from extremely poor to being much closer to the Laplace approximation which, itself, is markedly below the HMC coverage.

Scheme III is worth also considering in detail; even before applying MAVB, it has stronger performance than the Laplace approximation in that its curve has a much less poor “tail” (i.e. its worst blocks are around 25% too small vs around 60% for the Laplace approximation). MAVB provides some additional gains ensuring that most parameter blocks are only around 10% too small in terms of their variability. Scheme II is again somewhat intermediate; after applying MAVB, it is broadly comparable to the Laplace approximation.

To provide another interpretation of the role of MAVB, consider the accuracy measure in Section 4 that measures similarity between two distributions, averaged within and then across across parameter blocks: The Laplace approximation performs relatively well (90%). Scheme I performs poorly (43%) because it clearly fails to capture the pos-

Figure 4: Improvements from MAVB



Note: This figure plots the percentile of the gap between the standard deviations estimated via Hamiltonian Monte Carlo [HMC] and the approximate methods. The percentage gap, i.e. $(\text{sd}_k^{\text{Laplace}} - \text{sd}_k^{\text{HMC}}) / \text{sd}_k^{\text{HMC}} \cdot 100$, is shown. A negative value on the vertical axis indicates that the corresponding percentile has a smaller variance than HMC. A vertical shift upward of the line indicates the variance of the parameters has increased. The solid markers indicate the deciles and extremes of the distribution. The dashed line with hollow triangles represents the estimates without using MAVB. The red line with solid circles represents the results after using MAVB.

terior variance. MAVB increases this considerably (68%) although it still falls below the Laplace approximation. Scheme III, however, out-performs the Laplace approximation (95%) with a slight improvement from MAVB (97%).

Appendix E provides some additional results. First, it breaks apart Figure 3 by the type of random effect; the main implication is that the initial lack of variability from Scheme I is most pronounced for the fixed effects and random effects with small g_j (age, ethnicity, income). The improvements for MAVB for those random effects are large and resolves the much of the negative gap.

Second, it examines the linear predictor (i.e. $\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{z}_i^T \boldsymbol{\alpha}$). It shows that MAVB has little effect, although all schemes perform well. In addition to closely estimating the posterior mean (Scheme I has a bias of -0.002 vs HMC), the standard deviation is also fairly close (bias of -0.013 or about -2%), especially compared to the gaps seen in Figure 4. A conjecture would be that MAVB as implemented here has little impact on the linear predictor as it is more about building correlations between parameter blocks, but the “stronger” MAVB noted above might address such limitations.

5.3 Choosing an Optimal Model

Finally, I return to the substantive analysis in [Ghitza and Gelman \(2013\)](#). A key question when performing MRP is the complexity of the accompanying model. Even with the regularization implied by the hierarchical effects, it is still possible to over-fit to the survey sample ([Goplerud et al. 2018](#)). The reported analysis relies on Model 9 without exploring this possibility. The computational burden needed to estimate multiple models and thereby engage in model testing and checking is often onerous for the applied researcher. I thus use the ability to rapidly fit variational approximations to deploy a standard model comparison technique (cross-validation) and examine whether a model of intermediate complexity should be preferred. Table 3 reports a number of statistics on model fit.

Table 3: Cross-Validation to Choose Optimal Model

Method	Models									Time
	1	2	3	4	5	6	7	8	9	
2004 Election										
LOO	11659	11339	11296	11150	11163	11164	11155	11170	11129	2504
WAIC	11658	11336	11292	11107	11115	11113	11083	11074	10968	2503
VI-CV	28.617	28.518	28.508	28.488	28.498	28.499	28.499	28.508	28.508	21
2008 Election										
LOO	11651	11270	11236	11095	11112	11118	11113	11111	11094	2463
WAIC	11651	11268	11233	11060	11073	11075	11053	11030	10957	2462
VI-CV	26.963	26.846	26.839	26.813	26.817	26.821	26.824	26.829	26.820	20

Note This table reports statistics for model fit. The first two rows for each election report fit statistics on the model estimated via Hamiltonian Monte Carlo that approximate cross-validation; the “LOO” information criterion and the WAIC information criterion ([Gelman et al. 2014](#); [Vehtari et al. 2017](#)). The third row reports the average out-of-sample deviance from a model fit using Scheme I. For all statistics, smaller is better and the best value is bolded. The time in minutes for each row to be estimated is shown in the final column; this includes estimation time and the time needed to estimate the relevant fit statistic.

The first two rows (LOO and WAIC) are popular tools for deciding between non-nested Bayesian models ([Vehtari et al. 2017](#)). Details on their exact calculation can be found in the relevant articles ([Gelman et al. 2014](#); [Vehtari et al. 2017](#)), but both are designed to be approximations to cross-validation that do not require fitting the Bayesian model repeatedly.

Fortunately, both have diagnostics to assess whether the underlying approximations are reliable; unfortunately, the diagnostics tests fail in this setting. Almost all models report unacceptable violations of the underlying assumptions for both the LOO and WAIC, and the associated software explicitly encourages the user to resort to K -fold cross-validation. On the other hand, variational inference provides a fast approximate method. The final row of the table (VI-CV) reports the average deviance (twice the negative log-likelihood) of the held-out predictions after conducting 10-fold cross-validation where observations are allocated to each fold with equal probability using Scheme I. Formally, if observation i has a prediction \hat{p}_i , the individual deviance

is $-2[y_i \ln(\hat{p}_i) + (n_i - y_i) \ln(1 - \hat{p}_i)]$. Observations with $n_i = 0$ are excluded from the reported average. Model 4 is also selected if Schemes II or III are used.

The results are interesting and push against the decision to use Model 9; it finds that while Model 1 performs noticeably worse than all other models, it is not necessarily best to use the most complex model. Indeed, an intermediate model—Model 4—performs the best although the differences in the error are quantitatively small between Models 4 and 9.

Appendix E provides a more detailed exploration against a “Bayesian gold standard.” Using a new set of folds, it fits Models 1, 4, and 9 using ten-fold cross-validation and performing HMC and the Laplace approximation on each fold. This is extremely time intensive—taking around *ten* days to complete the whole process. It confirms that cross-validated HMC, Laplace approximation, and Scheme I all select Model 4. The out-of-sample predictions between Scheme I and HMC are highly correlated (0.998). This gives some confidence that the results of the variational method can be used in lieu of prohibitively expensive classical cross-validation. When it is too expensive to conduct such an analysis, relying on methods such as simulation-based calibration (e.g. Yao et al. 2018) may be a feasible way to assess whether the variational approximation “successfully” approximated the posterior.

Returning to Table 2, the major feature that distinguishes Model 4 from less complex models is interactions between the core random effects (age, ethnicity, income) and state. This matches a reasonable expectation from political science that demographics are likely to vary across state but the complex higher-order interactions between region and three-way-interactions do not seem to add much predictive power.

These results are useful to practitioners of MRP in three ways; first, complex hierarchical models can now be compared against other state-of-the-art machine learning methods versus relying a very simple model (analogous to Model 1) due to computational costs (Bisbee 2019; Ornstein 2020). Thus, it is an interesting and open question whether methods such as BART are actually superior for MRP tasks (Bisbee 2019) or whether properly specified complex hierarchical models can be competitive. Second, it suggests that interactions between demographics and state characteristics are important to include although the evidence for going extremely “deep” and adding many higher-order interactions appears more limited. Finally, even if one prefers to fit a Bayesian model for the final regression, the ability to quickly search between models allows the researcher to narrow down a set of plausible candidate models for final exploration and model testing.

6 Conclusion

This paper provided a new set of variational algorithms that, leveraging Polya-Gamma data augmentation (Polson et al. 2013), require only a mean-field assumption to estimate a logistic hierarchical regression with an arbitrary number and size of random effects. It provided multiple factorization assumptions; Scheme I required the independence of the fixed effects and each block of random effects whereas Scheme III relaxed that

assumption at the expense of increased computational cost. All methods seemed to quite accurately capture the posterior means in even complex models. As expected in both simulations and real data, Scheme I performed worse—especially in terms of understating posterior variance for many random effects.

The paper also provided a generic way to improve the performance of Scheme I, and Schemes II and III to a lesser extent. By leveraging the existence of a parameter expansion of the underlying model either by allowing the means of the random effects to be non-zero or by imposing some translation, one can use a marginal augmentation sampler to improve the posterior approximation. This procedure (“marginally augmented variational Bayes”; MAVB) showed promising performance when applied to Scheme I: It increased the variance of the estimated approximations to be closer to the samples drawn using a fully Bayesian procedure, although still remaining too small on real data. However, given its speed even on complex models, MAVB provides a cheap way to make Scheme I a more viable approximation to the true posterior. It is also worth noting that Scheme III performed very well—often beating the very popular Laplace approximation on both real and simulated data.

Future work could proceed in at least two directions. First, the algorithms here can be naturally extended to count and multinomial outcomes providing a more unified approach to variational estimation of non-linear hierarchical models. Extending the model to include a weakly informative prior such as [Huang and Wand \(2013\)](#) is also an important extension.

