
Frequentist Consistency of Variational Bayes

Yixin Wang

Department of Statistics
Columbia University

David M. Blei

Department of Computer Science and Statistics
Columbia University

Abstract

A key challenge for modern Bayesian statistics is how to perform scalable inference of posterior distributions. To address this challenge, variational Bayes (VB) methods have emerged as a popular alternative to the classical Markov chain Monte Carlo (MCMC) methods. VB methods tend to be faster while achieving comparable predictive performance. However, there are few theoretical results around VB. In this paper, we establish frequentist consistency and asymptotic normality of VB methods. Specifically, we connect VB methods to point estimates based on variational approximations, called frequentist variational approximations, and we use the connection to prove a variational Bernstein–Von Mises theorem. The theorem leverages the theoretical characterizations of frequentist variational approximations to understand asymptotic properties of VB. In summary, we prove that (1) the VB posterior converges to the Kullback–Leibler (KL) minimizer of a normal distribution, centered at the truth and (2) the corresponding variational expectation of the parameter is consistent and asymptotically normal. As applications of the theorem, we derive asymptotic properties of VB posteriors in Bayesian mixture models, Bayesian generalized linear mixed models, and Bayesian stochastic block models. We conduct a simulation study to illustrate these theoretical results.

Keywords: Bernstein–Von Mises, Bayesian inference, variational methods, consistency, asymptotic normality, statistical computing

1 Introduction

Bayesian modeling is a powerful approach for discovering hidden patterns in data. We begin by setting up a probability model of latent variables and observations. We incorporate prior knowledge by setting priors on latent variables and a functional form of the likelihood. Finally we infer the posterior, the conditional distribution of the latent variables given the observations.

For many modern Bayesian models, exact computation of the posterior is intractable and statisticians must resort to approximate posterior inference. For decades, Markov chain Monte Carlo (MCMC) sampling (Hastings, 1970; Gelfand and Smith, 1990; Robert and Casella, 2004) has maintained its status as the dominant approach to this problem. MCMC algorithms are easy to use and theoretically sound. In recent years, however, data sizes have soared. This challenges MCMC methods, for which convergence can be slow, and calls upon scalable alternatives. One popular class of alternatives is variational Bayes (VB) methods.

To describe VB, we introduce notation for the posterior inference problem. Consider observations $x = x_{1:n}$. We posit local latent variables $z = z_{1:n}$, one per observation, and global latent variables

$\theta = \theta_{1:d}$. This gives a joint,

$$p(\theta, z, x) = p(\theta) \prod_{i=1}^n p(z_i | \theta) p(x_i | z_i, \theta). \quad (1)$$

The posterior inference problem is to calculate the posterior $p(\theta, z | x)$.

This division of latent variables is common in modern Bayesian statistics.¹ In the Bayesian Gaussian mixture model (GMM) (Roberts et al., 1998), the component means, covariances, and mixture proportions are global latent variables; the mixture assignments of each observation are local latent variables. In the Bayesian generalized linear mixed model (GLMM) (Breslow and Clayton, 1993), the intercept and slope are global latent variables; the group-specific random effects are local latent variables. In the Bayesian stochastic block model (SBM) (Hofman and Wiggins, 2008), the cluster assignment probabilities and edge probabilities matrix are two sets of global latent variables; the node-specific cluster assignments are local latent variables. In the latent Dirichlet allocation (LDA) model (Blei et al., 2003), the topic-specific word distributions are global latent variables; the document-specific topic distributions are local latent variables. We will study all these examples below.

VB methods formulate posterior inference as an optimization (Jordan et al., 1999; Wainwright and Jordan, 2008; Blei et al., 2016). We consider a family of distributions of the latent variables and then find the member of that family that is closest to the posterior.

Here we focus on mean-field variational inference (though our results apply more widely). First, we posit a family of factorizable probability distributions on latent variables

$$\mathcal{Q}^{n+d} = \left\{ q : q(\theta, z) = \prod_{i=1}^d q_{\theta_i}(\theta_i) \prod_{j=1}^n q_{z_j}(z_j) \right\}.$$

This family is called *the mean-field family*. It represents a joint of the latent variables with $n + d$ (parametric) marginal distributions, $\{q_{\theta_1}, \dots, q_{\theta_d}, q_{z_1}, \dots, q_{z_n}\}$.

VB finds the member of the family closest to the exact posterior $p(\theta, z | x)$, where closeness is measured by KL divergence. Thus VB seeks to solve the optimization,

$$q^*(\theta, z) = \underset{q(\theta, z) \in \mathcal{Q}^{n+d}}{\operatorname{argmin}} \operatorname{KL}(q(\theta, z) || p(\theta, z | x)). \quad (2)$$

In practice, VB finds $q^*(\theta, z)$ by optimizing an alternative objective, *the evidence lower bound* (ELBO),

$$\operatorname{ELBO}(q(\theta, z)) = - \int q(\theta, z) \log \frac{q(\theta, z)}{p(\theta, z, x)} d\theta dz. \quad (3)$$

This objective is called the ELBO because it is a lower bound on the evidence $\log p(x)$. More importantly, the ELBO is equal to the negative KL plus $\log p(x)$, which does not depend on $q(\cdot)$. Maximizing the ELBO minimizes the KL (Jordan et al., 1999).

The optimum $q^*(\theta, z) = q^*(\theta)q^*(z)$ approximates the posterior, and we call it *the VB posterior*.² Though it cannot capture posterior dependence across latent variables, it has hope to capture each of their marginals. In particular, this paper is about the theoretical properties of the VB posterior $q^*(\theta)$, the VB posterior of θ . We will also focus on the corresponding expectation of the global variable, i.e., an estimate of the parameter. It is

$$\hat{\theta}_n^* := \int \theta \cdot q^*(\theta) d\theta.$$

We call θ^* *the variational Bayes estimate* (VBE).

¹In particular, our results are applicable to general models with *local* and *global* latent variables (Hoffman et al., 2013). The number of local variables z increases with the sample size n ; the number of global variables θ does not. We also note that the conditional independence of Equation (1) is not necessary for our results. But we use this common setup to simplify the presentation.

²For simplicity we will write $q(\theta, z) = \prod_{i=1}^d q(\theta_i) \prod_{j=1}^n q(z_j)$, omitting the subscript on the factors $q(\cdot)$. The understanding is that the factor is indicated by its argument.

VB methods are fast and yield good predictive performance in empirical experiments (Blei et al., 2016). However, there are few rigorous theoretical results. In this paper, we prove that (1) the VB posterior converges in total variation (TV) distance to the KL minimizer of a normal distribution centered at the truth and (2) the VBE is consistent and asymptotically normal.

These theorems are frequentist in the sense that we assume the data come from $p(z, x; \theta_0)$ with a true (nonrandom) θ_0 . We then study properties of the corresponding posterior distribution $p(\theta | x)$, when approximating it with variational inference. What this work shows is that the VB posterior is consistent even though the mean field approximating family can be a brutal approximation. In this sense, VB is a theoretically sound approximate inference procedure.

1.1 Main ideas

We describe the results of the paper. Along the way, we will need to define some terms: the variational frequentist estimate (VFE), the variational log likelihood, the VB posterior, the VBE, and the VB ideal. Our results center around the VB posterior and the VBE. (Table 1 contains a glossary of terms.)

The variational frequentist estimate (VFE) and the variational log likelihood. The first idea that we define is the *variational frequentist estimate* (VFE). It is a point estimate of θ that maximizes a local variational objective with respect to an optimal variational distribution of the local variables. (The VFE treats the variable θ as a parameter rather than a random variable.) We call the objective *the variational log likelihood*,

$$M_n(\theta; x) = \max_{q(z)} \mathbb{E}_{q(z)} [\log p(x, z | \theta) - \log q(z)]. \quad (4)$$

In this objective, the optimal variational distribution $q^\dagger(z)$ solves the local variational inference problem,

$$q^\dagger(z) = \operatorname{argmin}_q \text{KL}(q(z) || p(z | x, \theta)). \quad (5)$$

Note that $q^\dagger(z)$ implicitly depends on both the data x and the parameter θ .

With the objective defined, the VFE is

$$\hat{\theta}_n = \operatorname{argmax}_\theta M_n(\theta; x). \quad (6)$$

It is usually calculated with variational expectation maximization (EM) (Wainwright and Jordan, 2008; Ormerod and Wand, 2010), which iterates between the E step of Equation (5) and the M step of Equation (6). Recent research has explored the theoretical properties of the VFE for stochastic block models (Bickel et al., 2013), generalized linear mixed models (Hall et al., 2011), and Gaussian mixture models (Westling and McCormick, 2015).

We make two remarks. First, the maximizing variational distribution $q^\dagger(z)$ of Equation (5) is different from $q^*(z)$ in the VB posterior: $q^\dagger(z)$ is implicitly a function of individual values of θ , while $q^*(z)$ is implicitly a function of the variational distributions $q(\theta)$. Second, the variational log likelihood in Equation (4) is similar to the original objective function for the EM algorithm (Dempster et al., 1977). The difference is that the EM objective is an expectation with respect to the exact conditional $p(z | x)$, whereas the variational log likelihood uses a variational distribution $q(z)$.

Variational Bayes and ideal variational Bayes. While earlier applications of variational inference appealed to variational EM and the VFE, most modern applications do not. Rather they use VB, as we described above, where there is a prior on θ and we approximate its posterior with a global variational distribution $q(\theta)$. One advantage of VB is that it provides regularization through the prior. Another is that it requires only one type of optimization: the same considerations around updating the local variational factors $q(z)$ are also at play when updating the global factor $q(\theta)$.

To develop theoretical properties of VB, we connect the VB posterior to the variational log likelihood; this is a stepping stone to the final analysis. In particular, we define the VB *ideal posterior* $\pi^*(\theta | x)$,

$$\pi^*(\theta | x) = \frac{p(\theta) \exp\{M_n(\theta; x)\}}{\int p(\theta) \exp\{M_n(\theta; x)\} d\theta}. \quad (7)$$

Here the local latent variables z are constrained under the variational family but the global latent variables θ are not. Note that because it depends on the variational log likelihood $M_n(\theta; x)$, this distribution implicitly contains an optimal variational distribution $q^\dagger(z)$ for each value of θ ; see Equations (4) and (5).

Loosely, the VB ideal lies between the exact posterior $p(\theta|x)$ and a variational approximation $q(\theta)$. It recovers the exact posterior when $p(z|\theta, x)$ degenerates to a point mass and $q^\dagger(z)$ is always equal to $p(z|\theta, x)$; in that case the variational likelihood is equal to the log likelihood and Equation (7) is the posterior. But $q^\dagger(z)$ is usually an approximation to the conditional. Thus the VB ideal usually falls short of the exact posterior.

That said, the VB ideal is more complex than a simple parametric variational factor $q(\theta)$. The reason is that its value for each θ is defined by the optimization within $M_n(\theta; x)$. Such a distribution will usually lie outside the distributions attainable with a simple family.

In this work, we first establish the theoretical properties of the VB ideal. We then connect it to the VB posterior.

Variational Bernstein–Von Mises. We have set up the main concepts. We now describe the main results.

Suppose the data come from a true (finite-dimensional) parameter θ_0 . The classical Bernstein–Von Mises theorem says that, under certain conditions, the exact posterior $p(\theta|x)$ approaches a normal distribution, independent of the prior, as the number of observations tends to infinity. In this paper, we extend the theory around Bernstein–Von Mises to the variational posterior. Here we summarize our results.

