

Optimal Gaussian Approximations to the Posterior for Log-Linear Models with Diaconis–Ylvisaker Priors

James Johndrow^{*†} and Anirban Bhattacharya^{†§}

Abstract. In contingency table analysis, sparse data is frequently encountered for even modest numbers of variables, resulting in non-existence of maximum likelihood estimates. A common solution is to obtain regularized estimates of the parameters of a log-linear model. Bayesian methods provide a coherent approach to regularization, but are often computationally intensive. Conjugate priors ease computational demands, but the conjugate Diaconis–Ylvisaker priors for the parameters of log-linear models do not give rise to closed form credible regions, complicating posterior inference. Here we derive the optimal Gaussian approximation to the posterior for log-linear models with Diaconis–Ylvisaker priors, and provide convergence rate and finite-sample bounds for the Kullback–Leibler divergence between the exact posterior and the optimal Gaussian approximation. We demonstrate empirically in simulations and a real data application that the approximation is highly accurate, even for modest sample sizes. We also propose a method for model selection using the approximation. The proposed approximation provides a computationally scalable approach to regularized estimation and approximate Bayesian inference for log-linear models.

Keywords: credible region, conjugate prior, contingency table, Dirichet–Multinomial, Kullback–Leibler divergence, Laplace approximation.

1 Introduction

Contingency table analysis routinely relies on log-linear models, which represent the logarithm of cell probabilities as an additive model (Agresti, 2002). With the standard choice of Multinomial or Poisson likelihood, these are exponential family models, and are routinely fit through maximum likelihood estimation (Fienberg and Rinaldo, 2007). However, sparsity in the observed cell counts often makes maximum likelihood estimation infeasible (see Haberman (1974) and Bishop et al. (2007)) in practical applications. In such cases, regularization is often used to obtain unique parameter estimates (Park and Hastie, 2007; Zou and Hastie, 2005).

^{*}Department of Statistical Science, Duke University, Durham, North Carolina, USA, jj@stat.duke.edu

[†]Department of Statistics, Texas A&M University, College Station, Texas, USA, anirbanb@stat.tamu.edu

[‡]Dr. Johndrow gratefully acknowledges support from grant ES020619 from the National Institute of Environmental Health Sciences (NIEHS) of the U.S. National Institutes of Health.

[§]Dr. Bhattacharya acknowledges support for this project from the Office of Naval Research and the National Science Foundation (NSF DMS award # 1613193).

A common Bayesian approach to inference in high-dimensional contingency tables is to place a conjugate prior on the parameters of a graphical or hierarchical log-linear model, and an independent prior over the space of all such models (see e.g. Massam et al. (2009)). This leads to a standard model-averaged posterior (Hoeting et al., 1998), where all possible sparse log-linear models in the chosen class are weighted by their posterior evidence. Use of non-conjugate (e.g. Gaussian) priors with computation by Markov chain Monte Carlo (Gelfand and Smith, 1990) has also been proposed (Dellaportas and Forster, 1999). Although model averaging is generally considered ideal in high dimensional settings, computational algorithms for posterior inference scale exceedingly poorly in p . Since the smallest contingency table corresponding to cross-classification of p categorical variables has 2^p cells, the corresponding log-linear model has $2^p - 1$ free parameters, so the model space grows super-exponentially in p . Accordingly, posterior computation is essentially infeasible for $p > 15$, the largest case demonstrated to date in the literature (Dobra and Massam, 2010) to the best of our knowledge.

Alternatively, one can place a Gaussian prior on the parameters of a saturated log-linear model to induce Tikhonov type regularization, and then perform computation by Markov chain Monte Carlo. This approach is well-suited to situations in which the sample size is not tiny relative to the table dimension, but where zero counts nonetheless exist in some cells. In this case, data augmentation Gibbs samplers such as that proposed by Polson et al. (2013) provide for conditionally conjugate updates. However, this by itself is computationally intensive relative to alternatives such as elastic net (Zou and Hastie, 2005), and can suffer from poor mixing. In principle, a more scalable Bayesian approach for producing Tikhonov regularized point estimates would be to utilize the Diaconis–Ylvisaker conjugate prior (Diaconis and Ylvisaker, 1979) on the parameters of the log-linear model, which is essentially computation free. The main drawback is that the resulting posterior distribution is difficult to work with, lacking closed form expressions for even marginal credible intervals or fast algorithms for sampling from the posterior. An accurate and more tractable approximation to this posterior is therefore of practical interest.

Approximations to the posterior distribution have a long history in Bayesian statistics, with the Laplace approximation perhaps the most common and simple alternative (Tierney and Kadane, 1986; Shun and McCullagh, 1995). More sophisticated approximations, such as those obtained using variational methods (Attias, 1999) may in some cases be more accurate but require computation similar to that for generic EM algorithms. Moreover, there exist no theoretical guarantees of the approximation error in finite samples, and these approximations are known to be inadequate in relatively simple models (Wang and Titterton, 2004, 2005).

In this article, we propose a Gaussian approximation to the posterior for log-linear models with Diaconis–Ylvisaker priors. The approximation is shown to be the optimal Gaussian approximation to the posterior in the Kullback–Leibler divergence, and convergence rates to the exact posterior and a finite-sample Kullback–Leibler error bound are provided. The approximation is shown empirically to be accurate even for modest sample sizes; effectively, the empirical results suggest that the approximation is accurate enough to be used in place of the exact posterior within the range of sample sizes for which the posterior is sufficiently concentrated to be statistically useful. We

also propose a procedure to perform model selection using the posterior approximation computed on marginal tables of increasing order. The procedure performs comparably in model selection for graphical log-linear models compared to methods requiring vastly greater computational resources. A `Matlab` implementation of our procedure is available at <https://github.com/jamesjohndrow/dynormal-approx>.

2 Background

We first provide a brief review of exponential families. We then describe the family of conjugate priors for the natural parameter of an exponential family, referred to as Diaconis–Ylvisaker priors. We then provide more detailed background on log-linear models for Multinomial likelihoods and the associated Diaconis–Ylvisaker prior.

2.1 Exponential families

Following Diaconis and Ylvisaker (1979), let μ be a σ -finite measure defined on $(\mathbb{R}^p, \mathcal{B})$, where \mathcal{B} denotes all Borel sets on \mathbb{R}^p . Let $\text{supp}(\mu) = \{y \in \mathbb{R}^p : d\mu(y) > 0\}$ be the support of μ , and define \mathcal{Y} as the interior of the convex hull of $\text{supp}(\mu)$. For $\theta \in \mathbb{R}^p$, define $M(\theta) = \log \int_{\mathcal{Y}} e^{\theta^T y} d\mu(y)$, and let $\Theta = \{\theta \in \mathbb{R}^p : M(\theta) < \infty\}$, which we assume is an open set. We refer to Θ as the natural parameter space. The exponential family of probability measures $\{P(\cdot; \theta)\}$ indexed by a parameter $\theta \in \Theta$ is defined by

$$dP(y; \theta) = e^{\theta^T y - M(\theta)} d\mu(y), \quad \theta \in \Theta. \quad (1)$$

This family includes many of the probability distributions commonly used as sampling models in likelihood-based statistics. Diaconis and Ylvisaker (1979) develop the family of conjugate priors for the parameter θ of regular exponential family likelihoods. These Diaconis–Ylvisaker priors are given by

$$d\pi(\theta; n_0, y_0) = e^{n_0 y_0^T \theta - n_0 M(\theta)}, \quad n_0 \in \mathbb{R}, y_0 \in \mathbb{R}^d. \quad (2)$$

On observing data y consisting of n observations with sufficient statistics \bar{y} , the posterior is then also Diaconis–Ylvisaker, with parameters $n_0 + n, y_0 + \bar{y}$, i.e. $d\pi(\theta | y) = d\pi(\theta; n_0 + n, y_0 + \bar{y})$. In the sequel we focus on one member of the exponential family, the multinomial. In the natural parametrization, the multinomial likelihood gives rise to the log-linear model and the closely related multinomial logit model, which we now describe.

2.2 Log-linear models

Let $\mathcal{S}^d = \{(x_1, \dots, x_d) \in [0, 1]^d : \sum_{j=1}^d x_j \leq 1\}$ denote the d -dimensional unit simplex. Consider N independent samples from a categorical variable with $(d + 1)$ levels. We denote the levels of the variable by $0, 1, \dots, d$, without loss of generality. Let y_j denote the number of times the j th level is observed in the N samples and set $y = (y_0, y_1, \dots, y_d)^T$; clearly $\sum_{j=0}^d y_j = N$. The joint distribution of y is given by a multinomial distribution,

denoted $y \sim \text{Multinomial}(N, \pi)$, which is parametrized by $\pi = (\pi_1, \dots, \pi_d)^\top \in \mathcal{S}^d$, where π_j is the probability of observing the j th level for $j = 1, \dots, d$.