Second, the usefulness of MAVB should be explored both theoretically and in the context of other models. As noted earlier, there is nothing about using MAVB that is specific to logistic hierarchical models *per se*. Indeed, this idea of “improving” an approximation by pushing it through a Markov transition kernel can be generalized to a wide variety of MCMC samplers and models. It thus opens a question of which Markov transition density to use for other models that do not admit marginal augmentation. A reasonable conjecture is that as the mixing of the sampler improves, the transformed sample will be closer to the true posterior.

Supplemental Materials

The supporting information contains derivations of the variational algorithm (Appendix A), formal definitions and proofs of MAVB (Appendix B), results on accelerating CAVI using PX-VB and joint updates of certain parameters (Appendix C), additional simulations (Appendix D), and additional analyses on [Ghitza and Gelman \(2013\)](#) (Appendix E).

Open-source statistical software to implement the algorithms in this paper is available on GitHub as noted in the acknowledgements. Materials to replicate the analyses in the paper can be found at the following link: <https://doi.org/10.7910/DVN/DI191B>.

References

- 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.
- Bell, A. and Jones, K. (2015). “Explaining Fixed Effects: Random Effects Modeling of Time-Series Cross-Sectional and Panel Data.” *Political Science Research and Methods*, 3(1): 133–153.
- Bisbee, J. (2019). “BARP: Improving Mister P Using Bayesian Additive Regression Trees.” *American Political Science Review*, 113(4): 1060–1065.
- Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). “Variational Inference: A Review for Statisticians.” *Journal of the American Statistical Association*, 112(518): 859–877.
- Bürkner, P.-C. (2017). “brms: An R Package for Bayesian Multilevel Models Using Stan.” *Journal of Statistical Software*, 80(1): 1–28.
- Buttice, M. K. and Highton, B. (2013). “How Does Multilevel Regression and Post-stratification Perform with Conventional National Surveys?” *Political Analysis*, 21(4): 449–467.
- 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.
- Chung, Y., Gelman, A., Rabe-Hesketh, S., Liu, J., and Dorie, V. (2015). “Weakly Informative Prior for Point Estimation of Covariance Matrices in Hierarchical Models.” *Journal of Educational and Behavioral Statistics*, 40(2): 136–157.
- Clark, T. S. and Linzer, D. A. (2015). “Should I Use Fixed or Random Effects?” *Political Science Research and Methods*, 3(2): 399–408.
- Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. John Wiley & Sons.
- Faes, C., Ormerod, J. T., and Wand, M. P. (2011). “Variational Bayesian Inference for Parametric and Nonparametric Regression with Missing Data.” *Journal of the American Statistical Association*, 106(495): 959–971.
- Gao, Y., Kennedy, L., Simpson, D., and Gelman, A. (2020). “Improving Multilevel Regression and Poststratification with Structured Priors.” *Bayesian Analysis*, Advanced Access: doi 10.1214/20-BA1223.
- Gelman, A. and Hill, J. (2006). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.
- Gelman, A., Hwang, J., and Vehtari, A. (2014). “Understanding Predictive Information Criteria for Bayesian Models.” *Statistics and Computing*, 24(6): 997–1016.
- Gelman, A., Lax, J., Phillips, J., Gabry, J., and Trangucci, R. (2016). “Using Multilevel

- Regression and Poststratification to Estimate Dynamic Public Opinion.” *Unpublished manuscript*.
- Gelman, A. and Little, T. C. (1997). “Poststratification Into Many Categories Using Hierarchical Logistic Regression.” *Survey Methodology*, 23: 127—135.
- Gelman, A., van Dyk, D. A., Huang, Z., and Boscardin, J. W. (2008). “Using Redundant Parameterizations to Fit Hierarchical Models.” *Journal of Computational and Graphical Statistics*, 17(1): 95–122.
- Gerchinovitz, S., Ménard, P., and Stoltz, G. (2020). “Fano’s Inequality for Random Variables.” *Statistical Science*, 35: 178–201.
- Ghitza, Y. and Gelman, A. (2013). “Deep Interactions with MRP: Election Turnout and Voting Patterns Among Small Electoral Subgroups.” *American Journal of Political Science*, 57(3): 762–776.
- Giordano, R. J., Broderick, T., and Jordan, M. I. (2015). “Linear Response Methods for Accurate Covariance Estimates from Mean Field Variational Bayes.” In *Neural Information Processing Systems 2015*.
- Goplerud, M., Kuriwaki, S., Ratkovic, M., and Tingley, D. (2018). “Sparse Multilevel Regression (and Poststratification [sMRP]).” *Unpublished manuscript*.
- Hall, P., Johnstone, I. M., Ormerod, J. T., Wand, M. P., and Yu, J. C. (2019). “Fast and Accurate Binary Response Mixed Model Analysis via Expectation Propagation.” *Journal of the American Statistical Association*, 1–15.
- Hall, P., Ormerod, J. T., and Wand, M. P. (2011). “Theory of Gaussian Variational Approximation for a Poisson Mixed Model.” *Statistica Sinica*, 21(1): 369–389.
- Huang, A. and Wand, M. P. (2013). “Simple Marginally Noninformative Prior Distributions for Covariance Matrices.” *Bayesian Analysis*, 8(2): 439–452.
- Jaakkola, T. S. and Qi, Y. (2007). “Parameter Expanded Variational Bayesian Methods.” In *Neural Information Processing Systems 2007*.
- Jeon, M., Rijmen, F., and Rabe-Hesketh, S. (2017). “A Variational Maximization–Maximization Algorithm for Generalized Linear Mixed Models with Crossed Random Effects.” *Psychometrika*, 82(3): 693–716.
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., and Blei, D. M. (2017). “Automatic Differentiation Variational Inference.” *Journal of Machine Learning Research*, 18(14): 1–45.
- Lax, J. R. and Phillips, J. H. (2009a). “Gay Rights in the States: Public Opinion and Policy Responsiveness.” *American Political Science Review*, 103(3): 367–386.
- (2009b). “How Should We Estimate Public Opinion in The States?” *American Journal of Political Science*, 53(1): 107–121.
- (2012). “The Democratic Deficit in the States.” *American Journal of Political Science*, 56(1): 148–166.

- Liu, C., Rubin, D. B., and Wu, Y. N. (1998). “Parameter Expansion to Accelerate EM: The PX-EM Algorithm.” *Biometrika*, 85(4): 755–770.
- Liu, J. S. and Wu, Y. N. (1999). “Parameter Expansion for Data Augmentation.” *Journal of the American Statistical Association*, 94(448): 1264–1274.
- Menictas, M., Di Credico, G., and Wand, M. P. (2019). “Streamlined Variational Inference for Linear Mixed Models with Crossed Random Effects.” *arxiv preprint: 1910.01799*.
- Nash, J. C. and Varadhan, R. (2011). “Unifying Optimization Algorithms to Aid Software System Users: optimx for R.” *Journal of Statistical Software*, 43(9): 1–14.
- Ormerod, J. T. and Wand, M. P. (2012). “Gaussian Variational Approximate Inference for Generalized Linear Mixed Models.” *Journal of Computational and Graphical Statistics*, 21(1): 2–17.
- Ornstein, J. T. (2020). “Stacked Regression and Poststratification.” *Political Analysis*, 28(2): 293–301.
- Park, D. K., Gelman, A., and Bafumi, J. (2004). “Bayesian Multilevel Estimation with Poststratification: State-Level Estimates from National Polls.” *Political Analysis*, 12(4): 375–385.
- Polson, N. G., Scott, J. G., and Windle, J. (2013). “Bayesian Inference for Logistic Models Using Pólya–Gamma Latent Variables.” *Journal of the American Statistical Association*, 108(504): 1339–1349.
- Rabe-Hesketh, S. and Skrondal, A. (2008). *Multilevel and Longitudinal Modeling using STATA*. STATA Press.
- Rabe-Hesketh, S., Skrondal, A., and Pickles, A. (2004). “Generalized Multilevel Structural Equation Modeling.” *Psychometrika*, 69(2): 167–190.
- Ruiz, F. J. and Titsias, M. K. (2019). “A Contrastive Divergence for Combining Variational Inference and MCMC.” In *International Conference on Machine Learning*.
- Salimans, T., Kingma, D., and Welling, M. (2015). “Markov Chain Monte Carlo and Variational Inference: Bridging the Gap.” In *International Conference on Machine Learning*.
- Steenbergen, M. R. and Jones, B. S. (2002). “Modeling Multilevel Data Structures.” *American Journal of Political Science*, 46(1): 218–237.
- Stegmüller, D. (2013). “How Many Countries for Multilevel Modeling? A Comparison of Frequentist and Bayesian Approaches.” *American Journal of Political Science*, 57(3): 748–761.
- Tan, L. S. and Nott, D. J. (2013). “Variational Inference for Generalized Linear Mixed Models using Partially Noncentered Parametrizations.” *Statistical Science*, 28(2): 168–188.
- Tan, L. S. L. (2021). “Use of Model Reparametrization to Improve Variational Bayes.”

- Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(1): 30–57.
- Tausanovitch, C. and Warshaw, C. (2014). “Representation in Municipal Government.” *American Political Science Review*, 108: 605–641.
- Van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). “Super Learner.” *Statistical Applications in Genetics and Molecular Biology*, 6(1).
- Van Dyk, D. A. and Meng, X.-L. (2001). “The Art of Data Augmentation.” *Journal of Computational and Graphical Statistics*, 10(1): 1–50.
- Vehtari, A., Gelman, A., and Gabry, J. (2017). “Practical Bayesian Model Evaluation using Leave-One-Out Cross-Validation and WAIC.” *Statistics and Computing*, 27(5): 1413–1432.
- Wand, M. and Ripley, B. (2020). “KernSmooth: Functions for kernel smoothing for Wand & Jones (1995).” *R package version*, 2.23-18.
- Warshaw, C. and Rodden, J. (2012). “How Should We Measure District-Level Public Opinion on Individual Issues?” *The Journal of Politics*, 74(1): 203–219.
- Yao, Y., Vehtari, A., Simpson, D., and Gelman, A. (2018). “Yes, but Did It Work? Evaluating Variational Inference.” In *International Conference on Machine Learning*.
- Yin, M. and Zhou, M. (2018). “Semi-Implicit Variational Inference.” In *International Conference on Machine Learning*.
- Zhao, Y., Staudenmayer, J., Coull, B. A., and Wand, M. P. (2006). “General Design Bayesian Generalized Linear Mixed Models.” *Statistical Science*, 21(1): 35–51.

Appendix A: Derivation of Variational Algorithms


A.1 Alternative Notations: Gelman and Hill and Plate Diagram

Equation A.5 expresses the generative model in Equation 1 using the popular notation in Gelman and Hill (2006) where $\alpha_{j,g[i]}$ denotes the random effect j for group g of which i is a member.

$$\begin{aligned}
 y_i | \beta, \{\{\alpha_{j,g}\}_{g=1}^{g_j}\}_{j=1}^J &\sim \text{Binom}(n_i, p_i), \quad p_i = \frac{\exp(\psi_i)}{1 + \exp(\psi_i)}, \quad \psi_i = \mathbf{x}_i^T \beta + \sum_{j=1}^J \mathbf{z}_{i,j}^b \alpha_{j,g[i]} \\
 \alpha_{j,g} | \Sigma_j &\sim N(\mathbf{0}_{d_j}, \Sigma_j), \quad \Sigma_j \sim \text{IW}(\nu_j, \Phi_j) \quad \forall (j, g), \quad p(\beta) \propto 1
 \end{aligned}
 \tag{A.5}$$

A plate diagram is shown below: Three arrows from $\mathbf{z}_{i,j}^b$, $\alpha_{j,g}$ and $\mathbf{m}_{i,j}$ intersecting to the right of y_i denote picking out the specific random effect corresponding to observation i , i.e. $\alpha_{j,g[i]}$ in Gelman and Hill (2006) notation.

Figure A.5: Plate Diagram of Hierarchical Model



A.2 Derivation of Variational Algorithms

Equation A.6 provides the log-complete joint density of the model. I assume a flat prior on β in all of the subsequent derivations where Γ_d is the multivariate Gamma function.

$$\begin{aligned} \ln p(\mathbf{y}, \Omega, \beta, \alpha, \{\Sigma_j\}_{j=1}^J) &= \mathbf{s}^T [\mathbf{X}\beta + \mathbf{Z}\alpha] - \frac{1}{2} [\mathbf{X}\beta + \mathbf{Z}\alpha]^T \Omega [\mathbf{X}\beta + \mathbf{Z}\alpha] + \\ &\left(\sum_i \ln f_{PG}(\omega_i | n_i, 0) \right) + \sum_{j=1}^J \left[-(d_j g_j)/2 \ln(2\pi) - g_j/2 \ln(|\Sigma_j|) + \sum_{g=1}^{g_j} -\frac{1}{2} \alpha_{jg}^T \Sigma_j^{-1} \alpha_{jg} \right] + \\ &- \left(\sum_i n_i \right) \ln(2) + \sum_{j=1}^J \ln c(\nu_j, \Phi_j) - (\nu_j + d_j + 1)/2 \ln(|\Sigma_j|) - \frac{1}{2} \text{tr}(\Phi_j \Sigma_j^{-1}) \\ \text{where } \ln c(\nu_j, \Phi_j) &= \nu_j/2 \ln |\Phi_j| - (\nu_j d)/2 \ln(2) - \ln(\Gamma_d(\nu_j/2)) \end{aligned} \quad (\text{A.6})$$

A standard result in variational inference is that, under a mean-field assumption as imposed in Schemes I, II, and III, the optimal approximating distribution is proportional to the exponential of the expectation of Equation A.6 over all other parameters (e.g. Bishop 2006; Blei et al. 2017). From this, the updates can be derived as follows.⁵

- $q(\Omega)$. This factorizes into N independent Polya-Gamma variables:

$$q(\omega_i) \sim PG(\tilde{b}_i, \tilde{c}_i)$$

Under Scheme I and Scheme II, the approximating distribution is as follows.

$$\tilde{b}_i = n_i, \quad \tilde{c}_i = \sqrt{[\mathbf{x}_i^T \tilde{\mu}_\beta + \mathbf{z}_i^T \tilde{\mu}_\alpha]^2 + \mathbf{x}_i^T \tilde{\Lambda}_\beta \mathbf{x}_i + \mathbf{z}_i^T \tilde{\Lambda}_\alpha \mathbf{z}_i}$$

Under Scheme III, this depends on the covariance between α and β . I use $\Lambda_{\beta-\alpha}$ to denote the estimated covariance matrix from the variational approximation on the stacked vector $[\beta^T, \alpha^T]$. This is block diagonal under Schemes I and II.

$$\tilde{b}_i = n_i, \quad \tilde{c}_i = \sqrt{[\mathbf{x}_i^T \tilde{\mu}_\beta + \mathbf{z}_i^T \tilde{\mu}_\alpha]^2 + [\mathbf{x}_i^T, \mathbf{z}_i^T] \tilde{\Lambda}_{\beta-\alpha} \begin{bmatrix} \mathbf{x}_i \\ \mathbf{z}_i \end{bmatrix}}$$

I define the stacked expectation of the Polya-Gammas into a diagonal matrix as follows:

$$\tilde{\Lambda}_\Omega = \text{diag} \left(\frac{\tilde{b}_i}{2\tilde{c}_i} \tanh(\tilde{c}_i/2) \right)$$

⁵Note that a corresponding fully Bayesian Gibbs Sampler can be mostly read-off from these algorithms if one instead samples from the noted variational distribution *and* plugs in the sampled parameters instead of their expectations/variances.

- $q(\Sigma)$. This factorizes into J independent Inverse Wishart distributions whose parameters are as follows where the subscript jj denotes taking the sub vector or matrix corresponding to $\alpha_{j,g}$ that is stacked into α .

$$q(\Sigma_j) \sim \text{IW} \left(\tilde{\nu}_j, \tilde{\Phi}_j \right)$$

$$\tilde{\nu}_j = \nu_j + g_j, \quad \tilde{\Phi}_j = \Phi_j + \sum_{g=1}^{g_j} ([\tilde{\mu}_\alpha]_{jg}) ([\tilde{\mu}_\alpha]_{jg})^T + [\tilde{\Lambda}_\alpha]_{jg,jg}$$

Note that $E_{q(\Sigma_j)}[\Sigma_j^{-1}] = \tilde{\nu}_j [\tilde{\Phi}_j]^{-1}$

- $q(\beta, \alpha)$. The optimal approximating distribution under Scheme I, II, or III is multivariate normal. The exact details depend on the factorization scheme and are each enumerated below. The stability and convergence of the algorithm improves when all of the mean parameters $(\tilde{\mu}_\beta, \tilde{\mu}_{\alpha,j})$ are updated jointly.

– Scheme I:

$$q(\beta) \sim N \left(\tilde{\mu}_\beta, \tilde{\Lambda}_\beta \right), \quad \tilde{\Lambda}_\beta = \left(\mathbf{X}^T \tilde{\Lambda}_\Omega \mathbf{X} \right)^{-1}, \quad \tilde{\mu}_\beta = \tilde{\Lambda}_\beta \mathbf{X}^T \left(\mathbf{s} - \tilde{\Lambda}_\Omega \mathbf{Z} \tilde{\mu}_\alpha \right)$$

Recall that $\mathbf{Z}\alpha = \sum_{j=1}^J \mathbf{Z}_j \alpha_j$. The approximation distributions $q(\alpha_j)$ can be updated cyclically by iterating through $j \in \{1, \dots, J\}$ with each update as follows

$$q(\alpha_j) \sim N \left(\tilde{\mu}_{\alpha,j}, \tilde{\Lambda}_{\alpha,j} \right), \quad \tilde{\Lambda}_{\alpha,j} = \left(\mathbf{Z}_j^T \tilde{\Lambda}_\Omega \mathbf{Z}_j + [\mathbf{I}_{g_j} \otimes E_{q(\Sigma_j)}[\Sigma_j^{-1}]] \right)^{-1}$$

$$\tilde{\mu}_{\alpha,j} = \tilde{\Lambda}_{\alpha,j} \mathbf{Z}_j^T \left(\mathbf{s} - \tilde{\Lambda}_\Omega \mathbf{X} \tilde{\mu}_\beta - \sum_{\ell \in \{1, \dots, J\} \setminus j} \tilde{\Lambda}_\Omega \mathbf{Z}_\ell \tilde{\mu}_{\alpha,\ell} \right)$$

Note that $\tilde{\Lambda}_\alpha$ can be made by stacking $\tilde{\Lambda}_{\alpha,j}$ block diagonally. $\tilde{\mu}_\alpha$ is formed by stacking the $\tilde{\mu}_{\alpha,j}$ vertically.