- Lemma 1 shows that the VB ideal $\pi^*(\theta|x)$ is consistent and converges to a normal distribution around the VFE. If the VFE is consistent, the VB ideal $\pi^*(\theta|x)$ converges to a normal distribution whose mean parameter is a random vector centered at the true parameter. (Note the randomness in the mean parameter is due to the randomness in the observations x .)
- We next consider the point in the variational family that is closest to the VB ideal $\pi^*(\theta|x)$ in KL divergence. Lemma 2 and Lemma 3 show that this KL minimizer is consistent and converges to the KL minimizer of a normal distribution around the VFE. If the VFE is consistent (Bickel et al., 2013; Hall et al., 2011) then the KL minimizer converges to the KL minimizer of a normal distribution with a random mean centered at the true parameter.
- Lemma 4 shows that the VB posterior $q^*(\theta)$ enjoys the same asymptotic properties as the KL minimizers of the VB ideal $\pi^*(\theta|x)$.
- Theorem 5 is the *variational Bernstein–Von Mises theorem*. It shows that the VB posterior $q^*(\theta)$ is asymptotically normal around the VFE. Again, if the VFE is consistent then the VB posterior converges to a normal with a random mean centered at the true parameter. Further, Theorem 6 shows that the VBE $\hat{\theta}_n^*$ is consistent with the true parameter and asymptotically normal.
- Finally, we prove two corollaries. First, if we use a full rank Gaussian variational family then the corresponding VB posterior recovers the true mean and covariance. Second, if we use a mean-field Gaussian variational family then the VB posterior recovers the true mean and the marginal variance, but not the off-diagonal terms. The mean-field VB posterior is underdispersed.

Please refer to the appendix for full technical details and proofs.

References

- Bickel, P., Choi, D., Chang, X., Zhang, H., et al. (2013). Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *The Annals of Statistics*, 41(4):1922–1943.
- Blei, D., Kucukelbir, A., and McAuliffe, J. (2016). Variational inference: A review for statisticians. *Journal of American Statistical Association*.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Breslow, N. E. and Clayton, D. G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American statistical Association*, 88(421):9–25.
- Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39:1–38.
- 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.
- Hall, P., Pham, T., Wand, M. P., Wang, S. S., et al. (2011). Asymptotic normality and valid inference for Gaussian variational approximation. *The Annals of Statistics*, 39(5):2502–2532.
- Hastings, W. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97–109.
- Hoffman, M., Blei, D., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347.
- Hofman, J. and Wiggins, C. (2008). Bayesian approach to network modularity. *Physical Review Letters*, 100(25).
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233.
- Ormerod, J. T. and Wand, M. P. (2010). Explaining variational approximations. *The American Statistician*, 64(2):140–153.
- Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag, New York, NY.
- Roberts, S. J., Husmeier, D., Rezek, I., and Penny, W. (1998). Bayesian approaches to Gaussian mixture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1133–1142.
- Wainwright, M. J. and Jordan, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305.
- Westling, T. and McCormick, T. H. (2015). Establishing consistency and improving uncertainty estimates of variational inference through M-estimation. *arXiv preprint arXiv:1510.08151*.

Supplementary Materials

A Related work

This work draws on two themes. The first is the body of work on asymptotic properties of variational inference. [You et al. \(2014\)](#) and [Ormerod et al. \(2014\)](#) study variational Bayes for a classical Bayesian linear model. They use normal priors and spike-and-slab priors on the coefficients, respectively. [Wang and Titterton \(2005\)](#) and [Wang et al. \(2006\)](#) analyze variational Bayes in Bayesian mixture models with conjugate priors.

On the frequentist side, [Hall et al. \(2011a,b\)](#) establishes consistency of Gaussian variational EM estimates in a Poisson mixed-effects model with a single predictor and a grouped random intercept. [Celisse et al. \(2012\)](#) and [Bickel et al. \(2013\)](#) proved asymptotic normality of parameter estimates in the SBM under a mean field variational approximation. [Westling and McCormick \(2015\)](#) study the consistency of variational EM estimates in mixture models through a connection to M-estimation.

However, all these treatments of variational methods—either in a Bayesian or frequentist setting—are constrained to specific models and priors. Our work broadens this work by considering more general models. Moreover, the frequentist work focuses on estimation procedures under a variational approximation. We expand on this work by proving a variational Bernstein–Von Mises theorem, leveraging the frequentist results to analyze VB posteriors.

The second theme is the Bernstein–Von Mises theorem. The classical (parametric) Bernstein–Von Mises theorem roughly says that the posterior distribution of $\sqrt{n}(\theta - \theta_0)$ “converges”, under the true parameter value θ_0 , to $\mathcal{N}(X, 1/I(\theta_0))$, where $X \sim \mathcal{N}(0, 1/I(\theta_0))$ and $I(\theta_0)$ is the Fisher information ([Ghosh and Ramamoorthi, 2003](#); [Van der Vaart, 2000](#); [Le Cam, 1953](#); [Le Cam and Yang, 2012](#)). Early forms of this theorem date back to Laplace, Bernstein, and Von Mises ([Laplace, 1809](#); [Bernstein, 1917](#); [Von Mises, 1931](#)). A version also appears in [Lehmann and Casella \(2006\)](#). [Kleijn et al. \(2012\)](#) established the Bernstein–Von Mises theorem under model misspecification. Recent advances include extensions to semiparametric and nonparametric cases ([Bickel et al., 2012](#); [Castillo et al., 2014](#)). In particular, [Lu et al. \(2016\)](#) proved a Bernstein–Von Mises type result for Bayesian inverse problems, characterizing Gaussian approximations of probability measures with respect to the KL divergence. Below, we borrow proof ideas from [Lu et al. \(2016\)](#). But we move beyond the Gaussian approximation to establish frequentist consistency of variational Bayes.

B This paper

The rest of the paper is organized as follows. Appendix D characterizes theoretical properties of the VB ideal. Appendix E contains the central results of the paper. It first connects the VB ideal and the VB posterior. It then proves the variational Bernstein–Von Mises theorem, which characterizes the asymptotic properties of the VB posterior and VB estimate. Appendix F studies three models under this theoretical lens, illustrating how to establish consistency and asymptotic normality of specific VB estimates. Finally, Appendix G reports simulation studies to illustrate these theoretical results.

C Glossary of terms

Please refer to Table 1.

D The VB ideal

To study the VB posterior $q^*(\theta)$, we first study the VB ideal of Equation (7). In the next section we connect it to the VB posterior.

Recall the VB ideal is

$$\pi^*(\theta | x) = \frac{p(\theta) \exp(M_n(\theta; x))}{\int p(\theta) \exp(M_n(\theta; x)) d\theta},$$

Name	Definition
Variational log likelihood	$M_n(x; \theta) := \sup_{q(z) \in \mathcal{Q}^n} \int q(z) \log \frac{p(z, x \theta)}{q(z)} dz$
Variational frequentist estimate (VFE)	$\hat{\theta}_n := \operatorname{argmax}_{\theta} M_n(x; \theta)$
VB ideal	$\pi^*(\theta x) := \frac{p(\theta) \exp\{M_n(x; \theta)\}}{\int p(\theta) \exp\{M_n(x; \theta)\} d\theta}$
Evidence Lower Bound (ELBO)	$\operatorname{ELBO}(q(\theta, z)) := \int \int q(\theta) q(z) \log \frac{p(z, x, \theta)}{q(\theta) q(z)} d\theta dz$
VB posterior	$q^*(\theta) := \operatorname{argmax}_{q(\theta) \in \mathcal{Q}^d} \sup_{q(z) \in \mathcal{Q}^n} \operatorname{ELBO}(q(\theta, z))$
VB estimate (VBE)	$\hat{\theta}_n^* := \int \theta \cdot q^*(\theta) d\theta$

Table 1: Glossary of terms

where $M_n(\theta; x)$ is the variational log likelihood of Equation (4). If we embed the variational log likelihood $M_n(\theta; x)$ in a statistical model of x , this model has likelihood

$$\ell(\theta; x) \propto \exp(M_n(\theta; x)).$$

We call it the *frequentist variational model*. The VB ideal $\pi^*(\theta | x)$ is thus the classical posterior under the frequentist variational model $\ell(\theta; x)$; the VFE is the classical maximum likelihood estimate (MLE).

Consider the results around frequentist estimation of θ under variational approximations of the local variables z (Bickel et al., 2013; Hall et al., 2011b; Westling and McCormick, 2015). These works consider asymptotic properties of estimators that maximize $M_n(\theta; x)$ with respect to θ . We will first leverage these results to prove properties of the VB ideal and their KL minimizers in the mean field variational family \mathcal{Q}^d . Then we will use these properties to study the VB posterior, which is what is estimated in practice.

This section relies on the consistent testability and the local asymptotic normality (LAN) of $M_n(\theta; x)$ (defined later) to show the VB ideal is consistent and asymptotically normal. We will then show that its KL minimizer in the mean field family is also consistent and converges to the KL minimizer of a normal distribution in TV distance.

These results are not surprising. Suppose the variational log likelihood behaves similarly to the true log likelihood, i.e., they produce consistent parameter estimates. Then, in the spirit of the classical Bernstein–Von Mises theorem under model misspecification (Kleijn et al., 2012), we expect the VB ideal to be consistent as well. Moreover, the approximation through a factorizable variational family should not ruin this consistency—point masses are factorizable and thus the limiting distribution lies in the approximating family.

D.1 The VB ideal

The lemma statements and proofs adapt ideas from Ghosh and Ramamoorthi (2003); Van der Vaart (2000); Bickel and Yahav (1967); Kleijn et al. (2012); Lu et al. (2016) to the variational log likelihood. Let Θ be an open subset of \mathbb{R}^d . Suppose the observations $x = x_{1:n}$ are a random sample from the measure P_{θ_0} with density $\int p(x, z | \theta = \theta_0) dz$ for some fixed, nonrandom value $\theta_0 \in \Theta$. $z = z_{1:n}$ are local latent variables, and $\theta = \theta_{1:d} \in \Theta$ are global latent variables. We assume that the density maps $(\theta, x) \mapsto \int p(x, z | \theta) dz$ of the true model and $(\theta, x) \mapsto \ell(\theta; x)$ of the variational frequentist models are measurable. For simplicity, we also assume that for each n there exists a single measure that dominates all measures with densities $\ell(\theta; x), \theta \in \Theta$ as well as the true measure P_{θ_0} .

Assumption 1. We assume the following conditions for the rest of the paper:

1. (Prior mass) The prior measure with Lebesgue-density $p(\theta)$ on Θ is continuous and positive on a neighborhood of θ_0 . There exists a constant $M_p > 0$ such that $|(\log p(\theta))''| \leq M_p e^{|\theta|^2}$.
2. (Consistent testability) For every $\epsilon > 0$ there exists a sequence of tests ϕ_n such that

$$\int \phi_n(x) p(x, z | \theta_0) dz dx \rightarrow 0$$

and

$$\sup_{\theta: \|\theta - \theta_0\| \geq \epsilon} \int (1 - \phi_n(x)) \frac{\ell(\theta; x)}{\ell(\theta_0; x)} p(x, z | \theta_0) dz dx \rightarrow 0,$$

3. (Local asymptotic normality (LAN)) For every compact set $K \subset \mathbb{R}^d$, there exist random vectors Δ_{n, θ_0} bounded in probability and nonsingular matrices V_{θ_0} such that

$$\sup_{h \in K} |M_n(\theta + \delta_n h; x) - M_n(\theta; x) - h^\top V_{\theta_0} \Delta_{n, \theta_0} + \frac{1}{2} h^\top V_{\theta_0} h| \xrightarrow{P_{\theta_0}} 0,$$

where δ_n is a $d \times d$ diagonal matrix. We have $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. For $d = 1$, we commonly have $\delta_n = 1/\sqrt{n}$.

These three assumptions are standard for Bernstein–Von Mises theorem. The first assumption is a prior mass assumption. It says the prior on θ puts enough mass to sufficiently small balls around θ_0 . This allows for optimal rates of convergence of the posterior. The first assumption further bounds the second derivative of the log prior density. This is a mild technical assumption satisfied by most non-heavy-tailed distributions.

The second assumption is a consistent testability assumption. It says there exists a sequence of uniformly consistent (under P_{θ_0}) tests for testing $H_0 : \theta = \theta_0$ against $H_1 : \|\theta - \theta_0\| \geq \epsilon$ for every $\epsilon > 0$ based on the frequentist variational model. This is a weak assumption. For example, it suffices to have a compact Θ and continuous and identifiable $M_n(\theta; x)$. It is also true when there exists a consistent estimator T_n of θ . In this case, we can set $\phi_n := 1\{T_n - \theta \geq \epsilon/2\}$.