The log-linear model is a generalized linear model for multinomial likelihoods obtained by choosing the logistic link function, which also results in the natural exponential family parametrization. Define the logistic transformation $\ell : \mathbb{R}^d \rightarrow \mathcal{S}^d$ and its inverse log ratio transformation $\ell^{-1} : \mathcal{S}^d \rightarrow \mathbb{R}^d$ as

$$\pi_j = \frac{e^{\theta_j}}{1 + \sum_{l=1}^d e^{\theta_l}}, \quad \theta_j = \log(\pi_j/\pi_0), \quad (j = 1, \dots, d), \quad (3)$$

where $\pi_0 = 1 - \sum_{j=1}^d \pi_j$, and $\theta_0 = 0$. We shall write $\pi = \ell(\theta)$ and $\theta = \ell^{-1}(\pi) = \log(\pi/\pi_0)$, respectively, to denote the transformations in (3). Using (3), the multinomial likelihood in the log-linear parameterization can be expressed as

$$f(y | \theta) \propto \frac{\exp\left(\sum_{j=1}^d y_j \theta_j\right)}{\left(1 + \sum_{l=1}^d e^{\theta_l}\right)^N}. \quad (4)$$

An important motivating case is when $y = \text{vec}(\mathbf{n})$, with \mathbf{n} a contingency table arising from cross-classification of N independent observations on p categorical variables w_1, \dots, w_p . Suppose that the v th variable w_v has d_v many levels, so that the contingency table has $\prod_{v=1}^p d_v$ many cells, and y is a $(d+1)$ -dimensional vector of counts with $d = \prod_{v=1}^p d_v - 1$. We refer to the parametrization $\theta = \log(\pi/\pi_0)$ in the contingency table setting as the *identity* parametrization. Also of particular interest in this setting are reparametrizations of (3) that represent $\log \pi/\pi_0$ as an additive model involving parameters that correspond to interactions among w_1, \dots, w_p . Every identified parametrization of the log-linear model for the multinomial likelihood can be represented by

$$\log(\pi/\pi_0) = X\tilde{\theta}, \quad (5)$$

where X is a d by d non-singular binary matrix and $\tilde{\theta} \in \mathbb{R}^d$. In the simulations and application, we make a specific choice for X that corresponds to the *corner parametrization* of the log-linear model (Massam et al., 2009). We illustrate the identity and corner parameterizations through a 2^3 contingency table in Example 1 below. Details for the general case can be found in Section 6.1.

Example 1. Consider three binary variables w_1, w_2, w_3 , with $w_v \in \{0, 1\}$ for $v = 1, 2, 3$, and let

$$\psi_{i_1 i_2 i_3} = \text{pr}(w_1 = i_1, w_2 = i_2, w_3 = i_3), \quad (i_1, i_2, i_3) \in \{0, 1\}^3.$$

A 2^3 contingency table $\mathbf{n} = (n_{i_1 i_2 i_3})$ is obtained from the cross-classification of N independent observations on w_1, w_2, w_3 , with $n_{i_1 i_2 i_3}$ denoting the cell count for the cell (i_1, i_2, i_3) . Let $y = \text{vec}(\mathbf{n}) = (n_{000}, \dots, n_{111})^\top$ be the vectorized cell counts with $d = 7$. In the *identity* parametrization, the vector of log-linear parameters $\theta \in \mathbb{R}^7$ is given by

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \theta_4 \\ \theta_5 \\ \theta_6 \\ \theta_7 \end{pmatrix} = \log \begin{pmatrix} \pi_1/\pi_0 \\ \pi_2/\pi_0 \\ \pi_3/\pi_0 \\ \pi_4/\pi_0 \\ \pi_5/\pi_0 \\ \pi_6/\pi_0 \\ \pi_7/\pi_0 \end{pmatrix} = \log \begin{pmatrix} \psi_{001}/\psi_{000} \\ \psi_{010}/\psi_{000} \\ \psi_{011}/\psi_{000} \\ \psi_{100}/\psi_{000} \\ \psi_{101}/\psi_{000} \\ \psi_{110}/\psi_{000} \\ \psi_{111}/\psi_{000} \end{pmatrix}.$$

On the other hand, in the *corner* parametrization, we express

$$\begin{aligned} \theta = \log \begin{pmatrix} \psi_{001}/\psi_{000} \\ \psi_{010}/\psi_{000} \\ \psi_{011}/\psi_{000} \\ \psi_{100}/\psi_{000} \\ \psi_{101}/\psi_{000} \\ \psi_{110}/\psi_{000} \\ \psi_{111}/\psi_{000} \end{pmatrix} &= \begin{pmatrix} \tilde{\theta}_{001} \\ \tilde{\theta}_{010} \\ \tilde{\theta}_{001} + \tilde{\theta}_{010} + \tilde{\theta}_{011} \\ \tilde{\theta}_{100} \\ \tilde{\theta}_{001} + \tilde{\theta}_{100} + \tilde{\theta}_{101} \\ \tilde{\theta}_{010} + \tilde{\theta}_{100} + \tilde{\theta}_{110} \\ \tilde{\theta}_{001} + \tilde{\theta}_{010} + \tilde{\theta}_{100} + \tilde{\theta}_{011} + \tilde{\theta}_{101} + \tilde{\theta}_{110} + \tilde{\theta}_{111} \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \times \begin{pmatrix} \tilde{\theta}_{001} \\ \tilde{\theta}_{010} \\ \tilde{\theta}_{011} \\ \tilde{\theta}_{100} \\ \tilde{\theta}_{101} \\ \tilde{\theta}_{110} \\ \tilde{\theta}_{111} \end{pmatrix} \\ &= X\tilde{\theta}. \end{aligned}$$

The indexing of the elements of $\tilde{\theta}$ by binary indices is for ease of interpretation. Indeed, entries of $\tilde{\theta}$ with a single 1 in the binary index are main effects, those with two 1's are two-way interactions and $\tilde{\theta}_{111}$ is a three-way interaction term. The matrix X can be easily verified to be non-singular, so that the θ and $\tilde{\theta}$ parametrizations are equivalent, with $d = 7$ free parameters in either case.

2.3 Conjugate priors for log-linear models

We now present the Diaconis–Ylvisaker prior for the multinomial likelihood (4) and derive an optimal Gaussian approximation to the corresponding posterior in Kullback–Leibler divergence. Extensions to log-linear models with a non-identity parametrization (i.e., $X \neq I_d$ in (5)) is straightforward by invariance properties of the Kullback–Leibler divergence and are discussed subsequently. All proofs are deferred to Section 6.

For the multinomial likelihood (4), the Diaconis–Ylvisaker prior is obtained by applying the inverse logistic transformation ℓ^{-1} to a Dirichlet distribution, which is the conjugate prior for π (Gutiérrez-Pena and Smith, 1995; Consonni et al., 2004). Recall that $\pi_0 = 1 - \sum_{j=1}^d \pi_j$. The Dirichlet distribution $\mathcal{D}(\alpha)$ on \mathcal{S}^d with parameter vector $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_d)^\top$ has density

$$q(\pi; \alpha) = \frac{\Gamma(\sum_{j=0}^d \alpha_j)}{\prod_{j=0}^d \Gamma(\alpha_j)} \prod_{j=0}^d \pi_j^{\alpha_j - 1}, \quad \pi \in \mathcal{S}^d, \quad (6)$$

and corresponding probability measure $\mathcal{Q}(\cdot, \alpha)$ with $\mathcal{Q}(A, \alpha) = \int_A q(\pi; \alpha) d\pi$ for Borel subsets A of \mathcal{S}^d .

Proposition 1. *Suppose $\pi \sim \mathcal{D}(\alpha)$ and let $\theta = \log(\pi/\pi_0) \in \mathbb{R}^d$. Define $A = \sum_{j=0}^d \alpha_j$. Then θ has a density on \mathbb{R}^d given by*

$$p(\theta; \alpha) = \frac{\Gamma(\sum_{j=0}^d \alpha_j)}{\prod_{j=0}^d \Gamma(\alpha_j)} \frac{\exp(\sum_{j=1}^d \alpha_j \theta_j)}{(1 + \sum_{l=1}^d e^{\theta_l})^A}. \quad (7)$$

We write $\theta \sim \mathcal{LD}(\alpha)$ and use $\mathcal{P}(\cdot; \alpha)$ to denote the probability measure associated with the density (7), with $\mathcal{P}(B; \alpha) = \int_B p(\theta; \alpha) d\theta$ for Borel subsets B of \mathbb{R}^d . If a non-identity parametrization $\theta = X\tilde{\theta}$ as in (5) is employed, then we denote the induced distribution on $\tilde{\theta} = X^{-1}\theta$ by $\mathcal{P}_X(\cdot; \alpha)$ and the density by $p_X(\theta; \alpha)$.