- Scheme II: The update for $q(\beta)$ is unchanged from Scheme I. The update for $q(\alpha)$ is as follows.

$$q(\alpha) \sim N \left(\tilde{\mu}_\alpha, \tilde{\Lambda}_\alpha \right)$$

$$\tilde{\Lambda}_\alpha = \left(\mathbf{Z}^T \tilde{\Lambda}_\Omega \mathbf{Z} + \text{blockdiag} \left(\{ \mathbf{I}_{g_j} \otimes E_{q(\Sigma_j)}[\Sigma_j^{-1}] \}_{j=1}^J \right) \right)^{-1}$$

$$\tilde{\mu}_\alpha = \tilde{\Lambda}_\alpha \mathbf{Z}^T \left(\mathbf{s} - \tilde{\Lambda}_\Omega \mathbf{X} \tilde{\mu}_\beta \right)$$

- Scheme III: $q(\beta, \alpha)$ is no longer assumed to contain any independent components and is updated jointly.

$$q(\beta, \alpha) \sim N \left(\begin{bmatrix} \tilde{\mu}_\beta \\ \tilde{\mu}_\alpha \end{bmatrix}, \tilde{\Lambda}_{\beta-\alpha} \right), \quad \tilde{\Lambda}_{\beta-\alpha} = \left([\mathbf{X}, \mathbf{Z}]^T \tilde{\Lambda}_\Omega [\mathbf{X}, \mathbf{Z}] + \mathbf{T} \right)^{-1}$$

$$\mathbf{T} = \begin{pmatrix} \mathbf{0}_{p \times p} & \mathbf{0}_{p \times \sum_j d_j g_j} \\ \mathbf{0}_{\sum_j d_j g_j \times p} & \text{blockdiag}(\{\mathbf{I}_{g_j} \otimes E_{q(\Sigma_j)}[\Sigma_j^{-1}]\}_{j=1}^J) \end{pmatrix}$$

$$\begin{bmatrix} \tilde{\mu}_\beta \\ \tilde{\mu}_\alpha \end{bmatrix} = \tilde{\Lambda}_{\beta-\alpha} [\mathbf{X}, \mathbf{Z}]^T \mathbf{s}$$

The ELBO can be expressed as in Equation A.7 and A.8. It can be decomposed into two parts; the expectation of Equation A.6 and the differential entropy of the approximating distribution. The sum of the two terms is the ELBO, i.e. ELBO = LogComplete + Entropy.

$$\begin{aligned} \text{LogComplete} = & - \left(\sum_i n_i \right) \ln(2) + \mathbf{s}^T [\mathbf{X} \tilde{\mu}_\beta + \mathbf{Z} \tilde{\mu}_\alpha] + \\ & - \frac{1}{2} [\mathbf{X} \tilde{\mu}_\beta + \mathbf{Z} \tilde{\mu}_\alpha]^T \tilde{\Lambda}_\Omega [\mathbf{X} \tilde{\mu}_\beta + \mathbf{Z} \tilde{\mu}_\alpha] + \frac{1}{2} \text{tr} \left(\tilde{\Lambda}_\Omega \text{Var}(\mathbf{X} \beta + \mathbf{Z} \alpha) \right) + \\ & \sum_{i=1}^N E_{q(\omega_i)} [\ln f_{PG}(\omega_i | n_i, 0)] + \\ & \sum_{j=1}^J \left[-\frac{1}{2} \left[\sum_{g=1}^{g_j} [\tilde{\mu}_{jg}]^T (E_{q(\Sigma_j)} \Sigma_j^{-1}) \tilde{\mu}_{jg} + \text{tr} \left(E_{q(\Sigma_j)} [\Sigma_j^{-1}] [\tilde{\Lambda}_\alpha]_{jg} \right) \right] \right] + \\ & \sum_{j=1}^J \ln c(\nu_j, \Phi_j) - \frac{(\nu_j + d_j + 1)}{2} E_{q(\Sigma_j)} [\ln(|\Sigma_j|)] - \frac{1}{2} \text{tr} (\Phi_j E_{q(\Sigma_j)} [\Sigma_j^{-1}]) \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} \text{Entropy} = & \frac{1}{2} \ln [2\pi e |\tilde{\Lambda}_{\alpha-\beta}|] + \\ & \sum_{i=1}^N \frac{\tilde{b}_i \tilde{c}_i}{4} \tanh(\tilde{c}_i/2) - E_{q(\omega_i)} [\ln f_{PG}(\omega_i | \tilde{b}_i, 0)] - \tilde{b}_i \ln [\cosh(\tilde{c}_i/2)] + \\ & \sum_{j=1}^J -\ln c(\tilde{\nu}_j, \tilde{\Phi}_j) + \frac{\tilde{\nu}_j + d_j + 1}{2} E_{q(\Sigma_j)} [\ln |\Sigma_j|] + \frac{1}{2} \text{tr} \left(\tilde{\Phi}_j E_{q(\Sigma_j)} [\Sigma_j^{-1}] \right) \end{aligned} \quad (\text{A.8})$$

Note that neither Equation A.8 nor Equation A.7 are individually computable in closed form as each contains an intractable expectation of the log-density of the Polya-

Gamma density as $\tilde{b}_i = n_i$. Fortunately, the terms cancel and thus the ELBO is tractable. Some final intermediate results are necessary to calculate the ELBO.

- The entropy of a Polya-Gamma random variable. The key identity follows from Equation 5 in [Polson et al. \(2013\)](#).

$$E_{q(\omega|b,c)}[-\ln f_{PG}(\omega|b,c)] = E_{q(\omega|b,c)} \left[\frac{c^2\omega}{2} - \ln f_{PG}(\omega|b,0) - b \ln \cosh(c/2) \right]$$

- Some expectations over an Inverse Wishart distribution. The results are stated in [Tan and Nott \(2013\)](#) and follow from noting that if $\Sigma_j \sim \text{IW}(\nu, \Phi)$ then $\Sigma_j^{-1} \sim \text{Wishart}(\nu, \Phi^{-1})$. For $\Sigma_j \sim \text{IW}(\nu_0, \Phi_0)$ where Σ_j is $d \times d$:

$$E_{q(\Sigma_j)}[\Sigma_j^{-1}] = \nu_0[\Phi_0]^{-1}, \quad E_{q(\Sigma_j)}[\ln |\Sigma_j|] = \ln |\Phi_0| - \sum_{k=1}^d \psi((\nu_0 - k + 1)/2) - d \ln(2)$$

Appendix B: Derivation of Parameter Expansions and MAVB

This section contains a number of results on using MAVB. First, I prove Theorem 1. Second, I derive the implementation for a proper working prior ($\mu_j \sim N(\mathbf{0}, \tau^2 \mathbf{I})$) and note that the limiting case as $\tau^2 \rightarrow \infty$ corresponds to the result in the main text.

B.1 Definitions

First, I explicitly characterize the definition of parameter expansion more generally below.

Definition 3 (Parameter Expansion). *Drawing on the results from [Liu and Wu \(1999\)](#) and [Van Dyk and Meng \(2001\)](#), define a parameter expansion as follows: ξ is a d_ξ -dimensional parameter that is defined on some space $S_\xi \subseteq \mathbb{R}^{d_\xi}$. A parameter expansion of the original model $p(\mathbf{y}, \theta)$ is defined via a transformation function t_ξ generating expanded parameters θ^X satisfying the following conditions:*

- *Equivalence: The transformation preserves the likelihood of the observed data*

$$p(\mathbf{y}) = \int p(\mathbf{y}, \theta) d\theta = \int p^X(\mathbf{y}, \theta^X | \xi) d\theta^X \quad \forall \xi \in S_\xi$$

- *Reduction: t_ξ is a one-to-one and differentiable function for $\xi \in S_\xi$ where $t_\xi(\theta^X) = \theta$*

- *Null Value:* There exists some $\xi_{\text{Null}} \in \mathcal{S}_\xi$ such that $t_{\xi_{\text{Null}}}(\theta^X) = \theta^X$

The expanded model has an associated evidence lower bound conditional on ξ as follows:

$$\text{ELBO}_{q(\theta^X)}^{X-\xi} = E_{q(\theta^X)} [\ln p^X(\mathbf{y}, \theta^X | \xi)] - E_{q(\theta^X)} [\ln q(\theta^X)]$$

B.2 Proof of Theorem 1

Theorem 1 can be proved directly following results from [Liu and Wu \(1999\)](#) and a data processing inequality stated by a number of authors. I explicitly restate those here with notation adapted to this paper.