The last assumption is a local asymptotic normality assumption on $M_n(\theta; x)$ around the true value θ_0 . It says the frequentist variational model can be asymptotically approximated by a normal location model centered at θ_0 after a rescaling of δ_n^{-1} . This normalizing sequence δ_n determines the optimal rates of convergence of the posterior. For example, if $\delta_n = 1/\sqrt{n}$, then we commonly have $\theta - \theta_0 = O_p(1/\sqrt{n})$.

In the spirit of the last assumption, we perform a change-of-variable step:

$$\tilde{\theta} = \delta_n^{-1}(\theta - \theta_0). \tag{8}$$

We center θ at the true value θ_0 and rescale it by the reciprocal of the rate of convergence δ_n^{-1} . This ensures that the asymptotic distribution of $\tilde{\theta}$ is not degenerate, i.e., it does not converge to a point mass. We define $\pi_{\tilde{\theta}}^*(\cdot | x)$ as the density of $\tilde{\theta}$ when θ has density $\pi^*(\cdot | x)$:

$$\pi_{\tilde{\theta}}^*(\tilde{\theta} | x) = \pi^*(\theta_0 + \delta_n \tilde{\theta} | x) \cdot |\det(\delta_n)|.$$

Now we characterize the asymptotic properties of the VB ideal.

Lemma 1. *The VB ideal converges in total variation to a sequence of normal distributions,*

$$\|\pi_{\tilde{\theta}}^*(\cdot | x) - \mathcal{N}(\cdot; \Delta_{n, \theta_0}, V_{\theta_0}^{-1})\|_{\text{TV}} \xrightarrow{P_{\theta_0}} 0.$$

Proof sketch of lemma 1. This is a consequence of the classical finite-dimensional Bernstein–Von Mises theorem under model misspecification (Kleijn et al., 2012). Theorem 2.1 of Kleijn et al. (2012) roughly says that the posterior is consistent if the model is locally asymptotically normal around the true parameter value θ_0 . Here the true data generating measure is P_{θ_0} with density $\int p(x, z | \theta = \theta_0) dz$, while the frequentist variational model has densities $\ell(\theta; x), \theta \in \Theta$.

What we need to show is that the consistent testability assumption in Assumption 1 implies assumption (2.3) in Kleijn et al. (2012):

$$\int_{|\tilde{\theta}| > M_n} \pi_{\tilde{\theta}}^*(\tilde{\theta} | x) d\tilde{\theta} \xrightarrow{P_{\theta_0}} 0$$

for every sequence of constants $M_n \rightarrow \infty$. To show this, we mimic the argument of Theorem 3.1 of Kleijn et al. (2012), where they show this implication for the iid case with a common convergence rate for all dimensions of θ . See Appendix A for details. \square

This lemma says the VB ideal of the rescaled θ , $\tilde{\theta} = \delta_n^{-1}(\theta - \theta_0)$, is asymptotically normal with mean Δ_{n,θ_0} . The mean, Δ_{n,θ_0} , as assumed in Assumption 1, is a random vector bounded in probability. The asymptotic distribution $\mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}^{-1})$ is thus also random, where randomness is due to the data x being random draws from the true data generating measure \mathbb{P}_{θ_0} . We notice that if the VFE, $\hat{\theta}_n$, is consistent and asymptotically normal, we commonly have $\Delta_{n,\theta_0} = \delta_n^{-1}(\hat{\theta}_n - \theta_0)$ with $\mathbb{E}(\Delta_{n,\theta_0}) = 0$. Hence, the VB ideal will converge to a normal distribution with a random mean centered at the true value θ_0 .

D.2 The KL minimizer of the VB ideal

Next we study the KL minimizer of the VB ideal in the mean field variational family. We show its consistency and asymptotic normality. To be clear, the asymptotic normality is in the sense that the KL minimizer of the VB ideal converges to the KL minimizer of a normal distribution in TV distance.

Lemma 2. *The KL minimizer of the VB ideal over the mean field family is consistent: almost surely under P_{θ_0} , it converges to a point mass,*

$$\arg \min_{q(\theta) \in \mathcal{Q}^d} \text{KL}(q(\theta) || \pi^*(\theta | x)) \xrightarrow{d} \delta_{\theta_0}.$$

Proof sketch of lemma 2. The key insight here is that point masses are factorizable. Lemma 1 above suggests that the VB ideal converges in distribution to a point mass. We thus have its KL minimizer also converging to a point mass, because point masses reside within the mean field family. In other words, there is no loss, in the limit, incurred by positing a factorizable variational family for approximation.

To prove this lemma, we bound the mass of $B^c(\theta_0, \eta_n)$ under $q(\theta)$, where $B^c(\theta_0, \eta_n)$ is the complement of an η_n -sized ball centered at θ_0 with $\eta_n \rightarrow 0$ as $n \rightarrow \infty$. In this step, we borrow ideas from the proof of Lemma 3.6 and Lemma 3.7 in Lu et al. (2016). See Appendix B for details. \square

Lemma 3. *The KL minimizer of the VB ideal of $\tilde{\theta}$ converges to that of $\mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}^{-1})$ in total variation: under mild technical conditions on the tail behavior of \mathcal{Q}^d (see Assumption 2 in Appendix C),*

$$\left\| \arg \min_{q \in \mathcal{Q}^d} \text{KL}(q(\cdot) || \pi_{\tilde{\theta}}^*(\cdot | x)) - \arg \min_{q \in \mathcal{Q}^d} \text{KL}(q(\cdot) || \mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}^{-1})) \right\|_{\text{TV}} \xrightarrow{P_{\theta_0}} 0.$$

Proof sketch of lemma 3. The intuition here is that, if the two distributions are close in the limit, their KL minimizers should also be close in the limit. Lemma 1 says that the VB ideal of $\tilde{\theta}$ converges to $\mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}^{-1})$ in total variation. We would expect their KL minimizer also converges in some metric. This result is also true for the (full-rank) Gaussian variational family if rescaled appropriately.

Here we show their convergence in total variation. This is achieved by showing the Γ -convergence of the functionals of q : $\text{KL}(q(\cdot) || \pi_{\tilde{\theta}}^*(\cdot | x))$ to $\text{KL}(q(\cdot) || \mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}^{-1}))$, for parametric q 's. Γ -convergence is a classical tool for characterizing variational problems; Γ -convergence of functionals ensures convergence of their minimizers (Dal Maso, 2012; Braides, 2006). See Appendix C for proof details and a review of Γ -convergence. \square

We characterized the limiting properties of the VB ideal and their KL minimizers. We will next show that the VB posterior is close to the KL divergence minimizer of the VB ideal. Appendix E culminates in the main theorem of this paper – the variational Bernstein–Von Mises theorem – showing the VB posterior share consistency and asymptotic normality with the KL divergence minimizer of VB ideal.

E Frequentist consistency of variational Bayes

We now study the VB posterior. In the previous section, we proved theoretical properties for the VB ideal and its KL minimizer in the variational family. Here we first connect the VB ideal to the VB posterior, the quantity that is used in practice. We then use this connection to understand the theoretical properties of the VB posterior.

We begin by characterizing the optimal variational distribution in a useful way. Decompose the variational family as

$$q(\theta, z) = q(\theta)q(z),$$

where $q(\theta) = \prod_{i=1}^d q(\theta_i)$ and $q(z) = \prod_{i=1}^n q(z_i)$. Denote the prior $p(\theta)$. Note d does not grow with the size of the data. We will develop a theory around VB that considers asymptotic properties of the VB posterior $q^*(\theta)$.

We decompose the ELBO of Equation (3) into the portion associated with the global variable and the portion associated with the local variables,

$$\begin{aligned} \text{ELBO}(q(\theta)q(z)) &= \int \int q(\theta)q(z) \log \frac{p(\theta, z, x)}{q(\theta)q(z)} d\theta dz \\ &= \int \int q(\theta)q(z) \log \frac{p(\theta)p(z, x | \theta)}{q(\theta)q(z)} d\theta dz \\ &= \int q(\theta) \log \frac{p(\theta)}{q(\theta)} d\theta + \int q(\theta) \int q(z) \log \frac{p(z, x | \theta)}{q(z)} d\theta dz. \end{aligned}$$

The optimal variational factor for the global variables, i.e., the VB posterior, maximizes the ELBO. From the decomposition, we can write it as a function of the optimized local variational factor,

$$q^*(\theta) = \arg \max_{q(\theta)} \sup_{q(z)} \int q(\theta) \left(\log \left[p(\theta) \exp \left\{ \int q(z) \log \frac{p(z, x | \theta)}{q(z)} dz \right\} \right] - \log q(\theta) \right) d\theta. \quad (9)$$

One way to see the objective for the VB posterior is as the ELBO profiled over $q(z)$, i.e., where the optimal $q(z)$ is a function of $q(\theta)$ (Hoffman et al., 2013). With this perspective, the ELBO becomes a function of $q(\theta)$ only. We denote it as a functional $\text{ELBO}_p(\cdot)$:

$$\text{ELBO}_p(q(\theta)) := \sup_{q(z)} \int q(\theta) \left(\log \left[p(\theta) \exp \left\{ \int q(z) \log \frac{p(z, x | \theta)}{q(z)} dz \right\} \right] - \log q(\theta) \right) d\theta. \quad (10)$$

We then rewrite Equation (9) as $q^*(\theta) = \arg \max_{q(\theta)} \text{ELBO}_p(q(\theta))$. This expression for the VB posterior is key to our results.

E.1 KL minimizers of the VB ideal

Recall that the KL minimization objective to the ideal VB posterior is the functional $\text{KL}(\cdot || \pi^*(\theta | x))$. We first show that the two optimization objectives $\text{KL}(\cdot || \pi^*(\theta | x))$ and $\text{ELBO}_p(\cdot)$ are close in the limit. Given the continuity of both $\text{KL}(\cdot || \pi^*(\theta | x))$ and $\text{ELBO}_p(\cdot)$, this implies the asymptotic properties of optimizers of $\text{KL}(\cdot || \pi^*(\theta | x))$ will be shared by the optimizers of $\text{ELBO}_p(\cdot)$.

Lemma 4. *The negative KL divergence to the VB ideal is equivalent to the profiled ELBO in the limit: under mild technical conditions on the tail behavior of \mathcal{Q}^d (see for example Assumption 3 in Appendix D), for $q(\theta) \in \mathcal{Q}^d$,*

$$\text{ELBO}_p(q(\theta)) = -\text{KL}(q(\theta) || \pi^*(\theta | x)) + o_p(1).$$

Proof sketch of Lemma 4. We first notice that

$$- \text{KL}(q(\theta) \parallel \pi^*(\theta \mid x)) \quad (11)$$

$$= \int q(\theta) \log \frac{p(\theta) \exp(M_n(\theta; x))}{q(\theta)} d\theta \quad (12)$$

$$= \int q(\theta) \left(\log \left[p(\theta) \exp \left\{ \sup_{q(z)} \int q(z) \log \frac{p(z, x \mid \theta)}{q(z)} dz \right\} \right] - \log q(\theta) \right) d\theta. \quad (13)$$

Comparing Equation (13) with Equation (10), we can see that the only difference between $-\text{KL}(\cdot \parallel \pi^*(\theta \mid x))$ and $\text{ELBO}_p(\cdot)$ is in the position of $\sup_{q(z)}$. $\text{ELBO}_p(\cdot)$ allows for a single choice of optimal $q(z)$ given $q(\theta)$, while $-\text{KL}(\cdot \parallel \pi^*(\theta \mid x))$ allows for a different optimal $q(z)$ for each value of θ . In this sense, if we restrict the variational family of $q(\theta)$ to be point masses, then $\text{ELBO}_p(\cdot)$ and $-\text{KL}(\cdot \parallel \pi^*(\theta \mid x))$ will be the same.