It is immediate that $\mathcal{LD}(\alpha)$ is a conjugate family of prior distributions for the likelihood (4), with the posterior $\theta | y \sim \mathcal{LD}(\alpha + y)$. To obtain some preliminary insight into the distribution family $\mathcal{LD}(\alpha)$, we derive the mean and covariance in Proposition 2 below.

Proposition 2. *Let $\theta \sim \mathcal{LD}(\beta)$, with $\beta = (\beta_0, \beta_1, \dots, \beta_d)^\top$ and $\beta_j > 0$ for all j . Then,*

$$\begin{aligned} E(\theta_j) &= \psi(\beta_j) - \psi(\beta_0), \quad (j = 1, \dots, d), \\ \text{cov}(\theta_j, \theta_{j'}) &= \psi'(\beta_j) \delta_{jj'} + \psi'(\beta_0), \quad (j, j' = 1, \dots, d), \end{aligned}$$

where ψ and ψ' are the digamma and trigamma functions, respectively, and $\delta_{jj'} = 0$ if $j \neq j'$ and $\delta_{jj'} = 1$ otherwise.

The proof of Proposition 2 is established within the proof of Theorem 1 in Section 6. Assume the data y is generated from a Multinomial (N, π^0) distribution and let $\theta^0 = \log(\pi^0/\pi_0^0)$ be the true log-linear parameter, where $\pi_0^0 = 1 - \sum_{j=1}^d \pi_j^0$. If an $\mathcal{LD}(\alpha)$ prior is placed on θ , one can use Proposition 2 to show that the posterior mean $E(\theta | y)$ converges almost surely to θ^0 with increasing sample size, and the posterior covariance $\text{cov}(\theta | y)$ converges to the inverse Fisher information matrix as long as the entries of the prior hyperparameter α are suitably bounded. In fact, a Bernstein–von Mises type result can be established, showing that the posterior distribution approaches a Gaussian distribution, centered at the true parameter value and having covariance the inverse Fisher information matrix, in the total variation metric. We do not pursue such frequentist asymptotic validations further in this paper. Our goal rather is to provide a Gaussian approximation to the posterior distribution that can be used in practice, and provide finite sample bounds to the approximation error.

3 Main results

In this section, we provide an optimal Gaussian approximation to an $\mathcal{LD}(\beta)$ distribution (7) in the Kullback–Leibler divergence, i.e., we exhibit a vector $\mu^* \in \mathbb{R}^d$ and a

positive definite matrix Σ^* such that the Kullback–Leibler divergence between $\mathcal{LD}(\beta)$ and $\mathcal{N}(\mu^*, \Sigma^*)$ is the minimum among all Gaussian distributions. This result provides a readily available Gaussian approximation to the posterior distribution $\mathcal{LD}(\beta = \alpha + y)$ of the log-linear parameter θ in (4) with a Diaconis–Ylvisaker prior $\mathcal{LD}(\alpha)$. We also provide a non-asymptotic error bound for the Kullback–Leibler approximation. Using Pinsker’s inequality, the approximation error in the total variation distance can be bounded in finite samples.

For two probability measures $\nu \ll \nu^*$, we write

$$D(\nu \parallel \nu^*) = E_{\nu^*} \log d\nu/d\nu^*$$

to denote the Kullback–Leibler divergence between ν and ν^* .

Theorem 1. *Given $\beta_j > 0, j = 0, 1, \dots, d$, let $\beta = (\beta_0, \dots, \beta_d)^\top$, and define*

$$\mu_j^* = \psi(\beta_j) - \psi(\beta_0), \quad \sigma_{jj'}^* = \psi'(\beta_j)\delta_{jj'} + \psi'(\beta_0), \tag{8}$$

where ψ and ψ' denote the digamma and trigamma functions respectively. Define $\mu^* = (\mu_j^*) \in \mathbb{R}^d$ and $\Sigma^* = (\sigma_{jj'}^*) \in \mathbb{R}^{d \times d}$. Then,

$$D\left\{\mathcal{LD}(\beta) \parallel \mathcal{N}(\mu^*, \Sigma^*)\right\} = \inf_{\mu, \Sigma} D\left\{\mathcal{LD}(\beta) \parallel \mathcal{N}(\mu, \Sigma)\right\}, \tag{9}$$

where the infimum is over all $\mu \in \mathbb{R}^d$ and all $\Sigma > 0 \in \mathbb{R}^{d \times d}$. Further, if $\beta_j > 1/2$ for all $j = 0, 1, \dots, d$, then

$$D\left\{\mathcal{LD}(\beta) \parallel \mathcal{N}(\mu^*, \Sigma^*)\right\} < \frac{1}{2} \sum_{j=0}^d \frac{1}{\beta_j} + \frac{1}{6B}, \tag{10}$$

where $B = \sum_{j=0}^d \beta_j$.

The matrix Σ^* has a compound-symmetry structure and is therefore positive-definite. From Proposition 2, the parameters of the optimal Gaussian approximation μ^* and Σ^* are indeed the mean and covariance matrix of the $\mathcal{LD}(\beta)$ distribution. Equation (10) provides an upper-bound to the approximation error. In the posterior, $\beta_j = \alpha_j + y_j$ and $B = \sum_{j=0}^d \alpha_j + N$. The condition $\beta_j \geq 1/2$ is therefore satisfied whenever every category has at least one observation. Since

$$\mathbb{E}_y[\alpha_j + y_j] = \alpha_j + N\pi_j^0,$$

the approximation error is approximately in the order of $\sum_{j=0}^d (\pi_j^0 N)^{-1}$, where as before π_j^0 denotes the true probability of category j . In the best case where all the categories receive approximately equal probability, i.e., $\pi_j^0 \asymp (d+1)^{-1}$, the approximation error is $\mathcal{O}(d^2/N)$. However, the convergence rate in N can be slower if some of the π_j^0 s are very small. In other words, the higher the entropy of the data generating distribution, the worse the approximation is, although our simulations suggest that the approximation is

practicable even for moderate sample sizes and unbalanced category probabilities. When one considers that the eigenvalues of the covariance matrix enter into the constant in Berry–Esséen convergence rates, and that here the covariance of the data is given by $\text{diag}(\pi^0) - \pi^0(\pi^0)^\top$, it appears that a similar phenomenon is at work here.

The main idea behind our proof is to exploit the invariance of the Kullback–Leibler divergence under bijective transformations and transfer the domain of the problem from \mathbb{R}^d to \mathcal{S}^d . Since an $\mathcal{LD}(\beta)$ distribution is obtained from a Dirichlet $\mathcal{D}(\beta)$ distribution via the inverse log-ratio transform ℓ^{-1} , the problem of finding the best Gaussian approximation to $\mathcal{LD}(\beta)$ is equivalent to finding the best approximation to $\mathcal{D}(\beta)$ among a class of distributions obtained by applying the logistic transform to Gaussian random variables. If $\theta \sim N(\mu, \Sigma)$, the induced distribution on $\pi = \ell(\theta)$ is called a logistic normal distribution – denoted $\mathcal{L}(\mu, \Sigma)$ – and has density on \mathcal{S}^d given by

$$\tilde{q}(\pi; \mu, \Sigma) = (2\pi)^{-d/2} |\Sigma|^{-1/2} \left(\prod_{j=0}^d \pi_j \right)^{-1} \exp \left[-\frac{1}{2} \{ \log(\pi/\pi_0) - \mu \}^\top \Sigma^{-1} \{ \log(\pi/\pi_0) - \mu \} \right]. \quad (11)$$

The problem therefore boils down to calculating the Kullback–Leibler divergence between a Dirichlet density $q(\cdot; \beta)$ and a logistic normal density $\tilde{q}(\cdot; \mu, \Sigma)$ and optimizing the expression with respect to μ and Σ . The details are deferred to Section 6.

Once the approximation is derived in the identity parametrization, we appeal to the invariance of the Kullback–Leibler divergence under one-to-one transformations to obtain the corresponding approximation in a non-identity parameterization $\theta = X\tilde{\theta}$ as in (5) for any non-singular X . The result is stated below.