Lemma 1 (Liu and Wu – Theorem 1). *For a transformation function $t_\xi(\theta^X)$ satisfying the conditions in Definition 3, i.e. one-to-one and differentiable for a fixed ξ , the following result holds: Assume that $\mathbf{z} \sim p(\mathbf{z})$ and $\xi \sim p_0(\xi)$. Let ξ_0 be a random draw from the prior $p_0(\xi)$ and define $\mathbf{w} = t_{\xi_0}^{-1}(\mathbf{z})$.*

If $\xi_1 \sim p(\xi | \mathbf{w}) \propto p(t_\xi(\mathbf{w})) |J_\xi(\mathbf{w})| p_0(\xi)$, then $\mathbf{z}' = t_{\xi_1}(\mathbf{w})$ has the same distribution as \mathbf{z} .

Lemma 2 (Data Processing Inequality - Various). *If $\pi(\mathbf{y} | \mathbf{x})$ is a conditional distribution to generate \mathbf{y} given \mathbf{x} , the following identity holds for any choice of $\pi(\mathbf{y} | \mathbf{x})$ assuming that (i) $p(\mathbf{x})$, $q(\mathbf{x})$, $p(\mathbf{y})$, $q(\mathbf{y})$ are all proper densities and (ii) the KL-divergence between $q(\mathbf{x})$ and $p(\mathbf{x})$ is finite.*

$$\text{KL}[q(\mathbf{x}) \parallel p(\mathbf{x})] \geq \text{KL}[q(\mathbf{y}) \parallel p(\mathbf{y})]$$

$$\text{where } q(\mathbf{y}) = \int \pi(\mathbf{y} | \mathbf{x}) q(\mathbf{x}) d\mathbf{x}, \quad p(\mathbf{y}) = \int \pi(\mathbf{y} | \mathbf{x}) p(\mathbf{x}) d\mathbf{x}$$

The proof of the data processing inequality has been noted by a variety of authors (e.g. [Ruiz and Titsias 2019](#) citing [Cover and Thomas 2006](#)). The result also appears to follow from [Gerchinovitz et al. \(2020\)](#)'s result on f -divergences (Lemma 2.1). For completeness, I show a proof via a manipulation of the KL-divergence below where I assume that $\pi(\mathbf{y} | \mathbf{x})$ is a conditional density for sampling \mathbf{y} given \mathbf{x} . This assumes that $p(\mathbf{x})$, $p(\mathbf{y})$, $q(\mathbf{x})$, $q(\mathbf{y})$ are all well-defined and proper. I also assume the initial KL divergence between $p(\mathbf{x})$ and $q(\mathbf{x})$ is finite.

$$\text{KL}[p(\mathbf{x}) \parallel q(\mathbf{x})] = \int \ln \left(\frac{p(\mathbf{x})}{q(\mathbf{x})} \right) p(\mathbf{x}) d\mathbf{x} \tag{A.10a}$$

$$= \int \int \ln \left(\frac{p(\mathbf{x}) \pi(\mathbf{y} | \mathbf{x})}{q(\mathbf{x}) \pi(\mathbf{y} | \mathbf{x})} \right) p(\mathbf{x}) \pi(\mathbf{y} | \mathbf{x}) d\mathbf{x} d\mathbf{y} \tag{A.10b}$$

$$= E_{p(\mathbf{y})} \left[E_{p(\mathbf{x} | \mathbf{y})} \left[\ln \left(\frac{p(\mathbf{x}) \pi(\mathbf{y} | \mathbf{x})}{q(\mathbf{x}) \pi(\mathbf{y} | \mathbf{x})} \right) \right] \right] \tag{A.10c}$$

$$= E_{p(\mathbf{y})} \left[E_{p(\mathbf{x}|\mathbf{y})} \left[-\ln \left(\frac{q(\mathbf{x})\pi(\mathbf{y}|\mathbf{x})}{p(\mathbf{x})\pi(\mathbf{y}|\mathbf{x})} \right) \right] \right] \quad (\text{A.10d})$$

$$\geq E_{p(\mathbf{y})} \left[-\ln \left[E_{p(\mathbf{x}|\mathbf{y})} \left(\frac{q(\mathbf{x})\pi(\mathbf{y}|\mathbf{x})}{p(\mathbf{x})\pi(\mathbf{y}|\mathbf{x})} \right) \right] \right] \quad (\text{A.10e})$$

$$= E_{p(\mathbf{y})} \left[-\ln \left[\int \frac{q(\mathbf{x})\pi(\mathbf{y}|\mathbf{x})}{p(\mathbf{x})\pi(\mathbf{y}|\mathbf{x})} p(\mathbf{x}|\mathbf{y}) d\mathbf{x} \right] \right] \quad (\text{A.10f})$$

$$= E_{p(\mathbf{y})} \left[-\ln \left[\int \frac{q(\mathbf{y})q(\mathbf{x}|\mathbf{y})}{p(\mathbf{y})p(\mathbf{x}|\mathbf{y})} p(\mathbf{x}|\mathbf{y}) d\mathbf{x} \right] \right] \quad (\text{A.10g})$$

$$= E_{p(\mathbf{y})} \left[-\ln \left[\int \frac{q(\mathbf{y})}{p(\mathbf{y})} q(\mathbf{x}|\mathbf{y}) d\mathbf{x} \right] \right] = E_{p(\mathbf{y})} \left[-\ln \left[\frac{q(\mathbf{y})}{p(\mathbf{y})} \right] \right] \quad (\text{A.10h})$$

$$= E_{p(\mathbf{y})} \left[\ln \left(\frac{p(\mathbf{y})}{q(\mathbf{y})} \right) \right] = \text{KL}[p(\mathbf{y}) \parallel q(\mathbf{y})] \quad (\text{A.10i})$$

With these in mind, and assuming the posterior is proper, applying the procedure in Lemma 1 maintains the stationarity of the posterior. Thus, it satisfies the conditions in Lemma 2 and reduces the KL divergence between the new approximation $\tilde{q}(\boldsymbol{\theta})$ and the posterior target. Noting the following identity, Theorem 1 follows as $\ln p(\mathbf{y})$ is unchanged by the transformation as the posterior remains invariant. Note this proof works for other transition kernels that maintain the stationarity of the true posterior.

$$\text{KL}[q(\boldsymbol{\theta}) \parallel p(\boldsymbol{\theta}|\mathbf{y})] = E_{q(\boldsymbol{\theta})}[\ln q(\boldsymbol{\theta})] - E_{q(\boldsymbol{\theta})}[\ln p(\boldsymbol{\theta}|\mathbf{y})] = -\text{ELBO}_{q(\boldsymbol{\theta})} + \ln p(\mathbf{y}) \quad (\text{A.11})$$

B.3 Applying MAVB

To apply the above results to the specific case in the paper, I assume a conditionally conjugate prior on the expansion parameter, i.e. $p_0(\boldsymbol{\mu}_j) \sim N(\mathbf{0}, \tau^2 \mathbf{I}_{d_j})$. Applying Definition 2 of MAVB (i.e. following Theorem 1 in Liu and Wu 1999) gives the following Algorithm 3.

The limiting case can be found as $\tau^2 \rightarrow \infty$ or by applying the logic of Scheme 2.1 in Liu and Wu (1999). A MAVB corresponding to the optimal marginal augmentation scheme and a flat improper prior on $\boldsymbol{\mu}_j$ is shown in the main text's Algorithm 2. Putting a highly diffuse working prior, e.g. $\tau^2 = 10^4$, should lead to nearly identical results. Note that a highly *informative* working prior, e.g. $\tau^2 = 10^{-6}$, results in effectively no change from applying MAVB.

Algorithm 3 MAVB with a Proper Working Prior

Set the Number of Samples Desired: M

Estimate $q(\theta)$ **using** CAVI (Algorithm 1)

For m in $1, \dots, M$

1. Draw $\theta^{(m)} \sim q(\theta)$
2. Draw an expansion parameter from its working prior: $\mu_j^{(m)} \sim p_0(\mu_j)$
3. Create the transformed versions of the parameters as follows

$$\alpha_{j,g}^X = \alpha_{j,g}^{(m)} + \mu_j^{(m)}$$

$$\beta^X = \beta^{(m)} - \sum_{j=1}^J M_j \mu_j^{(m)}$$

4. Sample a second expansion parameter μ_j for each j

$$\tilde{\mu}_j^{(m)} | \alpha_{j,g}^X, \Sigma_j^{(m)} \sim N \left(\left[\left(\Sigma_j^{(m)} \right)^{-1} g_j + 1/\tau^2 \mathbf{I}_{d_j} \right]^{-1} \left[\Sigma_j^{(m)} \right]^{-1} \left(\sum_{g=1}^{g_j} \alpha_{j,g}^X \right), \left[\left(\Sigma_j^{(m)} \right)^{-1} g_j + 1/\tau^2 \mathbf{I}_{d_j} \right]^{-1} \right)$$

5. Adjust the draws to get the improved sample $\tilde{\theta}^{(m)}$

$$\tilde{\alpha}_{j,g}^{(m)} = \alpha_{j,g}^X - \tilde{\mu}_j^{(m)}, \quad \tilde{\beta}^{(m)} = \beta^X + \sum_{j=1}^J M_j \tilde{\mu}_j^{(m)}$$

Appendix C: Accelerating CAVI using PX-VB and Joint Updates

This section notes two simple ways to accelerate convergence of CAVI at limited computational cost. I first derive the methods and show their impact.