The only members of the variational family of $q(\theta)$ that admit finite limiting $-\text{KL}(q(\theta) \parallel \pi^*(\theta \mid x))$ are the ones that converge to point masses at rate δ_n , so we expect $\text{ELBO}_p(\cdot)$ and $-\text{KL}(\cdot \parallel \pi^*(\theta \mid x))$ to be close as $n \rightarrow \infty$. We prove this by bounding the remainder in the Taylor expansion of $M_n(\theta; x)$ by a sequence converging to zero in probability. See Appendix D for details. \square

E.2 The VB posterior

Appendix D characterizes the asymptotic behavior of the VB ideal $\pi^*(\theta \mid x)$ and their KL minimizers. Lemma 4 establishes the connection between the VB posterior $q^*(\theta)$ and the KL minimizers of the VB ideal $\pi^*(\theta \mid x)$. Recall $\arg \min_{q(\theta) \in \mathcal{Q}^d} \text{KL}(q(\theta) \parallel \pi^*(\theta \mid x))$ is consistent and converges to the KL minimizer of a normal distribution. We now build on these results to study the VB posterior $q^*(\theta)$.

Now we are ready to state the main theorem. It establishes the asymptotic behavior of the VB posterior $q^*(\theta)$.

Theorem 5. (Variational Bernstein-von-Mises Theorem)

1. The VB posterior is consistent: almost surely under P_{θ_0} ,

$$q^*(\theta) \xrightarrow{d} \delta_{\theta_0}.$$

2. The VB posterior is asymptotically normal in the sense that it converges to the KL minimizer of a normal distribution:

$$\left\| q_{\tilde{\theta}}^*(\cdot) - \arg \min_{q \in \mathcal{Q}^d} \text{KL}(q(\cdot) \parallel \mathcal{N}(\cdot; \Delta_{n, \theta_0}, \mathbf{V}_{\theta_0}^{-1})) \right\|_{\text{TV}} \xrightarrow{P_{\theta_0}} 0. \quad (14)$$

Here we transform $q^*(\theta)$ to $q_{\tilde{\theta}}^*(\tilde{\theta})$, which is centered around the true θ_0 and scaled by the convergence rate; see Equation (8).

Proof sketch of Theorem 5. This theorem is a direct consequence of Lemma 2, Lemma 3, Lemma 4. We need the same mild technical conditions on \mathcal{Q}^d as in Lemma 3 and Lemma 4. See Appendix E for details. \square

Given the convergence of the VB posterior, we can now establish the asymptotic properties of the VBE.

Theorem 6. (Asymptotics of the VBE)

Assume $\int |\theta|^2 \pi(\theta) d\theta < \infty$. Let $\hat{\theta}_n^* = \int \theta \cdot q_1^*(\theta) d\theta$ denote the VBE.

1. The VBE is consistent: under \mathbb{P}_{θ_0} ,

$$\hat{\theta}_n^* \xrightarrow{a.s.} \theta_0.$$

2. The VBE is asymptotically normal in the sense that it converges in distribution to the mean of the KL minimizer:³ if $\Delta_{n,\theta_0} \xrightarrow{d} X$ for some X ,

$$\delta_n^{-1}(\hat{\theta}_n^* - \theta_0) \xrightarrow{d} \int \tilde{\theta} \cdot \arg \min_{q \in \mathcal{Q}^d} \text{KL}(q(\tilde{\theta}) || \mathcal{N}(\tilde{\theta}; X, V_{\theta_0}^{-1})) d\tilde{\theta}.$$

Proof sketch of Theorem 6. As the posterior mean is a continuous function of the posterior distribution, we would expect the VBE is consistent given the VB posterior is. We also know that the posterior mean is the Bayes estimator under squared loss. Thus we would expect the VBE to converge in distribution to squared loss minimizer of the KL minimizer of the VB ideal. The result follows from a very similar argument from Theorem 2.3 of Kleijn et al. (2012), which shows that the posterior mean estimate is consistent and asymptotically normal under model misspecification as a consequence of the Bernstein–Von Mises theorem and the argmax theorem. See Appendix E for details. \square

We remark that Δ_{n,θ_0} , as in Assumption 1, is a random vector bounded in P_{θ_0} probability. The randomness is due to x being a random sample generated from P_{θ_0} .

In cases where VFE is consistent, like in all the examples we will see in Appendix F, Δ_{n,θ_0} is a zero mean random vector with finite variance. For particular realizations of x the value of Δ_{n,θ_0} might not be zero; however, because we scale by δ_n^{-1} , this does not destroy the consistency of VB posterior or the VBE.

E.3 Gaussian VB posteriors

We illustrate the implications of Theorem 5 and Theorem 6 on two choices of variational families: a full rank Gaussian variational family and a factorizable Gaussian variational family. In both cases, the VB posterior and the VBE are consistent and asymptotically normal with different covariance matrices. The VB posterior under the factorizable family is underdispersed.

Corollary 7. *Posit a full rank Gaussian variational family, that is*

$$\mathcal{Q}^d = \{q : q(\theta) = \mathcal{N}(m, \Sigma)\}, \quad (15)$$

with Σ positive definite. Then

1. $q^*(\theta) \xrightarrow{d} \delta_{\theta_0}$, almost surely under \mathbb{P}_{θ_0} .
2. $\|q_{\tilde{\theta}}^*(\cdot) - \mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}^{-1})\|_{\text{TV}} \xrightarrow{P_{\theta_0}} 0$.
3. $\hat{\theta}_n^* \xrightarrow{a.s.} \theta_0$.
4. $\delta_n^{-1}(\hat{\theta}_n^* - \theta_0) - \Delta_{n,\theta_0} = o_{P_{\theta_0}}(1)$.

Proof sketch of corollary 7. This is a direct consequence of Theorem 5 and Theorem 6. We only need to show that Lemma 3 is also true for the full rank Gaussian variational family. The last conclusion implies $\delta_n^{-1}(\hat{\theta}_n^* - \theta_0) \xrightarrow{d} X$ if $\Delta_{n,\theta_0} \xrightarrow{d} X$ for some random variable X . We defer the proof to Appendix F. \square

This corollary says that under a full rank Gaussian variational family, VB is consistent and asymptotically normal in the classical sense. It accurately recovers the asymptotic normal distribution implied by the local asymptotic normality of $M_n(\theta; x)$.

Before stating the corollary for the factorizable Gaussian variational family, we first present a lemma on the KL minimizer of a Gaussian distribution over the factorizable Gaussian family. We show

³The randomness in the mean of the KL minimizer comes from Δ_{n,θ_0} .

that the minimizer keeps the mean but has a diagonal covariance matrix that matches the precision. We also show the minimizer has a smaller entropy than the original distribution. This echoes the well-known phenomenon of VB algorithms underestimating the variance.

Lemma 8. *The factorizable KL minimizer of a Gaussian distribution keeps the mean and matches the precision:*

$$\operatorname{argmin}_{\mu_0 \in \mathbb{R}^d, \Sigma_0 \in \operatorname{diag}(d \times d)} \operatorname{KL}(\mathcal{N}(\cdot; \mu_0, \Sigma_0) \parallel \mathcal{N}(\cdot; \mu_1, \Sigma_1)) = \mu_1, \Sigma_1^*,$$

where Σ_1^* is diagonal with $\Sigma_{1,ii}^* = ((\Sigma_1^{-1})_{ii})^{-1}$ for $i = 1, 2, \dots, d$. Hence, the entropy of the factorizable KL minimizer is smaller than or equal to that of the original distribution:

$$\mathbb{H}(\mathcal{N}(\cdot; \mu_0, \Sigma_1^*)) \leq \mathbb{H}(\mathcal{N}(\cdot; \mu_0, \Sigma_1)).$$

Proof sketch of Lemma 8. The first statement is consequence of a technical calculation of the KL divergence between two normal distributions. We differentiate the KL divergence over μ_0 and the diagonal terms of Σ_0 and obtain the result. The second statement is due to the inequality of the determinant of a positive matrix being always smaller than or equal to the product of its diagonal terms (Amir-Moez and Johnston, 1969; Beckenbach and Bellman, 2012). In this sense, mean field variational inference underestimates posterior variance. See Appendix G for details. \square

The next corollary studies the VB posterior and the VBE under a factorizable Gaussian variational family.

Corollary 9. *Posit a factorizable Gaussian variational family,*

$$\mathcal{Q}^d = \{q : q(\theta) = \mathcal{N}(m, \Sigma,)\} \tag{16}$$

where Σ positive definite and diagonal. Then

1. $q^*(\theta) \xrightarrow{d} \delta_{\theta_0}$, almost surely under \mathbb{P}_{θ_0} .
2. $\|q_{\hat{\theta}}^*(\cdot) - \mathcal{N}(\cdot; \Delta_{n,\theta_0}, V_{\theta_0}'^{-1})\|_{\text{TV}} \xrightarrow{P_{\theta_0}} 0$,
where V' is diagonal and has the same diagonal entries as V_{θ_0} .
3. $\hat{\theta}_n^* \xrightarrow{a.s.} \theta_0$.
4. $\delta_n^{-1}(\hat{\theta}_n^* - \theta_0) - \Delta_{n,\theta_0} = o_{P_{\theta_0}}(1)$.

Proof of corollary 9. This is a direct consequence of Lemma 8, Theorem 5, and Theorem 6. \square

This corollary says that under the factorizable Gaussian variational family, VB is consistent and asymptotically normal in the classical sense. The rescaled asymptotic distribution for $\hat{\theta}$ recovers the mean but underestimates the covariance. This underdispersion is a common phenomenon we see in mean field variational Bayes.

F Applications

We proved consistency and asymptotic normality of the variational Bayes (VB) posterior (in total variation (TV) distance) and the variational Bayes estimate (VBE). We mainly relied on the prior mass condition, the local asymptotic normality of the variational log likelihood $M_n(x; \theta)$ and the consistent testability assumption of the data generating parameter.

We now apply this argument to three types of Bayesian models: Bayesian mixture models (Bishop, 2006; Murphy, 2012), Bayesian generalized linear mixed models (McCulloch and Neuhaus, 2001; Jiang, 2007), and Bayesian stochastic block models (Wang and Wong, 1987; Snijders and Nowicki, 1997; Mossel et al., 2012; Abbe and Sandon, 2015; Hofman and Wiggins, 2008). For each model class,

we illustrate how to leverage the known asymptotic results for frequentist variational approximations to prove asymptotic results for VB. We assume the prior mass condition for the rest of this section: the prior measure of a parameter θ with Lebesgue density $p(\theta)$ on Θ is continuous and positive on a neighborhood of the true data generating value θ_0 . For simplicity, we posit a mean field family for the local latent variables and a factorizable Gaussian variational family for the global latent variables.

F.1 Bayesian Mixture models

The Bayesian mixture model is a versatile class of models for density estimation and clustering (Bishop, 2006; Murphy, 2012).

Consider a Bayesian mixture of K unit-variance univariate Gaussians with means $\mu = \{\mu_1, \dots, \mu_K\}$. For each observation $x_i, i = 1, \dots, n$, we first randomly draw a cluster assignment c_i from a categorical distribution over $\{1, \dots, K\}$; we then draw x_i randomly from a unit-variance Gaussian with mean μ_{c_i} . The model is

$$\begin{aligned} \mu_k &\sim p_\mu, & k &= 1, \dots, K, \\ c_i &\sim \text{Categorical}(1/K, \dots, 1/K), & i &= 1, \dots, n, \\ x_i | c_i, \mu &\sim \mathcal{N}(c_i^\top \mu, 1), & i &= 1, \dots, n. \end{aligned}$$

For a sample of size n , the joint distribution is

$$p(\mu, c, x) = \prod_{i=1}^K p_\mu(\mu_i) \prod_{i=1}^n p(c_i) p(x_i | c_i, \mu).$$

Here μ is a K -dimensional global latent vector and $c_{1:n}$ are local latent variables. We are interested in inferring the posterior of the μ vector.

We now establish asymptotic properties of VB for Bayesian Gaussian mixture model (GMM).