Corollary 1. *If $\theta \sim \mathcal{LD}(\beta)$ then*

$$D(\mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(\cdot; X^{-1}\mu^*, X^{-1}\Sigma^*(X^{-1})^T)) = \inf_{\mu, \Sigma} D(\mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(\cdot; \mu, \Sigma)) \quad (12)$$

for any full-rank d by d matrix X . Moreover, the bound on the KL divergence as a function of β in (10) is attained for $D(\mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(\cdot; \mu^*, \Sigma^*))$

Thus, the best Gaussian approximation to the posterior (in the Kullback–Leibler sense) under the Diaconis–Ylviaker prior is given by $\mathcal{N}(X^{-1}\mu^*, X^{-1}\Sigma^*(X^{-1})^T)$ for any one-to-one linear transformation $\theta = X^{-1}\tilde{\theta}$. We refer to this as the optimal Normal (oN) approximation. In Section 7, we provide a series of simulation studies showing that the oN approximation outperforms the Laplace approximation empirically, and is accurate enough to serve as the basis of approximate inference for reasonable sample sizes.

4 Application: estimating pairwise dependence

One potential use of the approximation is to estimate pairwise dependence from 2-way marginal tables. Because the approximation is normal, one can also apply standard

methods for multiplicity control to obtain sparse solutions. To assess the usefulness of the approximation in this context, we perform a simulation study in which a sparse graphical log-linear model corresponding to the cell probabilities of a 2^8 table is generated in the following way:

1. Sample main effects from $\mathcal{N}(0, 1)$
2. Sample two-way interactions from $0.5\delta_0 + 0.5\mathcal{N}(0, 1)$
3. For $j > 2$, if all corresponding lower-order interactions are nonzero, sample a j -way interaction from $0.5\delta_0 + 0.5\mathcal{N}(0, 1)$. This ensures that the resulting model is graphical.

We then compute the true value of Cramer's V

$$\rho_V = \sum_{c_j=0}^1 \sum_{c_k=0}^1 \frac{(\pi_{c_j c_k}^{(jk)})^2}{\binom{j}{c_j} \pi_{c_j}^{(k)} \pi_{c_k}^{(j)}} - 1,$$

where $\pi^{(jk)} = \sum_{l \neq j, k} \sum_{c_l=1}^2 \pi_{c_1, c_2, \dots, c_p}$ are the true (j, k) marginal cell probabilities, and $\pi_c^{(j)} = \mathbb{P}[y_j = c]$. We then sample $N = 665$ observations from $\text{Multinomial}(N, \pi)$, compute the posterior approximation under the DY prior with prior hyperparameter $\alpha = 1/4$, and obtain posterior point estimates $\hat{\rho}_V$ for every pair of variables j, k . The choice of number of variables and sample size is identical to that of the Rochdale data analyzed in section 4.1. The simulation is repeated 100 times. Figure 1 shows the results, with true values of ρ_V plotted against $\hat{\rho}_V$. The estimates are quite accurate, and have correlation 0.94 with the truth.

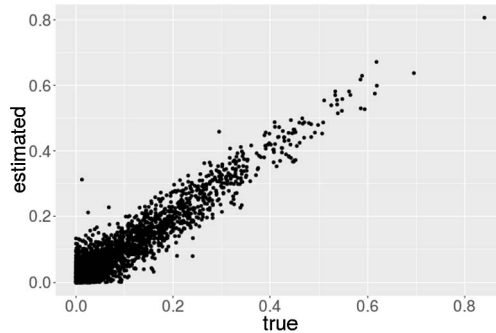


Figure 1: True values of ρ_V plotted against $\hat{\rho}_V$ estimated using posterior approximation under DY prior.

4.1 Analysis of pairwise dependence in Rochdale data

The Rochdale data is a well-studied dataset of eight binary variables. The purpose of the study, which was conducted on 665 households in Rochdale, UK, was to understand dependence among factors contributing to married women's economic activity. The eight

variables were wife employed (a), wife’s age > 38 (b), husband employed (c), one or more child < 4 years old (d), wife did not graduate high school (e), husband did not graduate high school (f), Asian origin (g), other household member (aside from husband and wife) working (h). This dataset was first analyzed by Whittaker (1990), and subsequently by numerous other authors. Here, we compare inference on marginal dependence using the optimal normal approximation to Copula Gaussian Graphical Models (CGGM), proposed by Dobra and Lenkoski (2011).

We follow the same procedure described above to obtain estimates of pairwise Cramer’s V and Gaussian approximation to the posterior under the DY prior with hyperparameter $\alpha = 1/4$ for two-way interaction effects in the corner parametrization for all of the $\binom{8}{2} = 28$ marginal tables. To obtain a sparse estimate of the interaction terms, we apply the method of Benjamini and Hochberg to Gaussian tail probabilities for the scaled two-way interaction terms $\theta_{jk}^*/\sigma_{jk}^*$ for every $j, k \in \{1, \dots, 8\}^2$, controlling FDR at 0.05. We view this as an exploratory effort to assess the statistical significance of the pairwise interactions. The scaled interaction terms, with horizontal lines indicating the threshold for significance, are shown in the left panel of Figure 2. In the right panel, $\hat{\rho}_V$ estimated using CGGM is plotted against $\hat{\rho}_V$ estimated using the normal approximation to the posterior under the DY prior, with variable pairs with interactions below the threshold indicated in blue. Clearly, there is strong agreement between the $\hat{\rho}_V$ in the two models, and the pairs excluded by the Benjamini–Hochberg procedure are those that correspond to the weakest interactions.




Figure 2: Left: values of $\theta_{jk}^*/\sigma_{jk}^*$. Right: estimated values of ρ_V from normal approximation plotted against estimated values of ρ_V from CGGMs.

The strongest pairwise interactions measured by $|\theta_{jk}^*/\sigma_{jk}^*|$ (with no thresholding) were

1. b:d (wife’s age > 38:wife did not graduate high school), $\hat{\theta}_{b:d}^*/\hat{\sigma}_{b:d}^* = -9.90$;
2. b:h (wife’s age > 38:other member of the household working), $\hat{\theta}_{b:h}^*/\hat{\sigma}_{b:h}^* = -8.58$;
3. a:d (wife employed:one or more child < 4 years old), $\hat{\theta}_{a:d}^*/\hat{\sigma}_{a:d}^* = -7.47$; and
4. e:f (wife did not graduate high school:husband did not graduate high school), $\hat{\theta}_{e:f}^*/\hat{\sigma}_{e:f}^* = 7.45$.

The magnitude (relative to the other facts) and signs make sense sociologically, keeping in mind the study was conducted several decades ago: older women were less likely to have graduated high school (b:h), in households where the woman is older, it is less likely that another member of the household works in addition to the husband (a:d),¹ the wife is less likely to work if there are young children at home (a:d); and people tend to sort by marriage into similar education levels (e:f). Both Whittaker (1990), pg 282, and Dobra and Lenkoski (2011), pg 984 list three of the four as among the top four strongest interactions (b:d, b:h, and e:f).

The four *weakest* effects were

1. b:f (wife's age > 38:husband did not graduate high school), $\hat{\theta}_{b:f}^*/\hat{\sigma}_{b:f}^* = 1.65$;
2. d:f (one or more child < 4 years old:husband did not graduate high school), $\hat{\theta}_{d:f}^*/\hat{\sigma}_{d:f}^* = -0.82$.
3. b:c (wife's age > 38:husband employed), $\hat{\theta}_{b:c}^*/\hat{\sigma}_{b:c}^* = 0.70$; and
4. c:h (husband employed:other member of the household working), $\hat{\theta}_{c:h}^*/\hat{\sigma}_{c:h}^* = 0.39$.

4.2 Posterior approximation with a 2^{16} table: National Long-Term Care Survey (NLTC) data

In principle, the normal approximation to the full posterior distribution for all of the log-linear model parameters can be used to perform inference on conditional independence as well as marginal independence. The main obstacle to this is that usually, sample sizes are insufficient for the posterior to be sufficiently concentrated that estimates of interaction terms in the corner parametrization are sufficiently low-variance to make the estimates reliable. In particular, since the full posterior distribution on the log-linear parameters does not enforce graphical or hierarchical constraints on the parameters of the log-linear model, a single noisy estimate of a higher-order interaction term can result in erroneously estimating large groups of variables to be conditionally dependent, when in fact many of them are conditionally independent of others given only a few of the variables. One could imagine numerous extensions to our approximation to combat this problem, for example, by using lower-dimensional parametrizations of the log-linear model.