C.1 Derivation of Acceleration Techniques

First, I derive a new application of PX-VB (Jaakkola and Qi 2007) to hierarchical models. The procedure is re-formulated below in the notation of this paper:

Lemma 3 (Parameter Expanded Variational Bayes - Jaakkola and Qi 2007). *Given some factorization assumption \mathcal{X} , the following procedure converges no slower than the associated CAVI algorithm and maintains a monotonic improvement of the ELBO.*

1. Perform one step of CAVI (e.g. Algorithm 1, Steps 1-4) giving $q(\boldsymbol{\theta})$ and $\text{ELBO}_{q(\boldsymbol{\theta})}$.
2. Noting $q(\boldsymbol{\theta}) \sim^d q(\boldsymbol{\theta}^X)$ when $\boldsymbol{\xi} = \boldsymbol{\xi}_{\text{Null}}$ and thus $\text{ELBO}_{q(\boldsymbol{\theta})}^{X-\boldsymbol{\xi}_{\text{Null}}} = \text{ELBO}_{q(\boldsymbol{\theta})}$, maximize the $\text{ELBO}_{q(\boldsymbol{\theta})}^{X-\boldsymbol{\xi}}$ over $\boldsymbol{\xi}$.

$$\hat{\boldsymbol{\xi}} = \arg \max_{\boldsymbol{\xi}} \text{ELBO}_{q(\boldsymbol{\theta})}^{X-\boldsymbol{\xi}} = \arg \max_{\boldsymbol{\xi}} E_{q(\boldsymbol{\theta})}[\ln p^X(\mathbf{y}, \boldsymbol{\theta} | \boldsymbol{\xi})] - E_{q(\boldsymbol{\theta})}[\ln q(\boldsymbol{\theta})]$$

Note that $\text{ELBO}_{q(\boldsymbol{\theta})}^{X-\hat{\boldsymbol{\xi}}} \geq \text{ELBO}_{q(\boldsymbol{\theta})}$.

3. Apply the reduction function to recover a distribution on the original, non-expanded space. Equivalently, transform $q(\boldsymbol{\theta})$ by applying a change-of-variables using $t_{\hat{\boldsymbol{\xi}}}(\boldsymbol{\theta})$.

$$q'(\boldsymbol{\theta}) = \int t_{\hat{\boldsymbol{\xi}}}(\boldsymbol{\theta}) q(\boldsymbol{\theta}) d\boldsymbol{\theta}$$

Note that $\text{ELBO}_{q'(\boldsymbol{\theta})} = \text{ELBO}_{q(\boldsymbol{\theta})}^{X-\hat{\boldsymbol{\xi}}}$ and $\text{ELBO}_{q'(\boldsymbol{\theta})} \geq \text{ELBO}_{q(\boldsymbol{\theta})}$.

The intuition is very similar to why parameter-expanded EM (PX-EM; Liu et al. 1998) is guaranteed to produce weakly faster convergence. After performing one step of CAVI (i.e. one set of updates to all approximating distributions in the original model, e.g. Algorithm 1), one notes that the model has been estimated *assuming* the expansion parameter is at its null value ($\boldsymbol{\xi} = \boldsymbol{\xi}_{\text{Null}}$), e.g. zero in the case of a location transformation. Thus, by optimizing over $\boldsymbol{\xi}$, one must weakly improve the objective and by applying the “reduction” function, one returns to a $q'(\boldsymbol{\theta})$ in the original variational family with an expansion parameter that is (implicitly) $\boldsymbol{\xi}_{\text{Null}}$. Being able to “move” in the unidentified space allows for faster convergence. As long as the optimal $\hat{\boldsymbol{\xi}} \neq \boldsymbol{\xi}_{\text{Null}}$ (e.g. not equal $\mathbf{0}$ for the location transformation), then applying PX-VB will increase the objective function and thus ensure faster convergence by decreasing the number of iterations required.

The specific algorithm for Scheme I is shown in Algorithm 4. Note that the only terms that involve μ_j is the prior on $\alpha_{j,g}^X$ as shown by the expanded model (Definition 1). Pleasingly, the closed form solution is quite simple: Center the random effects $\{\alpha_{j,g}\}_{g=1}^{g_j}$ to be mean zero and adjust β correspondingly.

Algorithm 4 Accelerating CAVI for Scheme I (Mean-Expansion)

Initialize as per Algorithm 1

For t in $1, \dots, T$

1. Perform one step of CAVI (i.e. Updates 1-4 from Algorithm 1) to get $q(\theta)$
2. Maximize $\text{ELBO}_{q(\theta)}^{X-\xi}$ over the expansion parameter $\{\mu_j\}_{j=1}^J$.

$$\{\hat{\mu}_j\}_{j=1}^J = \arg \max_{\{\mu_j\}_{j=1}^J} E_{q(\alpha), q(\Sigma_j)} \left[\sum_{j=1}^J \left[\sum_{g=1}^{g_j} -\frac{1}{2} (\alpha_{j,g} - \mu_j)^T \Sigma_j^{-1} (\alpha_{j,g} - \mu_j) \right] \right]$$

$$\hat{\mu}_j = \frac{1}{g_j} \sum_{g=1}^{g_j} E_{q(\alpha_{j,g})} [\alpha_{j,g}]$$

3. “Reduce” the variational parameters; note that this is guaranteed to result in a weakly higher ELBO than $q(\theta)$ from Step 1:

$$[\tilde{\mu}_{\alpha,j}]_g \leftarrow [\tilde{\mu}_{\alpha,j}]_g - \hat{\mu}_j, \quad \tilde{\mu}_\beta \leftarrow \tilde{\mu}_\beta + \sum_{j=1}^J M_j \hat{\mu}_j$$


4. Check for convergence.
-

A second way to accelerate the algorithm is to update certain parameter blocks jointly. Note that Algorithm 1 implies a cyclical rotation through each random effect j to update $\tilde{\mu}_{\alpha,j}$ and $\tilde{\Lambda}_{\alpha,j}$. Initial experiments showed that dramatic gains could be achieved by updating *all* $\tilde{\mu}_{\alpha,j}$ and $\tilde{\mu}_\beta$ jointly. This can be done very quickly using a sparse Cholesky decomposition. Given those mean parameters, the $\tilde{\Lambda}_{\alpha,j}$ can be updated cyclically. For models considered in this paper, this added little computational cost but dramatically improved convergence.

C.2 Impact of Acceleration

Figure A.6 shows the major limitation of the naive CAVI in Algorithm 1. I focus on Models 1, 4, and 9 for clarity. Even given 1,000 iterations, it still has not obtained convergence using the threshold of the ELBO changing by less than 10^{-8} or all parameters changing by less than 10^{-5} . Note that only the lines that are *dashed* reached convergence. It compares the two strategies for accelerating discussed above: (i) PX-VB (“Mean”, vs “None”) and (ii) jointly updating the mean parameters of $q(\beta, \alpha)$ (“Joint”, vs “Naive”).

Figure A.6: ELBO After Acceleration



Note: The horizontal axis (number of iterations) is shown on a logged scale to illustrate the differences more clearly. Dashed lines refer to an algorithm that converged within 1,000 iterations.

The PX-VB algorithm noted above does help performance by a faster increasing ELBO in most models. Jointly updating the mean parameters, however, has a much more considerable impact in all models leading to convergence in sometimes orders of magnitude fewer iterations.

Table A.4 shows the change in the ELBO at the final iteration, i.e. the 1000th iteration or convergence. We see that using PX-VB as defined above after jointly updating the parameters does little, while there is a clear improvement from using PX-VB when applying the naive cyclical updates in Algorithm 1. Even in the cases where PX-VB is insufficient to allow convergence, the final change in the ELBO is around $\times 10$ smaller (e.g. 10^{-5} vs 10^{-4}) after 1000 iterations.


Figure A.7 explains why this is the case; it plots the trajectory of the mean of the random effects at each iteration, i.e. $1/g_j \sum_{g=1}^{g_j} [\tilde{\mu}_{\alpha,j}]_g$. It shows starkly that there is a very slow decline towards zero for the naive cyclical algorithm without parameter expansion. By contrast, using PX-VB (Algorithm 4) or the joint updates resolves this

Table A.4: Final Change in ELBO by Estimation Method

Estimation Method	Year	Model 1	Model 4	Model 9
None - Joint	2004	-7.65	-8.15	-8.07
None - Joint	2008	-7.61	-8.05	-8.09
Mean - Joint	2004	-7.65	-8.15	-8.07
Mean - Joint	2008	-7.61	-8.05	-8.09
None - Naive	2004	-4.08	-4.19	-3.74
None - Naive	2008	-3.99	-4.02	-3.72
Mean - Naive	2004	-4.61	-5.32	-5.35
Mean - Naive	2008	-4.67	-5.34	-5.25

Note: The final change in the ELBO (i.e. at convergence or the 1000th iteration) is shown. All numbers can be interpreted as the final change is 10^x , i.e. \log_{10} of the change in the ELBO. Estimation Method is described in the main text and is the type of parameter expansion - the update method for the mean parameters of $q(\beta, \alpha)$.

problem as they are centered at every iteration.