Corollary 10. *Assume the data generating measure P_{μ_0} has density $\int p(\mu_0, c, x) dc$. Let $q^*(\mu)$ and μ^* denote the VB posterior and the VBE. Under regularity conditions (A1-A5) and (B1,2,4) of Westling and McCormick (2015), we have*

$$\left\| q^*(\mu) - \mathcal{N}\left(\mu_0 + \frac{Y}{\sqrt{n}}, \frac{1}{n} V_0(\mu_0)\right) \right\|_{\text{TV}} \xrightarrow{P_{\mu_0}} 0,$$

and

$$\sqrt{n}(\mu^* - \mu_0) \xrightarrow{d} \mathcal{N}(Y, V_0(\mu_0)),$$

where μ_0 is the true value of μ that generates the data. We have

$$Y \sim \mathcal{N}(0, V(\mu_0)),$$

$$V(\mu_0) = A(\mu_0)^{-1} B(\mu_0) A(\mu_0)^{-1},$$

$$A(\mu) = \mathbb{E}_{P_{\mu_0}} [D_\mu^2 m(\mu; x)],$$

$$B(\mu) = \mathbb{E}_{P_{\mu_0}} [D_\mu m(\mu; x) D_\mu m(\mu; x)^\top],$$

$$m(\mu; x) = \sup_{q(c) \in \mathcal{Q}^n} \int q(c) \log \frac{p(x, c | \mu)}{q(c)} dc.$$

The diagonal matrix $V_0(\mu_0)$ satisfies $(V_0(\mu_0)^{-1})_{ii} = (A(\mu_0))_{ii}$. The specification of Gaussian mixture model is invariant to permutation among K components; this corollary is true up to permutations among the K components.

Proof sketch for Corollary 10. The consistent testability condition is satisfied by the existence of a consistent estimate due to Theorem 1 of Westling and McCormick (2015). The local asymptotic normality is proved by a Taylor expansion of $m(\mu; x)$ at μ_0 . This result then follows directly from our Theorem 5 and Theorem 6 in Appendix E. The technical conditions inherited from Westling and McCormick (2015) allow us to use their Theorems 1 and 2 for properties around variational frequentist estimate (VFE). See Appendix H for proof details. \square

F.2 Bayesian Generalized linear mixed models

Bayesian generalized linear mixed models (GLMMS) are a powerful class of models for analyzing grouped data or longitudinal data (McCulloch and Neuhaus, 2001; Jiang, 2007).

Consider a Poisson mixed model with a simple linear relationship and group-specific random intercepts. Each observation reads (X_{ij}, Y_{ij}) , $1 \leq i \leq m, 1 \leq j \leq n$, where the Y_{ij} 's are non-negative integers and the X_{ij} 's are unrestricted real numbers. For each group of observations (X_{ij}, Y_{ij}) , $1 \leq j \leq n$, we first draw the random effect U_i independently from $\mathcal{N}(0, \sigma^2)$. We follow by drawing Y_{ij} from a Poisson distribution with mean $\exp(\beta_0 + \beta_1 X_{ij} + U_i)$. The probability model is

$$\begin{aligned}\beta_0 &\sim p_{\beta_0}, \\ \beta_1 &\sim p_{\beta_1}, \\ \sigma^2 &\sim p_{\sigma^2}, \\ U_i &\stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2), \\ Y_{ij} \mid X_{ij}, U_i &\sim \text{Poi}(\exp(\beta_0 + \beta_1 X_{ij} + U_i)).\end{aligned}$$

The joint distribution is

$$\begin{aligned}p(\beta_0, \beta_1, \sigma^2, U_{1:m}, Y_{1:m, 1:n} \mid X_{1:m, 1:n}) = \\ p_{\beta_0}(\beta_0) p_{\beta_1}(\beta_1) p_{\sigma^2}(\sigma^2) \prod_{i=1}^m \mathcal{N}(U_i; 0, \sigma^2) \times \prod_{i=1}^m \prod_{j=1}^n \text{Poi}(Y_{ij}; \exp(\beta_0 + \beta_1 X_{ij} + U_i)).\end{aligned}$$

We establish asymptotic properties of VB in Bayesian Poisson linear mixed models.

Corollary 11. Consider the true data generating distribution $P_{\beta_0^0, \beta_1^0, (\sigma^2)^0}$ with the global latent variables taking the true values $\{\beta_0^0, \beta_1^0, (\sigma^2)^0\}$. Let $q_{\beta_0}^*(\beta_0)$, $q_{\beta_1}^*(\beta_1)$, $q_{\sigma^2}^*(\sigma^2)$ denote the VB posterior of $\beta_0, \beta_1, \sigma^2$. Similarly, let $\beta_0^*, \beta_1^*, (\sigma^2)^*$ be the VBEs accordingly. Consider $m = O(n^2)$. Under regularity conditions (A1-A5) of Hall et al. (2011b), we have

$$\left\| q_{\beta_0}^*(\beta_0) q_{\beta_1}^*(\beta_1) q_{\sigma^2}^*(\sigma^2) - \mathcal{N}\left((\beta_0^0, \beta_1^0, (\sigma^2)^0) + \left(\frac{Z_1}{\sqrt{n}}, \frac{Z_2}{\sqrt{mn}}, \frac{Z_3}{\sqrt{n}}\right), \text{diag}(V_1, V_2, V_3)\right) \right\|_{\text{TV}} \stackrel{P_{\beta_0^0, \beta_1^0, (\sigma^2)^0}}{\rightarrow} 0,$$

where

$$\begin{aligned}Z_1 &\sim \mathcal{N}(0, (\sigma^2)^0), Z_2 \sim \mathcal{N}(0, \tau^2), Z_3 \sim \mathcal{N}(0, 2\{(\sigma^2)^0\}^2), \\ V_1 &= \exp(-\beta_0 + \frac{1}{2}(\sigma^2)^0) / \phi(\beta_1^0), \\ V_2 &= \exp(-\beta_0 + \frac{1}{2}\sigma^2) / \phi''(\beta_1^0), \\ V_3 &= 2\{(\sigma^2)^0\}^2, \\ \tau^2 &= \frac{\exp\{-(\sigma^2)^0/2 - \beta_0^0\} \phi(\beta_1^0)}{\phi''(\beta_1^0) \phi(\beta_1^0) - \phi'(\beta_1^0)^2}.\end{aligned}$$

Here $\phi(\cdot)$ is the moment generating function of X .

Also,

$$(\sqrt{m}(\beta_0^* - \beta_0^0), \sqrt{mn}(\beta_1^* - \beta_1^0), \sqrt{m}((\sigma^2)^* - (\sigma^2)^0)) \stackrel{d}{\rightarrow} \mathcal{N}((Z_1, Z_2, Z_3), \text{diag}(V_1, V_2, V_3)).$$

Proof sketch for Corollary 11. The consistent testability assumption is satisfied by the existence of consistent estimates of the global latent variables shown in Theorem 3.1 of Hall et al. (2011b). The local asymptotic normality is proved by a Taylor expansion of the variational log likelihood based on estimates of the variational parameters based on equations (5.18) and (5.22) of Hall et al. (2011b). The technical conditions inherited from Hall et al. (2011b) allow us to leverage their Theorem 3.1 for properties of the VFE. The result then follows directly from Theorem 5 and Theorem 6 in Appendix E. See Appendix I for proof details. \square

F.3 Bayesian stochastic block models

Stochastic block models are an important methodology for community detection in network data (Wang and Wong, 1987; Snijders and Nowicki, 1997; Mossel et al., 2012; Abbe and Sandon, 2015).

Consider n vertices in a graph. We observe pairwise linkage between nodes $A_{ij} \in \{0, 1\}$, $1 \leq i, j \leq n$. In a stochastic block model, this adjacency matrix is driven by the following process: first assign each node i to one of the K latent classes by a categorical distribution with parameter π . Denote the class membership as $Z_i \in \{1, \dots, K\}$. Then draw $A_{ij} \sim \text{Bernoulli}(H_{Z_i, Z_j})$. The parameter H is a symmetric matrix in $[0, 1]^{K \times K}$ that specifies the edge probabilities between two latent classes; the parameter π are the proportions of the latent classes. The Bayesian stochastic block model is

$$\begin{aligned} \pi &\sim p(\pi), \\ H &\sim p(H), \\ Z_i | \pi &\stackrel{iid}{\sim} \text{Categorical}(\pi), \\ A_{ij} | Z_i, Z_j, H &\stackrel{iid}{\sim} \text{Bernoulli}(H_{Z_i, Z_j}). \end{aligned}$$

The dependence in stochastic block model is more complicated than the Bayesian GMM or the Bayesian GLMM.

Before establishing the result, we reparameterize (π, H) by $\theta = (\omega, \nu)$, where $\omega \in \mathbb{R}^{K-1}$ is the log odds ratio of belonging to classes $1, \dots, K-1$, and $\nu \in \mathbb{R}^{K \times K}$ is the log odds ratio of an edge existing between all pairs of the K classes. The reparameterization is

$$\begin{aligned} \omega(a) &= \log \frac{\pi(a)}{1 - \sum_{b=1}^{K-1} \pi(b)}, & a = 1, \dots, K-1 \\ \nu(a, b) &= \log \frac{H(a, b)}{1 - H(a, b)}, & a, b = 1, \dots, K. \end{aligned}$$

The joint distribution is

$$p(\theta, Z, A) = \prod_{a=1}^{K-1} [e^{\omega(a)n_a} (1 + \sum_{a=1}^{K-1} e^{\omega(a)})^{-n}] \times \prod_{a=1}^K \prod_{b=1}^K [e^{\nu(a, b)O_{ab}} (1 + e^{\nu(a, b)})^{n_{ab}}]^{1/2},$$

where

$$\begin{aligned} n_a(Z) &= \sum_{i=1}^n \mathbf{1}\{Z_i = a\}, \\ n_{ab}(Z) &= \sum_{i=1}^n \sum_{j \neq i}^n \mathbf{1}\{Z_i = a, Z_j = b\}, \\ O_{ab}(A, Z) &= \sum_{i=1}^n \sum_{j \neq i}^n \mathbf{1}\{Z_i = a, Z_j = b\} A_{ij}. \end{aligned}$$

We now establish the asymptotic properties of VB for stochastic block models.

Corollary 12. Consider ν_0, ω_0 as true data generating parameters. Let $q_\nu^*(\nu), q_\omega^*(\omega)$ denote the VB posterior of ν and ω . Similarly, let ν^*, ω^* be the VBE. Then

$$\left\| q_\nu^*(\nu) q_\omega^*(\omega) - \mathcal{N} \left((\nu, \omega); (\nu_0, \omega_0) + \left(\frac{\Sigma_1^{-1} Y_1}{\sqrt{n} \lambda_0}, \frac{\Sigma_2^{-1} Y_2}{\sqrt{n}} \right), V_n(\nu_0, \omega_0) \right) \right\|_{\text{TV}} \xrightarrow{P_{\nu_0, \omega_0}} 0$$

where $\lambda_0 = \mathbb{E}_{P_{\nu_0, \omega_0}}(\text{degree of each node})$, $(\log n)^{-1} \lambda_0 \rightarrow \infty$. Y_1 and Y_2 are two zero mean random vectors with covariance matrices Σ_1 and Σ_2 , where Σ_1, Σ_2 are known functions of ν_0, ω_0 . The diagonal matrix $V(\nu_0, \omega_0)$ satisfies $V^{-1}(\nu_0, \omega_0)_{ii} = \text{diag}(\Sigma_1, \Sigma_2)_{ii}$. Also,

$$(\sqrt{n} \lambda_0 (\nu^* - \nu_0), \sqrt{n} (\omega^* - \omega_0)) \xrightarrow{d} \mathcal{N}((\Sigma_1^{-1} Y_1, \Sigma_2^{-1} Y_2), V(\nu_0, \omega_0)),$$

The specification of classes in stochastic block model (SBM) is permutation invariant. So the convergence above is true up to permutation with the K classes. We follow [Bickel et al. \(2013\)](#) to consider the quotient space of (v, ω) over permutations.