Our aim in this section is only to show that the hurdle to applying the approximation to high-dimensional tables is not computational, but rather inferential. We show that the approximation can be computed for the 2^{16} table of NLTC data (see Dobra and Lenkoski (2011)) in minimal computation time (< 5 minutes on a circa 2014 MacBook Pro). A key to computing the approximation is that the $2^{16} \times 2^{16}$ matrix X is very sparse binary, and so is its inverse. It is likely that an explicit expression for the inverse can be found, but the inverse is so structured that it is numerically invertible in Matlab

¹Here, the wife's age is probably a proxy for the husband's age, with older parents more likely to have a traditional household in which only the husband works

in under one minute, so we simply used the numerical inverse. Figure 3 shows all 120 estimated values $\theta_{jk}^*/\sigma_{jk}^*$ – estimated using the marginal tables – and a histogram of θ^*/σ^* from the Gaussian approximation under the DY prior with $\alpha = 1/2$ for all 65,535 parameters of the log-linear model in the corner parametrization. The vertical lines in the latter indicate thresholds of significance using the Benjamini–Hochberg FDR procedure. Granular interpretation is difficult, but some general interpretation of the variables is evident. First, there appears to be weaker pairwise dependence between the first six variables than among the other variables. There are actually two distinct variable groups in the data: the first six are “Activities of Daily Life,” and the latter 10 “Instrumental activities of daily living,” so the stronger dependence between variables in the second group may be scientifically meaningful. Another general interpretation is that over half of the parameter estimates exceed the threshold, which suggests high dependence. Indeed, the median probability CGGM of Dobra and Lenkoski (2011) corresponds to a graphical model that is less than 50 percent sparse, and has no marginal independence, suggesting complex dependence between the 16 variables. Matlab code for the analysis is available at <https://github.com/jamesjohndrow/dynormal-approx>.




Figure 3: Left: value of $|\theta_{jk}^*/\sigma_{jk}^*|$ for every pair of variables in the NLTC data. Right: histogram of θ^*/σ^* for all 65,536 parameters in the log-linear model for NLTC data.

5 Discussion

Outside of linear models, conjugate priors are often non-standard or their multivariate generalizations are difficult to work with. This hampers uncertainty quantification because it is difficult to obtain posterior credible regions for parameters under such priors. Given that automatic and coherent quantification of uncertainty through the posterior is one of the chief advantages of a fully Bayesian approach, this limitation is a significant problem. The optimal Gaussian approximation to the posterior for log-linear models with Diaconis–Ylvisaker conjugate priors derived here offers a highly accurate and essentially computation-free approximation to posterior credible regions for this important

class of models. Interestingly, this Gaussian approximation is not the Laplace approximation, and it is faster to compute and offers a better approximation to the posterior than the Laplace approximation. If similar results could be obtained for the posterior in other models, it suggests that the Laplace approximation may not be an appropriate default Gaussian approximation to the posterior. The theoretical result provided here can be easily extended to cases where some categories cannot co-occur, i.e. cases of structural zeros in contingency tables. Extensions to model selection using our approximation are also available by the penalized credible region approach. It seems reasonable that the strategy used here to obtain optimality and convergence rate guarantees could be extended to a larger class of generalized linear models by studying the properties of multivariate Gaussian distributions under inverse link transformations. This may also present a strategy for obtaining approximate credible intervals for parameters in the Bayesian model averaging context for generalized linear models with conjugate priors.

6 Proofs

6.1 Additional log-linear model details

The discussion here largely follows Massam et al. (2009) and Lauritzen (1996) in its presentation. Let V be the set of variables that will be collected into a contingency table. Let $\mathcal{I}_\gamma, \gamma \in V$ denote the set of possible levels of values of γ . Without loss of generality, we can take this set to be a finite collection of sequential nonnegative integers. Let $\mathcal{I} = \times_{\gamma \in V} \mathcal{I}_\gamma$ be the set of all possible combinations of levels of the variables in V . Every cell i of the contingency table corresponds to an element of V ; thus $|\mathcal{I}| = d + 1$, where d is defined as in the main text.

Following Lauritzen (1996), define a cell of the contingency table as $i = (i_\gamma, \gamma \in V)$, and let $\pi(i) = \text{pr}[y_1 = i_1, \dots, y_p = i_p]$. For any $E \subset V$, let $i_E = (i_\gamma, \gamma \in E)$ be the cell of the E -marginal table corresponding to the values in i of the variables in E . Finally, designate the “base” cell $i^* = (0, 0, \dots, 0)$. Thus, every i can be written as $i = (i_E, i_{E^c}^*)$, where E is the subset of V on which $i \neq 0$. Then, the log-linear model in the corner parametrization is given by

$$\log \frac{\pi(i_E, i_{E^c}^*)}{\pi(i^*)} = \sum_{F \subseteq_{\neq \emptyset} E} \theta_F(i_F),$$

where for any $F \subset V$, $\theta_F(i_F)$ is a parameter corresponding the variables in F taking the values in i_F , and the notation $\subseteq_{\neq \emptyset}$ means all subsets excluding the empty set. Refer to Proposition 2.1 in Letac and Massam (2012) for a result showing how the model can be expressed in the form in (5).

6.2 Proof of Proposition 1

This is readily seen by the change of variable theorem; one only needs some work to calculate the Jacobian term for the change of variable. The matrix of partial derivatives

$J = (\partial\theta_j/\partial\pi_r)_{jr}$ is given by

$$\frac{\partial\theta_j}{\partial\pi_j} = \frac{1 - \sum_{l \neq j} \pi_l}{\pi_j(1 - \sum_{l=1}^d \pi_l)}, \quad \frac{\partial\theta_j}{\partial\pi_r} = -\frac{1}{1 - \sum_{l=1}^d \pi_l}, \quad (1 \leq j \neq r \leq d).$$

Write $J = U + uu^\top$, where $u = (1 - \sum_{l=1}^d \pi_l)^{-1/2}(1, -1, \dots, -1)^\top$ and $U = \text{Diag}(1/\pi_1, \dots, 1/\pi_d)$. We then have $|J| = |U|(1 + u^\top U^{-1}u)$ and therefore,

$$|J|^{-1} = \pi_1 \dots \pi_d \left(1 - \sum_{l=1}^d \pi_l\right) = \frac{e^{\sum_{i=1}^d \theta_i}}{(1 + \sum_{l=1}^d e^{\theta_l})^{d+1}}.$$

The proof is concluded by noting that $p(\theta; \alpha) = q(\ell(\theta); \alpha) |J|^{-1}$. \square

6.3 Proof of main results

We first state some preparatory results that are used to prove the main results.

Preliminaries. The following identity for the Gamma function is well known (see, e.g., Abramowitz and Stegun (1964)). For $z > 0$,

$$\log \Gamma(z) = \frac{\log(2\pi)}{2} + \left(z - \frac{1}{2}\right) \log z - z + R(z), \quad (13)$$

where $0 < R(z) < 1/(12z)$.

The digamma function $\psi(z) = \frac{d}{dz} \log \Gamma(z) = \frac{\Gamma'(z)}{\Gamma(z)}$ satisfies $\psi(z+1) = \psi(z) + 1/z$ for any $z > 0$. We use the following bound for the digamma function from Lemma 1 of Chen and Qi (2003). For any $z > 0$,

$$\frac{1}{2z} - \frac{1}{12z^2} < \psi(z+1) - \log z < \frac{1}{2z}. \quad (14)$$

The trigamma function $\psi'(z) = \frac{d}{dz} \psi(z)$ is the derivative of the digamma function. We derive a simple bound for the trigamma function that is used in the sequel.

Lemma 1. For any $z > 1/3$,

$$\frac{1}{z} < \psi'(z) < \frac{1}{z} + \frac{1}{z^2}. \quad (15)$$

The condition $z > 1/3$ is only required for the upper bound.

Proof. From Chen and Qi (2003), the trigamma function admits a series expansion

$$\psi'(z) = \sum_{j=0}^{\infty} \frac{1}{(z+j)^2}$$

valid for any $z > 0$. The function $t \mapsto t^{-2}$ is monotonically decreasing on $(0, \infty)$ and hence $x^{-2} > \int_x^{x+1} t^{-2} dt$ for any $x > 0$. Therefore, for any $z > 0$, $\psi'(z) >$

$\sum_{j=0}^{\infty} \int_{z+j}^{z+j+1} t^{-2} dt = \int_z^{\infty} t^{-2} dt = z^{-1}$. For the upper bound, we use Lemma 1 of Chen and Qi (2003) which states that $1/z - \psi'(z+1) > 1/(2z^2) - 1/(6z^3)$ for any $z > 0$. Since $\psi(z+1) = \psi(z) + 1/z$, $\psi'(z+1) = \psi'(z) - 1/z^2$, which yields $\psi'(z) - 1/z < 1/z^2 - 1/(2z^2) + 1/(6z^3) = 1/(2z^2) + 1/(6z^3)$ for any $z > 0$. The conclusion follows since $1/(6z^3) < 1/(2z^2)$ for any $z > 1/3$. \square

Finally, we state a useful result in Lemma 2.