Figure A.7: Drift in $\bar{\alpha}_j$ 

Appendix D: Additional Simulations

D.1 Varying Magnitude of Parameters

I show the model in the main text where the standard deviation of the fixed effects (i.e. $\beta \sim N(0, \sigma_\beta^2 \mathbf{I})$) and the random effects (i.e. $\alpha_{1,g[i]} \sim N(0, \sigma_\alpha^2)$) varies. In the main text, I only report results where $\sigma_\beta = 0.2$ and $\sigma_\alpha = 1$. The accuracy and coverage are shown below as they are more discriminating across methods.

The story is broadly similar to the main text; in cases where σ_α is large, Schemes I and II undercover. Applying MAVB improves this and puts coverage to near nominal levels. It similarly has large benefits to accuracy. The one exception is when the distribution of the true fixed effects is wide (i.e. $\sigma_\beta = 1$). In this case, even after applying MAVB, the coverage of the fixed effects is somewhat poor (0.90) although above ADVI. Accuracy also remains noticeably below the Laplace approximation. In those cases (i.e. $\sigma_\beta = 1$), it is worth noting that the linear predictor has a very wide distribution (ψ_i has a 5-95% interval of around -5.9 to 5.5) and the probability distribution of p_i is highly bimodal around 0 and 1. When $\sigma_\beta = 0.2$, as in the main text, the scale is more plausible for a logistic model with the 5-95% interval of the linear predictor being roughly between -2.7 and 2.3—with a much more uniform distribution of probabilities p_i . This is roughly the scale of the linear predictors observed in bootstrapped data.

D.2 Bootstrapping Ghitza and Gelman

To address this and see how the variational schemes perform on real data, I conducted a simulation based on [Ghitza and Gelman \(2013\)](#). I fit Model 1— $J = 4$ with random effects for age, income, ethnicity and state—and take the parameter estimates from the Laplace approximation as the “ground truth”. Given those estimates, I generate 100 datasets with the observed covariates from 2004 and simulate a binary outcome for each of the 4,080 observations. Table [A.5](#) reports three measures of the performance; and reports the same figures as in Table [1](#).

The results are similar to those in the main text. Scheme I without MAVB outperforms ADVI (Mean Field) in terms of Bias, RMSE, and coverage while having approximately the same accuracy. The comparison against ADVI (Full Rank) is more complex; it has a noticeably higher RMSE and bias, although its accuracy is noticeably better than Scheme I. Applying MAVB to Scheme I results in a large improvement in accuracy (around 10%) and makes it comparable to ADVI (Full Rank). Scheme III continues to perform well having a high accuracy (94-95%) and outperforms all other approximate methods including both versions of ADVI. Its performance is slightly improved (1-2%) by MAVB.

D.3 Checking for Non-Convexities

Models estimated by variational inference have the possibility of getting stuck in local optima. The default settings in the variational algorithms use a deterministic initialization based on an EM algorithm where the random effects are replaced by a fixed ridge

Table A.5: Results from Bootstrapped Simulations

		Bias		RMSE		Accuracy		Coverage	
		FE	RE	FE	RE	FE	RE		
	Laplace	0.004	-0.003	0.017	0.015	0.929	0.916	0.977	0.992
	HMC							0.977	0.994
	ADVI (MF)	-0.005	0.002	0.083	0.048	0.631	0.829	0.802	0.974
	ADVI (FR)	-0.074	0.027	0.295	0.159	0.814	0.861	0.940	0.975
	Scheme I	0.003	-0.003	0.016	0.014	0.659	0.836	0.913	0.985
	Scheme II	0.003	-0.003	0.016	0.013	0.659	0.918	0.912	0.991
	Scheme III	0.003	-0.003	0.016	0.011	0.943	0.959	0.973	0.992
MAVB +	Scheme I	0.003	-0.003	0.016	0.015	0.791	0.936	0.932	0.991
	Scheme II	0.002	-0.003	0.016	0.013	0.793	0.943	0.930	0.991
	Scheme III	0.003	-0.003	0.016	0.011	0.951	0.965	0.970	0.992

Note: This reports the bias (Bias), root mean squared error (RMSE) of the estimated posterior means against those estimated from HMC. The distance between the distributions (Accuracy) and frequentist coverage (Coverage) are reported; see the main text for an explanation of these measures. The statistics are disaggregated by fixed (FE) and random effects (RE). All results are created using all relevant parameters in each simulation and then averaged across one hundred simulations. ADVI (MF) uses the mean-field approximation; ADVI (FR) uses the full rank approximation in [Kucukelbir et al. \(2017\)](#).

prior. I examine a (reasonable) random initialization where (i) the variance parameters are set to zero (i.e. $\mathbf{\Lambda}_{\alpha-\beta} = \mathbf{0}$) and (ii) the mean parameters are randomly sampled from standard independent Gaussians and $\tilde{\Phi}_j$ is drawn from a standard Inverse Wishart.

Figure A.9 shows the results of the EM starting values (in red) versus 100 random initializations using Scheme I on Models 1, 4, and 9. It is clear that all converge to the same value of the ELBO. The maximum difference between any of the random initializations and the EM initialized model is less than 10^{-7} .

In my other experiments, I did not observe issues with non-convexity although it might occur in other settings or with poor starting values. Thus, it is reasonable for the researcher to check for this in their domain-specific setting.

Appendix E: Additional Analyses for Ghitza & Gelman

Figure A.10 shows disaggregated results on the run time for the models shown in Figure 1 by breaking the time into that used for estimating the parameters via CAVI and drawing samples and transforming them using MAVB. As noted in the main text, the time for MAVB is the time to draw 4,000 samples to make it comparable to the output of HMC. It shows that MAVB never consumes more than a minute or two of time—although its cost does grow in Schemes II and III when the number of parameters is huge.

Figure A.11 shows the mean time of each stage, averaged across years, for the nine models in the main paper. It shows that the main increase in run-time for Schemes I, II, and III comes from updating $q(\beta, \alpha)$ —although there is a slight increase in the

Polya-Gamma update time. Further, it shows that the increase in the size of α imposes noticeably more significant costs on Schemes II and III.

Figure A.12 shows the results of applying MAVB to the percentage gap on the absolute value of the means between the approximate methods and Hamiltonian Monte Carlo. As before, the distribution of values is over the statistic aggregated across groups. Note that the scale is much smaller than the corresponding figure on standard deviations in the main text (Figure 3) especially for the fixed effects.

Figure A.13 shows the average (percentage) discrepancy for each type of random effect. As noted in the main text, it shows that there is severe underestimation of variability for random effects with small numbers of groups—especially age, income and ethnicity. Fortunately, applying MAVB to Scheme I corrects some of this discrepancy giving it variability that is closer to the baseline of Hamiltonian Monte Carlo and beating the Laplace approximation on certain blocks. Scheme III remains systematically close to HMC and clearly out-performs the Laplace approximation on this measure.


Figure A.14 shows the correlation for the linear predictor (i.e. $\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{z}_i^T \boldsymbol{\alpha}$) against HMC—showing both mean and standard deviation. I focus only on the observations where $n_i > 0$ as `brms` does not use observations where $n_i = 0$ in estimating the model. For the mean, there is very close correspondence as should be expected from the tight correspondence of the means shown elsewhere. For the standard deviation, there is more variable performance. Scheme I has a small bias -0.013 (or 2%), a high correlation (0.96) but clearly quite wide variability. Schemes II and III are much closer (0.99 correlation; visually quite tight) if slightly systematically under-estimating variability. MAVB adds little benefit here.

Figure A.15 shows the result of a second cross-validation where all models (HMC, Laplace, VI [Scheme I]) are fit on ten folds and out-of-sample predictions are generated. It reports the deviance, as defined in the main text. All methods select Model 4 over Model 9 although less clearly in the case of Scheme I in 2008. Note that the folds here are separate from those in the main text; thus, it is reassuring that Model 4 is again selected by cross-validation. “Stan (LinPred)” refers draws of the expected outcome (i.e. the linear predictor draws pushed through the logistic link and divided by n_i to get a probability; `posterior_linpred(..., transform = TRUE)`). “Stan(PostPred)” refers to averaged draws of the posterior predictive distribution (`posterior_predict(...)`). The two measures are nearly identical with a mean absolute error of 0.003 on the probability scale.