Proof of Corollary 12. The consistent testability assumption is satisfied by the existence of consistent estimates by Lemma 1 of [Bickel et al. \(2013\)](#). The local asymptotic normality,

$$\begin{aligned} & \sup_{q(z) \in \mathcal{Q}^K} \int q(z) \log \frac{p(A, z | v_0 + \frac{t}{\sqrt{n^2 \rho_n}}, \omega + \frac{s}{\sqrt{n}})}{q(z)} dz \\ &= \sup_{q(z) \in \mathcal{Q}^K} \int q(z) \log \frac{p(A, z | v_0, \omega_0)}{q(z)} dz + s^\top Y_1 + t^\top Y_2 - \frac{1}{2} s^\top \Sigma_1 s - \frac{1}{2} t^\top \Sigma_2 t + o_P(1), \quad (17) \end{aligned}$$

for $(v_0, \omega_0) \in \mathcal{T}$ for compact \mathcal{T} with $\rho_n = \frac{1}{n} \mathbb{E}(\text{degree of each node})$, is established by Lemma 2 of [Bickel et al. \(2013\)](#). The result then follows directly from our Theorem 5 and Theorem 6 in Appendix E. \square

G Simulation studies

We illustrate the implication of Theorem 5 and Theorem 6 by simulation studies on Bayesian GLMM ([McCullagh, 1984](#)). We also study the VB posteriors of latent Dirichlet allocation (LDA) ([Blei et al., 2003](#)). This is a model that shares similar structural properties with SBM but has no consistency results established for its VFE.

We use two automated inference algorithms offered in Stan, a probabilistic programming system ([Carpenter et al., 2015](#)): VB through automatic differentiation variational inference (ADVI) ([Kucukelbir et al., 2016](#)) and Hamiltonian Monte Carlo (HMC) simulation through No-U-Turn sampler (NUTS) ([Hoffman and Gelman, 2014](#)). We note that optimization algorithms used for VB in practice only find local optima.


In both cases, we observe the VB posteriors get closer to the truth as the sample size increases; when the sample size is large enough, they coincide with the truth. They are underdispersed, however, compared with HMC methods.

G.1 Bayesian Generalized Linear Mixed Models


We consider the Poisson linear mixed model studied in Appendix F. Fix the group size as $n = 10$. We simulate data sets of size $N = (50, 100, 200, 500, 1000, 2000, 5000, 10000, 20000)$. As the size of the data set grows, the number of groups also grows; so does the number of local latent variables $U_i, 1 \leq i \leq m$. We generate a four-dimensional covariate vector for each $X_{ij}, 1 \leq i \leq m, 1 \leq j \leq n$, where the first dimension follows i.i.d $\mathcal{N}(0, 1)$, the second dimension follows i.i.d $\mathcal{N}(0, 25)$, the third dimension follows i.i.d Bernoulli(0.4), and the fourth dimension follows i.i.d. Bernoulli(0.8). We wish to study the behaviors of coefficient efficient for underdispersed/overdispersed continuous covariates and balanced/imbalanced binary covariates. We set the true parameters as $\beta_0 = 5$, $\beta_1 = (0.2, -0.2, 2, -2)$, and $\sigma^2 = 2$.

Figure 1 shows the boxplots of VB posteriors for β_0, β_1 , and σ^2 . All VB posteriors converge to their corresponding true values as the size of the data set increases. The box plots present rather few outliers; the lower fence, the box, and the upper fence are about the same size. This suggests normal VB posteriors. This echoes the consistency and asymptotic normality concluded from Theorem 5. The VB posteriors are underdispersed, compared to the posteriors via HMC. This also echoes our conclusion of underdispersion in Theorem 5 and Lemma 8.


Regarding the convergence rate, VB posteriors of all dimensions of β_1 quickly converge to their true value; the VB posteriors center around their true values as long as $N \geq 1000$. The convergence of VB posteriors of slopes for continuous variables (β_{11}, β_{12}) are generally faster than those for binary ones (β_{13}, β_{14}). The VB posterior of σ^2 shares a similarly fast convergence rate. The VB posterior of the intercept β_0 , however, struggles; it is away from the true value until the data set size




(a) Posterior of β_0




(b) Posterior of σ^2




(c) Posterior of β_{11}



(d) Posterior of β_{12}



(e) Posterior of β_{13}



(f) Posterior of β_{14}

Figure 1: VB posteriors and HMC posteriors of Poisson Generalized Linear Mixed Model versus size of datasets. VB posteriors are consistent and asymptotically normal but underdispersed than HMC posteriors. β_0 and σ^2 converge to the truth slower than β_1 does. They echo our conclusions in Theorem 5 and Corollary 11.

hits $N = 20000$. This aligns with the convergence rate inferred in Corollary 11, \sqrt{mn} for β_1 and \sqrt{m} for β_0 and σ^2 .

Computation wise, VB takes orders of magnitude less time than HMC. The performance of VB posteriors is comparable with that from HMC when the sample size is sufficiently large; in this case, we need $N = 20000$.

G.2 Latent Dirichlet Allocation

Latent Dirichlet Allocation (LDA) is a generative statistical model commonly adopted to describe word distributions in documents by latent topics.

Given M documents, each with $N_m, m = 1, \dots, M$ words, composing a vocabulary of V words, we assume K latent topics. Consider two sets of latent variables: topic distributions for document m , $(\theta_m)_{K \times 1}, m = 1, \dots, M$ and word distributions for topic k , $(\phi_k)_{V \times 1}, k = 1, \dots, K$. The generative process is

$$\begin{aligned} \theta_m &\sim p_\theta, & m &= 1, \dots, M, \\ \phi_k &\sim p_\phi, & k &= 1, \dots, K, \\ z_{m,j} &\sim \text{Mult}(\theta_m), & j &= 1, \dots, N_m, m = 1, \dots, M, \\ w_{m,j} &\sim \text{Mult}(\phi_{z_{m,j}}), & j &= 1, \dots, N_m, m = 1, \dots, M. \end{aligned}$$

The first two rows are assigning priors assigned to the latent variables. $w_{m,j}$ denotes word j of document m and $z_{m,j}$ denotes its assigned topic.

We simulate a data set with $V = 100$ sized vocabulary and $K = 10$ latent topics in $M = (10, 20, 50, 100, 200, 500, 1000)$ documents. Each document has N_m words where $N_m \stackrel{iid}{\sim} \text{Poi}(100)$. As the number of documents M grows, the number of document-specific topic vectors θ_m grows while the number of topic-specific word vectors ϕ_k stays the same. In this sense, we consider $\theta_m, m = 1, \dots, M$ as local latent variables and $\phi_k, k = 1, \dots, K$ as global latent variables. We are interested in the VB posteriors of global latent variables $\phi_k, k = 1, \dots, K$ here. We generate the data sets with true values of θ and ϕ , where they are random draws from $\theta_m \stackrel{iid}{\sim} \text{Dir}((1/K)_{K \times 1})$ and $\phi_k \stackrel{iid}{\sim} \text{Dir}((1/V)_{V \times 1})$.

Figure 2 presents the Kullback-Leibler (KL) divergence between the $K = 10$ topic-specific word distributions induced by the true ϕ_k 's and the fitted $\hat{\phi}_k$'s by VB and HMC. This KL divergence equals to $\text{KL}(\text{Mult}(\phi_k^0) \parallel \text{Mult}(\hat{\phi}_k)) = \sum_{i=1}^V \phi_{ki}^0 (\log \phi_{ki}^0 - \log \hat{\phi}_{ki})$, where ϕ_{ki}^0 is the i th entry of the true k th topic and $\hat{\phi}_{ki}$ is the i th entry of the fitted k th topic.

Figure 2a shows that VB posterior (dark blue) mean KL divergences of all $K = 10$ topics get closer to 0 as the number of documents M increase, faster than HMC (light blue). We become very close to the truth as the number of documents M hits 1000. Figure 2b⁴ shows that the boxplots of VB posterior mean KL divergences get closer to 0 as M increases. They are underdispersed compared to HMC posteriors. These align with our understanding of how VB posterior behaves in Theorem 5.

Computation wise, again VB is orders of magnitude faster than HMC. In particular, optimization in VB in our simulation studies converges within 10,000 steps.


H Discussion

Variational Bayes (VB) methods are a fast alternative to Markov chain Monte Carlo (MCMC) for posterior inference in Bayesian modeling. However, few theoretical guarantees have been established. This work proves consistency and asymptotic normality for variational Bayes (VB) posteriors. The convergence is in the sense of total variation (TV) distance converging to zero in probability. In addition, we establish consistency and asymptotic normality of variational Bayes estimate (VBE). The result is frequentist in the sense that we assume a data generating distribution driven by some fixed nonrandom true value for global latent variables.


These results rest on ideal variational Bayes and its connection to frequentist variational approximations. Thus this work bridges the gap in asymptotic theory between the frequentist variational approximation, in particular the variational frequentist estimate (VFE), and variational Bayes. It also assures us that variational Bayes as a popular approximate inference algorithm bears some theoretical soundness.

We present our results in the classical VB framework but the results and proof techniques are more generally applicable. Our results can be easily generalized to more recent developments of VB beyond Kullback-Leibler (KL) divergence, f -divergence or α -divergence for example. We could

⁴We only show boxplots for Topic 2 here. The boxplots of other topics look very similar.



(a) Posterior mean KL divergence of the $K = 10$ topics



(b) Boxplots of posterior KL divergence of Topic 2

Figure 2: Mean of KL divergence between the true topics and the fitted VB and HMC posterior topics versus size of datasets. (a) VB posteriors (dark blue) converge to the truth; they are very close to the truth as we hit $M = 1000$ documents. (b) VB posteriors are consistent but underdispersed compared to HMC posteriors (light blue). These align with our conclusions in Theorem 5.

also allow for model misspecification, as long as the variational loglikelihood $M_n(\theta; x)$ under the misspecified model still enjoys local asymptotic normality.

There are several interesting avenues for future work. The variational Bernstein–Von Mises theorem developed in this work is parametric; its parameters are of finite dimension. One direction is to develop a semiparametric or nonparametric counterpart. A second direction is to characterize the posterior contraction rates of VB posteriors. Finally, we characterized the asymptotics of an optimization problem, assuming that we obtain the global optimum. Though our simulations corroborated the theory, VB optimization typically finds a local optimum. Theoretically characterizing these local optima requires further study of the loss surface.

References

- Abbe, E. and Sandon, C. (2015). Community detection in general stochastic block models: fundamental limits and efficient recovery algorithms. *arXiv preprint arXiv:1503.00609*.
- Amir-Moez, A. and Johnston, G. (1969). On the product of diagonal elements of a positive matrix. *Mathematics Magazine*, 42(1):24–26.
- Beckenbach, E. F. and Bellman, R. (2012). *Inequalities*, volume 30. Springer Science & Business Media.
- Bernstein, S. N. (1917). *Theory of Probability*.
- Bickel, P., Choi, D., Chang, X., Zhang, H., et al. (2013). Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *The Annals of Statistics*, 41(4):1922–1943.
- Bickel, P., Kleijn, B., et al. (2012). The semiparametric Bernstein–von Mises theorem. *The Annals of Statistics*, 40(1):206–237.
- Bickel, P. J. and Yahav, J. A. (1967). Asymptotically pointwise optimal procedures in sequential analysis.
- Bishop, C. M. (2006). Pattern recognition. *Machine Learning*, 128.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Braides, A. (2006). A handbook of Γ -convergence. *Handbook of Differential Equations: stationary partial differential equations*, 3:101–213.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., and Riddell, A. (2015). Stan: a probabilistic programming language. *Journal of Statistical Software*.
- Castillo, I., Nickl, R., et al. (2014). On the Bernstein–von Mises phenomenon for nonparametric Bayes procedures. *The Annals of Statistics*, 42(5):1941–1969.
- Celisse, A., Daudin, J.-J., Pierre, L., et al. (2012). Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electronic Journal of Statistics*, 6:1847–1899.
- Dal Maso, G. (2012). *An introduction to Γ -convergence*, volume 8. Springer Science & Business Media.
- Ghosh, J. and Ramamoorthi, R. (2003). *Bayesian Nonparametrics*. Springer Series in Statistics. Springer.
- Hall, P., Ormerod, J. T., and Wand, M. (2011a). Theory of gaussian variational approximation for a Poisson mixed model. *Statistica Sinica*, pages 369–389.
- Hall, P., Pham, T., Wand, M. P., Wang, S. S., et al. (2011b). Asymptotic normality and valid inference for Gaussian variational approximation. *The Annals of Statistics*, 39(5):2502–2532.
- Hoffman, M., Blei, D., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347.
- Hoffman, M. D. and Gelman, A. (2014). The No-U-Turn sampler. *JMLR*, 15(1):1593–1623.
- Hofman, J. and Wiggins, C. (2008). Bayesian approach to network modularity. *Physical Review Letters*, 100(25).
- Jiang, J. (2007). *Linear and generalized linear mixed models and their applications*. Springer Science & Business Media.
- Kleijn, B., Van der Vaart, A., et al. (2012). The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381.
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., and Blei, D. M. (2016). Automatic differentiation variational inference. *arXiv preprint arXiv:1603.00788*.