Lemma 2. *Let $X \in \mathbb{R}^d$ be a random vector with $EX = \mu_X$ and $\text{var}(X) = \Sigma_X$. For $\mu \in \mathbb{R}^d$ and $d \times d$ positive definite matrix Σ , the mapping*

$$(\mu, \Sigma) \mapsto g(\mu, \Sigma) = \log |\Sigma| + E(X - \mu)^T \Sigma^{-1} (X - \mu) \tag{16}$$

attains its minima when $\mu = \mu_X$ and $\Sigma = \Sigma_X$. The minimum value of the objective function $g(\mu_X, \Sigma_X) = \log |\Sigma_X| + d$.

Proof. To start with, $E\{(X - \mu_X)^T \Sigma_X^{-1} (X - \mu_X)\} = \text{tr}[E\{(X - \mu_X)(X - \mu_X)^T \Sigma_X^{-1}\}] = \text{tr}(\mathbf{I}_d) = d$ and hence $g(\mu_X, \Sigma_X) = \log |\Sigma_X| + d$. Fix $\mu \in \mathbb{R}^d$ and Σ positive definite. We can write

$$\begin{aligned} E\{(X - \mu)\Sigma^{-1}(X - \mu)\} &= \text{tr}[E\{(X - \mu)(X - \mu)^T \Sigma^{-1}\}] \\ &= \text{tr}[E\{(X - \mu_X)(X - \mu_X)^T \Sigma^{-1}\} + (\mu_X - \mu)\Sigma^{-1}(\mu_X - \mu)] \\ &= \text{tr}(\Sigma_X \Sigma^{-1}) + (\mu_X - \mu)^T \Sigma^{-1}(\mu_X - \mu). \end{aligned}$$

Therefore,

$$g(\mu, \Sigma) - g(\mu_X, \Sigma_X) = \text{tr}(\Sigma_X \Sigma^{-1}) + (\mu_X - \mu)^T \Sigma^{-1}(\mu_X - \mu) - d - \log |\Sigma_X \Sigma^{-1}|.$$

The above quantity is non-negative since it equals $2D\{N(\mu_X, \Sigma_X) \parallel N(\mu, \Sigma)\}$, i.e., twice the Kullback–Leibler divergence between $N(\mu_X, \Sigma_X)$ and $N(\mu, \Sigma)$. Since μ and Σ were arbitrary, the first part is proved. The second part has been already proved at the beginning. \square

Proof of Theorem 1 and Corollary 1. We can now give a proof of Theorem 1. Recall the Dirichlet density q from (6) and the logistic normal density \tilde{q} from (11). We shall write $q(\pi)$ and $\tilde{q}(\pi)$ in place of $q(\pi \mid \beta)$ and $\tilde{q}(\pi \mid \mu, \Sigma)$ henceforth for brevity. From (6) and (11),

$$\begin{aligned} \log \frac{q(\pi)}{\tilde{q}(\pi)} &= \log B_\beta + \frac{d \log(2\pi)}{2} + \sum_{j=0}^d \beta_j \log \pi_j + \frac{\log |\Sigma|}{2} \\ &\quad + \frac{1}{2} \{ \log(\pi/\pi_0) - \mu \}^T \Sigma^{-1} \{ \log(\pi/\pi_0) - \mu \}, \end{aligned}$$

where

$$B_\beta = \frac{\Gamma(\sum_{j=0}^d \beta_j)}{\prod_{j=0}^d \Gamma(\beta_j)}.$$

Observe that μ and Σ appear only in the last two terms in the right hand side of the above display. Invoking Lemma 2, it is therefore evident that $D(q \parallel \tilde{q}) = E_q \log(q/\tilde{q})$ is minimized when $\mu^* = E_q \log(\pi/\pi_0)$ and $\Sigma^* = \text{var}_q\{\log(\pi/\pi_0)\}$, and the minimum value of the Kullback–Leibler divergence is

$$\log B_\beta + \sum_{j=0}^d \beta_j E_q \log \pi_j + \frac{d}{2} \{1 + \log(2\pi)\} + \frac{\log |\Sigma^*|}{2}. \quad (17)$$

Using standard properties of the Dirichlet distribution or Exponential family differential identities, with $\beta = \sum_{j=0}^d \beta_j$,

$$E_q \log \pi_j = \psi(\beta_j) - \psi(\beta), \quad j = 0, 1, \dots, d, \quad (18)$$

$$\text{cov}_q(\log \pi_j, \log \pi_l) = \psi'(\beta_j) \delta_{jl} - \psi'(\beta), \quad j, l = 0, 1, \dots, d. \quad (19)$$

Therefore, $\mu_j^* = E_q \log \pi_j - E_q \log \pi_0 = \psi(\beta_j) - \psi(\beta_0)$ for $j = 1, \dots, d$. Next, $\sigma_{jj'}^* = \text{cov}_q(\log \pi_j - \log \pi_0, \log \pi_{j'} - \log \pi_0) = \delta_{jj'} \psi'(\beta_j) + \psi'(\beta_0)$ for $j, j' = 1, \dots, d$. The expressions for μ^* and Σ^* are identical to (8), proving the first part of the theorem. Note this also establishes Proposition 2.

We now proceed to bound each term in the expression for the minimum Kullback–Leibler divergence in (17); refer to them by T_1, T_2, T_3 and T_4 respectively. First, we have,

$$\begin{aligned} T_1 &:= \log B_\beta = \log \Gamma(\beta) - \sum_{j=0}^d \log \Gamma(\beta_j) \\ &< -\frac{d \log(2\pi)}{2} + \left(\beta \log \beta - \sum_{j=0}^d \beta_j \log \beta_j \right) - \frac{1}{2} \left(\log \beta - \sum_{j=0}^d \log \beta_j \right) + \frac{1}{12\beta}. \end{aligned} \quad (20)$$

In the above display, we used (13) to bound $\log \Gamma(\beta)$ from above and $\log \Gamma(\beta_j)$ s from below. The $(-\beta)$ term in upper bound to $\log \Gamma(\beta)$ cancels out the $(-\sum_{j=0}^d \beta_j)$ contribution from the lower bounds to the $\log \Gamma(\beta_j)$ s. Next,

$$\begin{aligned} T_2 &:= \sum_{j=0}^d \beta_j E_q \pi_j = \sum_{j=0}^d \beta_j \{\psi(\beta_j) - \psi(\beta)\} \\ &= \sum_{j=0}^d \beta_j \{\psi(\beta_{j+1}) - \psi(\beta + 1)\} - \sum_{j=0}^d \beta_j \left(\frac{1}{\beta_j} - \frac{1}{\beta} \right) \\ &= \left\{ \sum_{j=0}^d \beta_j \psi(\beta_{j+1}) - \beta \psi(\beta) \right\} - d \\ &< \left(\sum_{j=0}^d \beta_j \log \beta_j - \beta \log \beta \right) - \frac{d}{2} + \frac{1}{12\beta}. \end{aligned} \quad (21)$$

In the first line of the above display, we used (18). From the first to the second line, we used the identity $\psi(z + 1) = \psi(z) + 1/z$. From the second to the third line, we only use $\sum_{j=0}^d \beta_j = \beta$. From the third to the fourth line, we made use of the bound (14) for the digamma function ψ . From the upper bound in (14), $\beta_j \psi(\beta_{j+1}) < \beta_j \log \beta_j + 1/2$ and hence $\sum_{j=0}^d \beta_j \psi(\beta_{j+1}) < \sum_{j=0}^d \beta_j \log \beta_j + (d + 1)/2$. From the lower bound in (14), $\beta \psi(\beta) > \beta \log \beta + 1/2 - 1/(12\beta)$.

Finally, from (19), we can write $\Sigma^* = D + \psi'(\beta_0)\mathbf{1}\mathbf{1}^\top$, with $D = \text{diag}(\psi'(\beta_1), \dots, \psi'(\beta_d))$. Using the fact $|X + uv^\top| = |X|(1 + v^\top X^{-1}u)$, we obtain

$$|\Sigma^*| = \left\{ 1 + \sum_{j=1}^d \psi'(\beta_0)/\psi'(\beta_j) \right\} \left\{ \prod_{j=1}^d \psi'(\beta_j) \right\} = \left\{ \sum_{j=0}^d \frac{\psi'(\beta_0)}{\psi'(\beta_j)} \right\} \left\{ \prod_{j=1}^d \psi'(\beta_j) \right\}.$$

From Lemma 1, $\psi'(\beta_j) > 1/\beta_j$, implying

$$\begin{aligned} T_4 &:= \frac{\log |\Sigma^*|}{2} = \frac{1}{2} \left[\log \left\{ \sum_{j=0}^d \frac{\psi'(\beta_0)}{\psi'(\beta_j)} \right\} + \sum_{j=1}^d \log \psi'(\beta_j) \right] \\ &< \frac{1}{2} \left\{ \log \beta + \sum_{j=0}^d \log \psi'(\beta_j) \right\}. \end{aligned} \tag{22}$$

Recalling $T_3 = d\{1 + \log(2\pi)\}/2$ and substituting the bounds for T_1, T_2 and T_4 from (20), (21) and (22) in (17), we obtain, after plenty of cancellations,

$$\begin{aligned} \sum_{j=1}^4 T_j &< \frac{1}{2} \sum_{j=0}^d \log\{\beta_j \psi'(\beta_j)\} + \frac{1}{6\beta} \\ &< \frac{1}{2} \sum_{j=0}^d \frac{1}{\beta_j} + \frac{1}{6\beta}. \end{aligned}$$

From the first to the second line, we invoked Lemma 1 to bound $\beta_j \psi'(\beta_j) < 1 + 1/\beta_j$ and used $\log(1 + x) < x$ for $x > 0$. We have obtained the desired bound, concluding the proof.