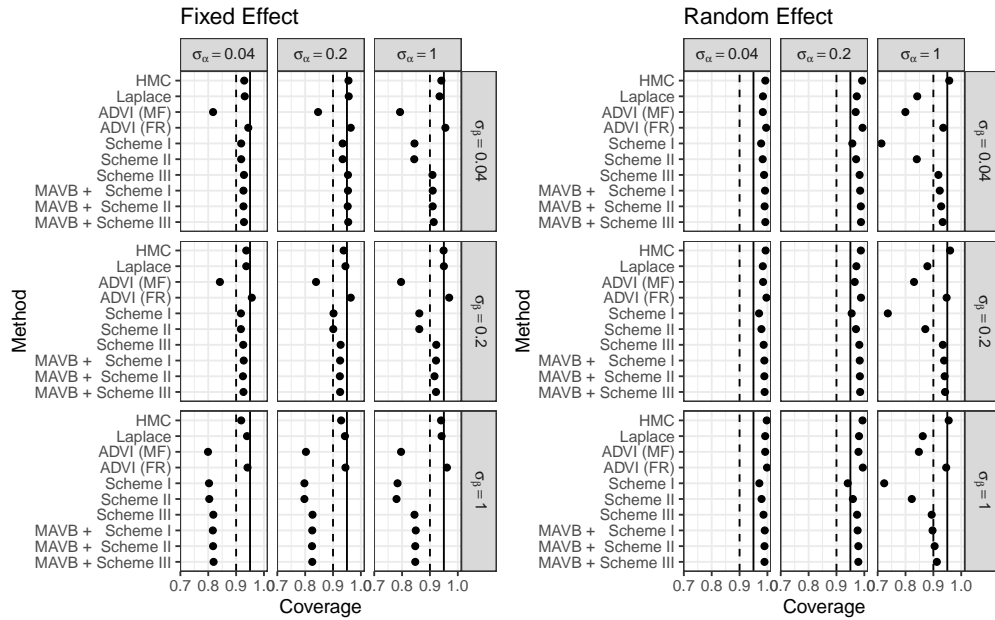
Finally, Figure A.16 shows the correlation between the out-of-sample predictions between HMC and the Laplace and Scheme I methods. As expected from the main results, the Laplace predictions are nearly identical to the HMC predictions while the Scheme I are highly correlated ($\rho = 0.998$) but clearly somewhat more noisy.

Figure A.8: Varying Size of Parameters

(a) Accuracy




(b) Coverage




Note: The solid line indicates 0.95 and the dashed line indicates 0.90.

Figure A.9: Examining Different Initializations




Note: The red line indicates a model initialized via EM; the black lines indicate random initializations.

Figure A.10: Disaggregating Run Time between CAVI and MAVB




Note: Each figure plots the run-time of Schemes I-III. The reported times are averaged across the 2004 and 2008 elections. The time is broken up into Total, CAVI (parameter estimation), and MAVB (4,000 samples plus transformation). Model 1-9 are described in Table 2. All models are fit on a computer with 16 GB of RAM.

Figure A.11: Time of CAVI by Scheme and Stage



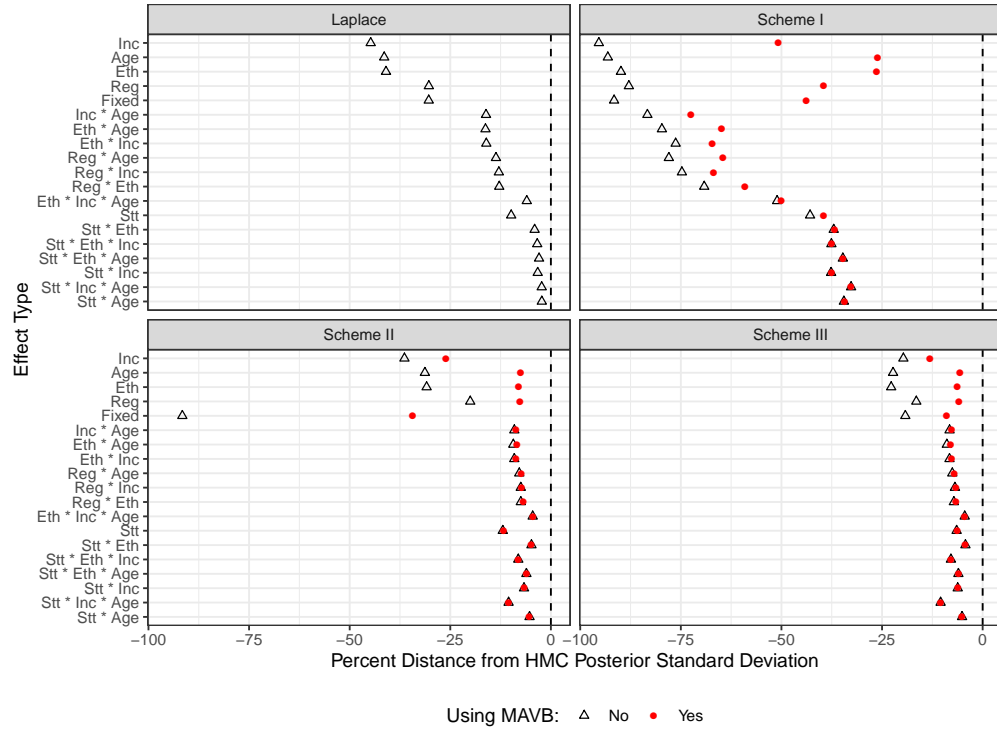
Note: The average time in seconds per each stage of CAVI is noted. “Beta/Alpha” refers to updating $q(\beta, \alpha)$, “Poly-Gamma” refers to $q(\{\omega_i\}_{i=1}^N)$, and “Sigma” refers to $q(\{\Sigma_j\}_{j=1}^J)$.

Figure A.12: Improvements from MAVB on Mean



Note: This figure plots the percentile of the percentage gap between the absolute value of the means estimated via Hamiltonian Monte Carlo [HMC] and the approximate methods. A negative value on the vertical axis indicates that the corresponding percentile has a smaller variance than HMC. A vertical shift upward of the line indicates the variance of the parameters has increased. The solid markers indicate the deciles and extremes of the distribution. The dashed line with hollow triangles represents the estimates without using MAVB. The red line with solid circles represents the results after using MAVB.

Figure A.13: Disaggregated Improvements from MAVB



Note: This figure plots the average discrepancy for the standard deviation of each random effect estimated by an approximation method from that coming from Hamiltonian Monte Carlo [HMC]. Negative values indicate that HMC has a larger standard deviation. Hollow triangles represent the estimates without MAVB. Filled circles represents the estimates using MAVB.

Figure A.14: Analysis of Linear Predictor

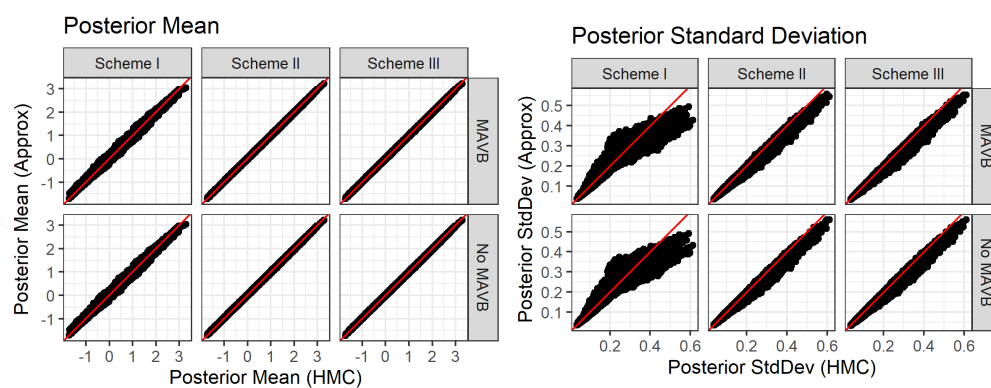
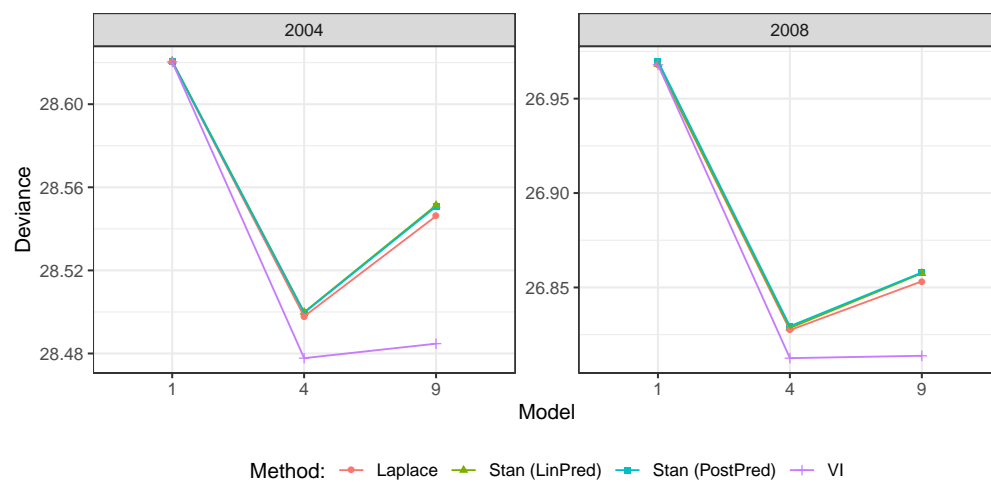



Figure A.15: “Gold Standard” CV



Note: This plots the mean out-of-sample deviance.

Figure A.16: Correlation of Out-of-Sample Predictions



Note: The out-of-sample predicted probabilities p_i are shown. A red line demarks the 45-degree line.