- Laplace, P. (1809). Memoire sur les integrales definies et leur application aux probabilites, et specialement a la recherche du milieu qu'il faut choisir entre les resultats des observations. *Memoires presentes a l'Academie des Sciences, Paris*.
- Le Cam, L. (1953). On some asymptotic properties of maximum likelihood estimates and related bayes' estimates. *Univ. Calif. Publ. in Statist.*, 1:277–330.
- Le Cam, L. and Yang, G. L. (2012). *Asymptotics in statistics: some basic concepts*. Springer Science & Business Media.
- Lehmann, E. L. and Casella, G. (2006). *Theory of point estimation*. Springer Science & Business Media.
- Lu, Y., Stuart, A. M., and Weber, H. (2016). Gaussian approximations for probability measures on \mathbf{R}^d . *arXiv preprint arXiv:1611.08642*.
- McCullagh, P. (1984). Generalized linear models. *European Journal of Operational Research*, 16(3):285–292.
- McCulloch, C. E. and Neuhaus, J. M. (2001). *Generalized linear mixed models*. Wiley Online Library.
- Mossel, E., Neeman, J., and Sly, A. (2012). Stochastic block models and reconstruction. *arXiv preprint arXiv:1202.1499*.
- Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
- Ormerod, J. T., You, C., and Muller, S. (2014). A variational Bayes approach to variable selection. Technical report, Citeseer.
- Snijders, T. A. and Nowicki, K. (1997). Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of classification*, 14(1):75–100.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- Von Mises, R. (1931). *Wahrscheinlichkeitsrechnung*.
- Wang, B. and Titterington, D. (2005). Inadequacy of interval estimates corresponding to variational bayesian approximations. In *AISTATS*.
- Wang, B., Titterington, D., et al. (2006). Convergence properties of a general algorithm for calculating variational bayesian estimates for a normal mixture model. *Bayesian Analysis*, 1(3):625–650.
- Wang, Y. J. and Wong, G. Y. (1987). Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*, 82(397):8–19.
- Westling, T. and McCormick, T. H. (2015). Establishing consistency and improving uncertainty estimates of variational inference through M-estimation. *arXiv preprint arXiv:1510.08151*.
- You, C., Ormerod, J. T., and Müller, S. (2014). On variational Bayes estimation and variational information criteria for linear regression models. *Australian & New Zealand Journal of Statistics*, 56(1):73–87.

Proofs

A Proof of Lemma 1

What we need to show here is that our consistent testability assumption implies assumption (2.3) in [Kleijn et al. \(2012\)](#):

$$\int_{\tilde{\theta} > M_n} \pi_{\tilde{\theta}}^*(\tilde{\theta} | x) d\tilde{\theta} \xrightarrow{P_{\theta_0}} 0$$

for every sequence of constants $M_n \rightarrow \infty$, where $\tilde{\theta} = \delta_n^{-1}(\theta - \theta_0)$.

This is a consequence of a slight generalization of Theorem 3.1 of [Kleijn et al. \(2012\)](#). That theorem shows this implication for the iid case with a common \sqrt{n} -convergence rate for all dimensions of θ . Specifically, they rely on a suitable test sequence under misspecification of uniform exponential power around the true value θ_0 to split the posterior measure.

To show this implication in our case, we replace all \sqrt{n} by δ_n^{-1} in the proofs of Theorem 3.1, Theorem 3.3, Lemma 3.3, Lemma 3.4 of [Kleijn et al. \(2012\)](#). We refer the readers to [Kleijn et al. \(2012\)](#) and omit the proof here.

B Proof of Lemma 2

We first perform a change of variable step regarding the mean field variational family. In light of Lemma 1, we know that the VB ideal degenerates to a point mass at the rate of δ_n^{-1} . We need to assume a variational family that degenerates to points masses at the same rate as the ideal VB posterior. This is because if the variational distribution converges to a point mass faster than $\pi(\theta | x)$, then the KL divergence between them will converge to $+\infty$. This makes the KL minimization meaningless as n increases.⁵

To avoid this pathology, we assume a variational family for the rescaled and re-centered θ , $\check{\theta} := \delta_n^{-1}(\theta - \mu)$, for some $\mu \in \Theta$. This is a centered and scaled transformation of θ , centered to an arbitrary μ . (In contrast, the previous transformation $\tilde{\theta}$ was centering θ at the true θ_0 .) With this transformation, the variational family is

$$q_{\check{\theta}}(\check{\theta}) = q(\mu + \delta_n \check{\theta}) |\det(\delta_n)|, \quad (18)$$

where $q(\cdot)$ is the original mean field variational family. We will overload the notation in Section 1 and write this transformed family as $q(\theta)$ and the corresponding family \mathcal{Q}^d .

In this family, for each fixed μ , θ degenerates to a point mass at μ as n goes to infinity. μ is not necessarily equal to the true value θ_0 . We allow μ to vary throughout the parameter space Θ so that \mathcal{Q}^d does not restrict what the distributions degenerate to. \mathcal{Q}^d only constrains that the variational distribution degenerates to some point mass at the rate of δ_n . This step also does not restrict the applicability of the theoretical results. In practice, we always have finite samples with fixed n , so assuming a fixed variational family for $\check{\theta}$ as opposed to θ amounts to a change of variable $\theta = \mu + \delta_n \check{\theta}$.

Next we show consistency of the KL minimizer of the VB ideal.

To show consistency, we need to show that the mass of the KL minimizer

$$q^\ddagger := \arg \min_{q(\theta) \in \mathcal{Q}^d} \text{KL}(q(\theta) || \pi^*(\theta | x))$$

concentrates near θ_0 as $n \rightarrow \infty$. That is,

$$\int_{B(\theta_0, \xi_n)} q^\ddagger(\theta) d\theta \xrightarrow{P_{\theta_0}} 1, \quad (19)$$

for some $\xi_n \rightarrow 0$ as $n \rightarrow \infty$. This implies

$$q^\ddagger(\theta) \stackrel{d}{\rightarrow} \delta_{\theta_0}$$

⁵Equation (36) and Equation (106) in the proof exemplify this claim.

by the Slutsky's theorem.

To begin with, we first claim that

$$\limsup_{n \rightarrow \infty} \min \text{KL}(q(\theta) || \pi^*(\theta | x)) \leq M, \quad (20)$$

for some constant $M > 0$, and

$$\int_{\mathbb{R}^d \setminus K} q^\ddagger(\theta) d\theta \rightarrow 0, \quad (21)$$

where K is the compact set assumed in the local asymptotic normality condition.

The first claim says that the limiting minimum KL divergence is upper bounded. The intuition is that a choice of $q(\theta)$ with $\mu = \theta_0$ would have a finite KL divergence in the limit. This is because (roughly) $\pi^*(\theta | x)$ converges to a normal distribution centered at θ_0 with rate δ_n , so it suffices to have a $q(\theta)$ that shares the same center and the same rate of convergence.

The second claim says that the restriction of $q^\ddagger(\theta)$ to the compact set K , due to the set compactness needed in the local asymptotic normality (LAN) condition, will not affect our conclusion in the limit. This is because the family of \mathcal{Q}^d we assume has a shrinking-to-zero scale. In this way, as long as μ resides within K , $q^\ddagger(\theta)$ will eventually be very close to its renormalized restriction to the compact set K , $q^{\ddagger,K}(\theta)$, where

$$q^{\ddagger,K}(\theta) = \frac{q^\ddagger(\theta) \cdot \mathbb{1}_\theta(K)}{\int q^\ddagger(\theta) \cdot \mathbb{1}_\theta(K) d\theta}.$$

We will prove these claims at the end.

To show $\int_{B(\theta_0, \xi_n)} q^{\ddagger,K}(\theta) d\theta \xrightarrow{P_{\theta_0}} 1$, we both upper bound and lower bound this integral. This step mimicks the Step 2 in the proof of Lemma 3.6 along with Lemma 3.7 in [Lu et al. \(2016\)](#).

We first upper bound the integral using the LAN condition,

$$\begin{aligned} & \int q^{\ddagger,K}(\theta) M_n(\theta; x) d\theta \\ &= \int q^{\ddagger,K}(\theta) \end{aligned} \quad (22)$$

$$\times \left[M_n(\theta_0; x) + \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n, \theta_0} - \frac{1}{2} [\delta_n^{-1}(\theta - \theta_0)]^\top V_{\theta_0} [\delta_n^{-1}(\theta - \theta_0)] + o_P(1) \right] d\theta \quad (23)$$

$$\leq M_n(\theta_0; x) - C_1 \sum_{i=1}^d \frac{\eta^2}{\delta_{n,ii}^2} \int_{B(\theta_0, \eta)^c} q^{\ddagger,K}(\theta) d\theta + o_P(1), \quad (24)$$

for large enough n and $\eta \ll 1$ and some constant $C_1 > 0$. The first equality is due to the LAN condition. The second inequality is due to the domination of quadratic term for large n .

Then we lower bound the integral using our first claim. By

$$\limsup_{n \rightarrow \infty} \text{KL}(q^{\ddagger,K}(\theta) || \pi^*(\theta | x)) \leq M,$$

we can have

$$\int q^{\ddagger,K}(\theta) M_n(\theta; x) d\theta \geq M_n(\theta_0; x) - M_0, \quad (25)$$

for some large constant $M_0 > M$.

This step is due to a couple of steps of technical calculation and the LAN condition. To show this implication, we first rewrite the KL divergence as follows.

$$\text{KL}(q^{\ddagger,K}(\theta) || \pi^*(\theta | x)) \quad (26)$$

$$= \int q^{\ddagger,K}(\theta) \log q^{\ddagger,K}(\theta) d\theta - \int q^{\ddagger,K}(\theta) \log \pi^*(\theta | x) d\theta \quad (27)$$

$$= \sum_{i=1}^d \int [\delta_{n,ii}^{-1} q_{h,i}^{\ddagger,K}(h)] \log [\delta_{n,ii}^{-1} q_{h,i}^{\ddagger,K}(h)] \delta_{n,ii} dh - \int q^{\ddagger,K}(\theta) \log \pi^*(\theta | x) d\theta \quad (28)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}^{\ddagger,K}(h)) - \int q^{\ddagger,K}(\theta) \log \pi^*(\theta | x) d\theta, \quad (29)$$

where this calculation is due to the form of the \mathcal{Q}^d family we assume and a change of variable of $h = \delta_n^{-1}(\theta - \mu)$; h is in the same spirit as $\check{\theta}$ above. Notation wise, μ is the location parameter specific to $q^{\ddagger,K}(\theta)$ and $\mathbb{H}(q_h(h))$ denotes the entropy of distribution q_h .

We further approximate the last term by the LAN condition.

$$\int q^{\ddagger,K}(\theta) \log \pi^*(\theta | x) d\theta \quad (30)$$

$$= \int q^{\ddagger,K}(\theta) \log \frac{p(\theta) \exp(M_n(\theta; x))}{\int p(\theta) \exp(M_n(\theta; x)) d\theta} d\theta \quad (31)$$

$$= \int q^{\ddagger,K}(\theta) \log p(\theta) d\theta + \int q(\theta) M_n(\theta; x) d\theta - \log \int p(\theta) \exp(M_n(\theta; x)) d\theta \quad (32)$$

$$= \int q^{\ddagger,K}(\theta) \log p(\theta) d\theta + \int q^{\ddagger,K}(\theta) M_n(\theta; x) d\theta - \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + \log \det(\delta_n) + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \right]. \quad (33)$$

This first equality is due to the definition of $\pi^*(\theta | x)$. The second equality is due to $\int q^{\ddagger,K}(\theta) d\theta = 1$. The third equality is due to Laplace approximation and the LAN condition.