Now, to show Corollary 1, we make use of the fact that the KL divergence D is invariant under one-to-one transformations of θ to conclude that for any full rank matrix X ,

$$D\left\{ \mathcal{LD}(\beta) \parallel \mathcal{N}(\mu, \Sigma) \right\} = D\left\{ \mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(X\mu, X^\top \Sigma X) \right\}. \tag{23}$$

So

$$\inf_{\mu, \Sigma} \left\{ \mathcal{LD}(\beta) \parallel \mathcal{N}(\mu, \Sigma) \right\} = \inf_{\tilde{\mu}, \tilde{\Sigma}} D\left\{ \mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(\tilde{\mu}, \tilde{\Sigma}) \right\}. \tag{24}$$

Since the infimum on the left side in (24) is attained by μ^*, Σ^* , we have by (23) that

$$D(\mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(\cdot; X\mu^*, X^\top \Sigma^* X)) = \inf_{\mu, \Sigma} D(\mathcal{P}_X(\cdot; \beta) \parallel \mathcal{N}(\cdot; \mu, \Sigma)),$$

which gives Corollary 1.

7 Simulations assessing accuracy of approximation

We conducted several simulation studies to assess the performance of the approximation in Theorem 1 and Corollary 1. In each study, we simulated 100 realizations from

$$\pi \sim \mathcal{D}(a, \dots, a), \quad y \sim \text{Multinomial}(N, \pi), \quad (25)$$

with the posterior of π under a Dirichlet $\mathcal{D}(a, \dots, a)$ prior being $\mathcal{D}(y_1 + a, \dots, y_d + a)$. We chose the dimension d to be 2^8 , corresponding to a $p = 8$ -way contingency table for binary variables. To obtain a simulation-based approximation to the posterior for $\theta = \log(\pi/\pi_0)$ under the Diaconis–Ylvisaker prior, we sampled mc many π values from the $\mathcal{D}(y_1 + a, \dots, y_d + a)$ posterior and then transformed to $\theta = \ell^{-1}(\pi)$ to obtain posterior samples of θ ; we refer to this procedure as the Monte Carlo approximation. We also computed a Laplace approximation to the posterior under the Diaconis–Ylvisaker prior, which is given by Normal $(\hat{\theta}_{MAP}, \mathcal{I}(\hat{\theta}_{MAP})^{-1})$, where $\hat{\theta}_{MAP}$ is the *maximum a-posteriori* estimate of θ and $\mathcal{I}(\theta)$ is the Fisher information matrix evaluated at θ . The maximum a-posteriori estimate $\hat{\theta}_{MAP}$ was computed by the Newton–Raphson method.

We compare the accuracy of the proposed Gaussian approximation to the Monte Carlo procedure and the Laplace approximation. In addition to the identity parameterization, i.e., $X = I_d$ in (5), we also consider the corner parameterization given by $\log(\pi/\pi_0) = X\tilde{\theta}$ for an appropriate X matrix; see Section 6.1 for more details. For the Monte Carlo samples, each sample of θ is transformed to $\tilde{\theta}$ via $X^{-1}\theta = \tilde{\theta}$. For the normal approximations $\theta \sim \text{Normal}(\mu, \Sigma)$, the corresponding approximate posterior is given by $\tilde{\theta} \sim \text{Normal}(X^{-1}\mu, X^{-1}\Sigma X^{-1})$.

We conduct simulations for different values of N (250, 1000, and 10,000) and a (1 and $1/d$). We then assess performance in several ways.

- Proportion of variation unexplained, measured by $\sqrt{\sum_{j=1}^d (\theta - \theta_0)^2 / \text{sd}(\theta_0)}$, where θ_0 is the true value of θ (or $\tilde{\theta}$, as appropriate).
- Coverage of 95 percent posterior credible intervals for θ or $\tilde{\theta}$.
- The standardized loss in the Frobenius norm for estimates of Σ , the posterior covariance, given by $\|\hat{\Sigma} - \Sigma\|_F / \|\Sigma\|_F$, where $\|S\|_F$ is the Frobenius norm of S . Note that the covariance in Theorem 1 is exactly the posterior covariance, so this measure is computed only for the simulation and Laplace approximations.
- The value of the Kolmogorov–Smirnov statistic for comparing the Monte Carlo empirical measure $\frac{1}{mc} \sum_{t=1}^{mc} \delta_{\theta_t}$ to the normal approximation from Theorem 1, Normal (μ, Σ) .
- The computation time required to compute each posterior approximation.

Table 1 shows unexplained variation for the Laplace approximation, the Monte Carlo approximation for $mc = 10^3, 10^4, 10^5$, and 10^6 , and the optimal normal approximation.

As expected, the optimal normal approximation outperforms the Laplace approximation. Moreover, it is comparable to the Monte Carlo approximation at every sample size and for all of the values of mc considered. Performance for all approximations is noticeably better in the corner parametrization than the identity parametrization.

	Laplace	$mc = 10^3$	$mc = 10^4$	$mc = 10^5$	$mc = 10^6$	oN
identity, N=250	1.08	0.98	0.98	0.98	0.98	0.98
corner, N=250	0.85	0.81	0.81	0.81	0.81	0.81
identity, N=1000	0.84	0.77	0.77	0.77	0.77	0.77
corner, N=1000	0.67	0.61	0.61	0.61	0.61	0.61
identity, N=10,000	0.40	0.35	0.35	0.35	0.35	0.35
corner, N=10,000	0.31	0.27	0.27	0.27	0.27	0.27

Table 1: $\sqrt{\sum_{j=1}^d (\theta - \theta_0)^2 / \text{sd}(\theta_0)}$ for different values of mc , different sample sizes, and two parametrizations. Results are averaged over 100 replicate simulations for each sample size.

Table 2 shows coverage of approximate 95 percent credible intervals for the Laplace approximation, optimal Normal approximation, and the Monte Carlo approximation. The intervals derived using the Laplace approximation are universally too wide. Nominal coverage for the Monte Carlo approximation is insensitive to the value of mc in the range tested, and is slightly high at the two smaller sample sizes. The optimal normal approximation has the best coverage; in all cases it is between 0.94 and 0.96 and for $N = 10,000$ the coverage is 0.95 in both parametrizations.

	Laplace	$mc = 10^3$	$mc = 10^4$	$mc = 10^5$	$mc = 10^6$	oN
identity, N=250	0.95	0.97	0.97	0.97	0.97	0.96
corner, N=250	1.00	0.96	0.96	0.96	0.96	0.96
identity, N=1000	0.98	0.96	0.96	0.96	0.96	0.96
corner, N=1000	1.00	0.94	0.94	0.94	0.94	0.94
identity, N=10,000	1.00	0.95	0.95	0.95	0.95	0.95
corner, N=10,000	1.00	0.95	0.95	0.95	0.95	0.95

Table 2: Coverage of 95% posterior credible intervals.

Table 3 shows dependence of $\|\widehat{\Sigma} - \Sigma\|_F / \|\Sigma\|_F$ on mc for the two different parametrizations and three sample sizes considered. Note that Σ is known exactly since $\Sigma = \Sigma^*$, the posterior covariance under the DY prior. The main point of this table is to demonstrate the relatively large number of Monte Carlo samples required to obtain reasonably small error in estimation of the posterior covariance. Even with 10^5 samples the relative error is on the 1 percent range. Thus, compound linear hypothesis testing and computation of credible regions is very inefficient using the Monte Carlo method.

Table 4 shows the computation time in seconds for each of the three approximations. The Laplace approximation is fast, requiring about 0.03-0.04 seconds to compute at all sample sizes. The optimal normal approximation is about an order of magnitude faster, with the computation time arising mainly in computing the polygamma functions

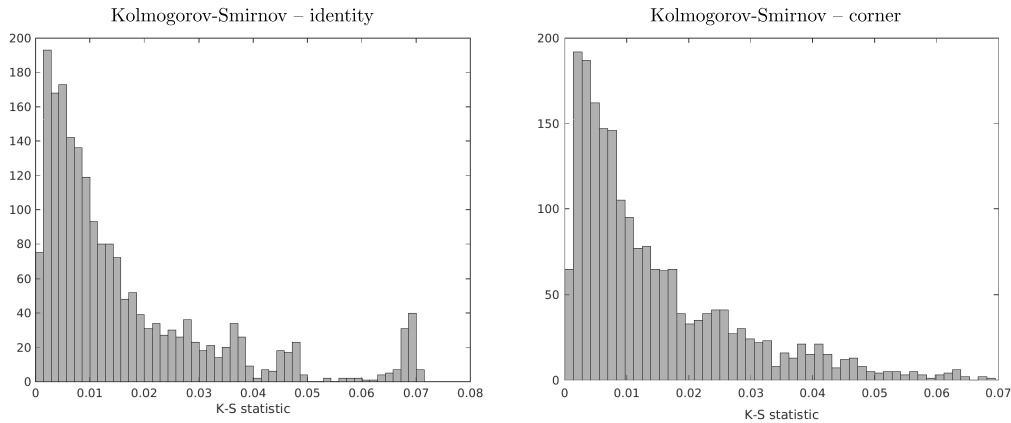
	$mc = 10^3$	$mc = 10^4$	$mc = 10^5$	$mc = 10^6$
identity, N=250	0.0982	0.0328	0.0093	0.0032
corner, N=250	0.0923	0.0290	0.0086	0.0029
identity, N=1000	0.1045	0.0330	0.0103	0.0035
corner, N=1000	0.0882	0.0277	0.0087	0.0029
identity, N=10,000	0.1231	0.0397	0.0118	0.0040
corner, N=10,000	0.0861	0.0280	0.0084	0.0027

Table 3: $\|\widehat{\Sigma} - \Sigma\|_F / \|\Sigma\|_F$ for different sample sizes and values of mc .

	Laplace	$mc = 10^6$	oN
N=250	0.037	32.652	0.003
N=1000	0.031	31.980	0.003
N=10,000	0.035	32.338	0.003

Table 4: Average time (seconds) to compute each approximation, averaged over 100 replicate simulations for each sample size.

and matrix multiplications. The Monte Carlo approximation is about four orders of magnitude slower than the optimal Normal approximation. Here, only $mc = 10^6$ is considered because of the non-negligible error in the posterior covariance for smaller samples; the algorithm scales linearly in mc so for $mc = 10^5$ the required time would be approximately 3 seconds. Only about 100 samples could be obtained in the 0.003 seconds required to compute the optimal normal approximation.

Figure 4: Distribution of Kolmogorov-Smirnov statistics comparing $\frac{1}{mc} \sum_{t=1}^{mc} \delta_{\theta_t}$ to the oN approximation for 20 randomly selected entries of θ and over 100 replicate simulations (entries of θ were re-selected for each replicate).

Results in the previous tables make clear that the optimal normal approximation is superior to the other approximations considered in terms of point estimation, estimation of 95 percent credible intervals, covariance estimation, and computation time.

However, it is possible that differences between the optimal normal approximation and the exact posterior exist in the tails of the distribution. To assess this, we compare the empirical measure of the Monte Carlo approximation using $mc = 10^6$ samples to the optimal normal approximation by computing the Kolmogorov–Smirnov (KS) statistic for the marginal distributions of 20 randomly selected entries of θ . The entries considered were re-selected for each of the 100 replicate simulations and for each of the three sample sizes. Shown in Figure 4 are histograms of these KS statistics in the corner and identity parametrizations. Most are less than 0.02, and none are greater than 0.07. Considering that the KS statistic is a point estimate of the total variation distance between distributions, this indicates that the optimal normal approximation is an excellent approximation to the posterior marginals. Moreover, we cannot rule out the possibility of residual Monte Carlo error in the marginals from the Monte Carlo approximation, which may account for part of the observed discrepancy.

References

- Abramowitz, M. and Stegun, I. A. (1964). *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. 55. Courier Corporation. 214
- Agresti, A. (2002). *Categorical Data Analysis*, volume 359. John Wiley & Sons. MR1914507. doi: <https://doi.org/10.1002/0471249688>. 201
- Attias, H. (1999). “Inferring parameters and structure of latent variable models by variational Bayes.” In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 21–30. Morgan Kaufmann Publishers Inc. 202
- Bishop, Y. M., Fienberg, S. E., and Holland, P. W. (2007). *Discrete Multivariate Analysis: Theory and Practice*. Springer Science & Business Media. MR2344876. 201
- Chen, C.-P. and Qi, F. (2003). “The best lower and upper bounds of harmonic sequence.” *RGMA Research Report Collection*, 6(2). 214, 215
- Consonni, G., Veronese, P., and Gutiérrez-Peña, E. (2004). “Reference priors for exponential families with simple quadratic variance function.” *Journal of Multivariate Analysis*, 88(2): 335–364. 205
- Dellaportas, P. and Forster, J. J. (1999). “Markov chain Monte Carlo model determination for hierarchical and graphical log-linear models.” *Biometrika*, 86(3): 615–633. 202
- Diaconis, P. and Ylvisaker, D. (1979). “Conjugate priors for exponential families.” *The Annals of Statistics*, 7(2): 269–281. 202, 203
- Dobra, A. and Lenkoski, A. (2011). “Copula Gaussian graphical models and their application to modeling functional disability data.” *The Annals of Applied Statistics*, 5(2A): 969–993. MR2840183. doi: <https://doi.org/10.1214/10-A0AS397>. 210, 211, 212
- Dobra, A. and Massam, H. (2010). “The mode oriented stochastic search (MOSS) algo-

- rithm for log-linear models with conjugate priors.” *Statistical Methodology*, 7(3): 240–253. MR2643600. doi: <https://doi.org/10.1016/j.stamet.2009.04.002>. 202
- Fienberg, S. E. and Rinaldo, A. (2007). “Three centuries of categorical data analysis: Log-linear models and maximum likelihood estimation.” *Journal of Statistical Planning and Inference*, 137(11): 3430–3445. MR2363267. doi: <https://doi.org/10.1016/j.jspi.2007.03.022>. 201
- Gelfand, A. E. and Smith, A. F. (1990). “Sampling-based approaches to calculating marginal densities.” *Journal of the American Statistical Association*, 85(410): 398–409. 202
- Gutiérrez-Pena, E. and Smith, A. (1995). “Conjugate parameterizations for natural exponential families.” *Journal of the American Statistical Association*, 90(432): 1347–1356. 205
- Haberman, S. J. (1974). “Log-linear models for frequency tables derived by indirect observation: Maximum likelihood equations.” *The Annals of Statistics*, 911–924. 201
- Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1998). “Bayesian model averaging.” In *In Proceedings of the AAAI Workshop on Integrating Multiple Learned Models*, 77–83. Citeseer. 202
- Lauritzen, S. L. (1996). *Graphical models*. Oxford University Press. 213
- Letac, G. and Massam, H. (2012). “Bayes factors and the geometry of discrete hierarchical loglinear models.” *The Annals of Statistics*, 40(2): 861–890. MR2985936. doi: <https://doi.org/10.1214/12-AOS974>. 213
- Massam, H., Liu, J., and Dobra, A. (2009). “A conjugate prior for discrete hierarchical log-linear models.” *The Annals of Statistics*, 37(6A): 3431–3467. 202, 204, 213
- Park, M. Y. and Hastie, T. (2007). “L1-regularization path algorithm for generalized linear models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4): 659–677. 201
- 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. MR3174712. doi: <https://doi.org/10.1080/01621459.2013.829001>. 202
- Shun, Z. and McCullagh, P. (1995). “Laplace approximation of high dimensional integrals.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 749–760. MR1354079. 202
- Tierney, L. and Kadane, J. B. (1986). “Accurate approximations for posterior moments and marginal densities.” *Journal of the American Statistical Association*, 81(393): 82–86. 202
- Wang, B. and Titterton, D. (2004). “Lack of consistency of mean field and variational Bayes approximations for state space models.” *Neural Processing Letters*, 20(3): 151–170. 202

- Wang, B. and Titterton, D. (2005). “Inadequacy of interval estimates corresponding to variational Bayesian approximations.” *Proc. 10th Int. Workshop Artificial Intelligence and Statistics*, 373–380. [202](#)
- Whittaker, J. (1990). “Graphical models in applied multivariate statistics.” [210](#), [211](#)
- Zou, H. and Hastie, T. (2005). “Regularization and variable selection via the elastic net.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2): 301–320. [201](#), [202](#)

Acknowledgments

The authors thank David Dunson for useful conversations and comments during the preparation of this manuscript.