Going back to the KL divergence, this approximation gives

$$\text{KL}(q^{\ddagger,K}(\theta) || \pi^*(\theta | x)) \quad (34)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}^{\ddagger,K}(h)) - \int q^{\ddagger,K}(\theta) \log \pi^*(\theta | x) d\theta \quad (35)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}^{\ddagger,K}(h)) - \int q^{\ddagger,K}(\theta) \log p(\theta) d\theta - \int q^{\ddagger,K}(\theta) M_n(\theta; x) d\theta + \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + \log \det(\delta_n) + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \right] \quad (36)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}^{\ddagger,K}(h)) - \int q^{\ddagger,K}(\theta) \log p(\theta) d\theta - \int q^{\ddagger,K}(\theta) M_n(\theta; x) d\theta + \frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1). \quad (37)$$

The first equality is exactly Equation (29). The second equality is due to Equation (33). The third equality is due to the cancellation of the two $\log \det(\delta_n)$ terms. This exemplifies why we assumed the convergence rate of the \mathcal{Q}^d family in the first place; we need to avoid the KL divergence going to infinity.

By

$$\limsup_{n \rightarrow \infty} \text{KL}(q^{\ddagger,K}(\theta) || \pi^*(\theta | x)) \leq M,$$

we have

$$\int q^{\ddagger,K}(\theta) M_n(\theta; x) d\theta \quad (38)$$

$$\geq -M + \sum_{i=1}^d \mathbb{H}(q_{h,i}^{\ddagger,K}(h)) - \int q^{\ddagger,K}(\theta) \log p(\theta) d\theta + \frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \quad (39)$$

$$\geq -M_0 + M_n(\theta_0; x) + o_P(1) \quad (40)$$

for some constant $M_0 > 0$. This can be achieved by choosing a large enough M_0 to make the last inequality true. This is doable because all the terms does not change with n except $\int q^{\ddagger,K}(\theta) \log p(\theta) d\theta$. And we have $\limsup_{n \rightarrow \infty} \int q^{\ddagger,K}(\theta) \log p(\theta) d\theta < \infty$ due to our prior mass condition.

Now combining Equation (40) and Equation (24), we have

$$M_n(\theta_0; x) - C_1 \sum_{i=1}^d \frac{\eta^2}{\delta_{n,ii}^2} \int_{B(\theta_0, \eta)^c} q^{\ddagger, K}(\theta) d\theta + o_P(1) \geq -M_0 + M_n(\theta_0; x).$$

This gives

$$\int_{B(\theta_0, \eta)^c} q^{\ddagger, K}(\theta) d\theta + o_P(1) \leq \frac{M_0 \cdot (\min_i \delta_{n,ii})^2}{C_2 \eta^2},$$

for some constant $C_2 > 0$. The right side of the inequality will go to zero as n goes to infinity if we choose $\eta = \sqrt{M_0(\min_i \delta_{n,ii})}/C_2 \rightarrow 0$. That is, we just showed Equation (19) with $\xi_n = \eta$.

We are now left to show the two claims we made at the beginning.

To show Equation (20), it suffices to show that there exists a choice of $q(\theta)$ such that

$$\limsup_{n \rightarrow \infty} \text{KL}(q(\theta) || \pi^*(\theta | x)) < \infty.$$

We choose $\tilde{q}(\theta) = \prod_{i=1}^d N(\theta_i; \theta_{0,i}, \delta_{n,ii}^2 v_i)$ for $v_i > 0, i = 1, \dots, d$. We thus have

$$\text{KL}(\tilde{q}(\theta) || \pi^*(\theta | x)) \tag{41}$$

$$= \sum_{i=1}^d \frac{1}{2} \log(v_i) + \frac{d}{2} + d \log(2\pi) - \int \tilde{q}(\theta) \log p(\theta) d\theta - \int \tilde{q}(\theta) M_n(\theta; x) d\theta - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \tag{42}$$

$$= \sum_{i=1}^d \frac{1}{2} \log(v_i) + \frac{d}{2} + d \log(2\pi) - \log p(\theta_0) - M_n(\theta_0; x) - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \tag{43}$$

$$= \sum_{i=1}^d \frac{1}{2} \log(v_i) + \frac{d}{2} + d \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + C_6 + o_P(1), \tag{44}$$

for some constant $C_6 > 0$. The finiteness of limsup is due to the last term being bounded in the limit. The first equality is due to the same calculation as in Equation (43). The third equality is due to the cancellation of the two $M_n(\theta_0; x)$ terms and the two $p(\theta_0)$ terms; this renders the whole term independent of n . The second equality is due to the limit of $\tilde{q}(\theta)$ concentrating around θ_0 . Specifically, we expand $\log p(\theta)$ to the second order around θ_0 ,

$$\int \tilde{q}(\theta) \log p(\theta) d\theta \tag{45}$$

$$= \log p(\theta_0) + \int \tilde{q}(\theta) \left[(\theta - \theta_0)(\log p(\theta_0))' + \frac{(\theta - \theta_0)^2}{2} \int_0^1 (\log p(\xi\theta + (1-\xi)\theta_0))'' (1-\xi)^2 d\xi \right] d\theta \tag{46}$$

$$\leq \log p(\theta_0) + \frac{1}{2!} \max_{\xi \in [0,1]} \left\{ \int \tilde{q}(\theta) (\theta - \theta_0)^2 (\log p(\xi\theta + (1-\xi)\theta_0))'' d\theta \right\} \tag{47}$$

$$\leq \log p(\theta_0) + \frac{M_p}{\sqrt{(2\pi)^d \det(\delta_n^2) \prod_i v_i}} \int_{\mathbb{R}^d} |\theta|^2 e^{(|\theta| + |\theta_0|)^2} \cdot e^{-\frac{1}{2}\theta^\top (\delta_n V \delta_n)^{-1} \theta} d\theta \tag{48}$$

$$\leq \log p(\theta_0) + \frac{M_p}{\sqrt{(2\pi)^d \det(\delta_n^2) \prod_i v_i}} e^{\theta_0^2} \int_{\mathbb{R}^d} |\theta|^2 e^{-\frac{1}{2}\theta^\top [(\delta_n V \delta_n)^{-1} - 2I_d] \theta} d\theta \tag{49}$$

$$\leq \log p(\theta_0) + C_3 M_p e^{\theta_0^2} \max_d (\delta_{n,ii}^2) \det(V^{-1} - 2\delta_n^2)^{-1} \tag{50}$$

$$\leq \log p(\theta_0) + C_4 \max_d (\delta_{n,ii}^2) \tag{51}$$

where $\max_d (\delta_{n,ii}^2) \rightarrow 0$ and $C_3, C_4 > 0$. The first equality is due to Taylor expansion with integral form residuals. The second inequality is due to the first order derivative terms equal to zero and taking

the maximum of the second order derivative. The third inequality is due to the prior mass condition where we assume the second derivative of $\log p(\theta)$ is bounded by $M_p e^{|\theta|^2}$ for some constant $M_p > 0$. The fourth inequality is pulling $e^{\theta_0^2}$ out of the integral. The fifth inequality is due to rescaling θ by its covariance matrix and appealing to the mean of a Chi-squared distribution with d degrees of freedom. The sixth (and last) inequality is due to $\det(V^{-1} - 2\delta_n^2)^{-1} > 0$ for large enough n .

We apply the same Taylor expansion argument to the $\int \tilde{q}(\theta)M_n(\theta; x)d\theta$.

$$\int_{K_n} \tilde{q}(\theta)M_n(\theta; x) d\theta \quad (52)$$

$$= M_n(\theta_0; x) + \int_{K_n} \tilde{q}(\theta) \left[\delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n, \theta_0} + \frac{1}{2}(\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) + o_P(1) \right] d\theta \quad (53)$$

$$\leq M_n(\theta_0; x) + \frac{1}{2} \text{Tr}(V_{\theta_0} V) + o_P(1) \quad (54)$$

$$\leq M_n(\theta_0; x) + C_6 + o_P(1) \quad (55)$$

where K_n is a compact set. The first equality is due to the LAN condition. The second inequality is due to $\tilde{q}(\theta)$ centered at θ_0 with covariance $\delta_n V \delta_n$. The third inequalities are true for $C_6 > 0$.

For the set outside of this compact set K_n , we consider for a general choice of q distribution, $\tilde{q}(\theta) = \mathcal{N}(\theta; \theta_0 + \Delta_{n, \theta_0}, \delta_n V_{\theta_0} \delta_n)$, of which $\tilde{q}(\theta) = \prod_{i=1}^d \mathcal{N}(\theta_i; \theta_{0,i}, \delta_{n,ii}^2 v_i)$ we work with is a special case.

$$\int_{\mathbb{R}^d \setminus K_n} \tilde{q}(\theta)(\log p(\theta) + M_n(\theta; x)) d\theta \quad (56)$$

$$\leq C_7 \int_{\mathbb{R}^d \setminus K_n} \mathcal{N}(\theta; \theta_0 + \Delta_{n, \theta_0}, \delta_n V_{\theta_0} \delta_n)(\log p(\theta) + M_n(\theta; x)) d\theta \quad (57)$$

$$\leq C_8 [\det(\delta_n)^{-1} \log(\det(\delta_n)^{-1})] \int_{\mathbb{R}^d \setminus K_n} \mathcal{N}(\tilde{\theta}; \Delta_{n, \theta_0}, V_{\theta_0}) \log \pi^*(\tilde{\theta} | x) \det(\delta_n) d\tilde{\theta} \quad (58)$$

$$\leq C_9 \log(\det(\delta_n)^{-1}) \int_{\mathbb{R}^d \setminus K_n} [\pi^*(\tilde{\theta} | x) + o_P(1)] \log \pi^*(\tilde{\theta} | x), V_{\theta_0}) d\tilde{\theta} \quad (59)$$

$$\leq C_{10} \log(\det(\delta_n)^{-1}) \int_{\mathbb{R}^d \setminus K_n} [\mathcal{N}(\tilde{\theta}; \Delta_{n, \theta_0}, V_{\theta_0}) + o_P(1)] \log \mathcal{N}(\tilde{\theta}; \Delta_{n, \theta_0}, V_{\theta_0}) d\tilde{\theta} \quad (60)$$

$$\leq o_P(1) \quad (61)$$

for some $C_7, C_8, C_9, C_{10} > 0$. The first inequality is due to $\tilde{q}(\theta)$ centered at θ_0 and with rate of convergence δ_n . The second inequality is due to a change of variable $\tilde{\theta} = \delta_n^{-1}(\theta - \theta_0)$. The third inequality is due to Lemma 1. The fourth inequality is due to Lemma 1 and Theorem 2 in [Piera and Parada \(2009\)](#). The fifth inequality is due to a choice of fast enough increasing sequence of compact sets K_n .

The lower bound of $\int \tilde{q}(\theta)(\log p(\theta) + M_n(\theta; x))d\theta$ can be derived with exactly the same argument. Our first claim Equation (20) is thus proved.

To show our second claim Equation (21), we first denote $B(\mu, M)$ as the largest ball centered at μ and contained in the compact set K . We know by the construct of \mathcal{Q}^d — \mathcal{Q}^d has a shrinking-to-zero scale — that for each $\epsilon > 0$, there exists an N such that for all $n > N$ we have $\int_{\|\theta - \mu\| > M} q(\theta) d\theta < \epsilon$. Therefore, we have

$$\int_{\mathbb{R}^d \setminus K} q^\ddagger(\theta) d\theta \leq \int_{\mathbb{R}^d \setminus B(\mu, M)} q^\ddagger(\theta) d\theta \leq \epsilon.$$

C Proof of Lemma 3

To show the convergence of optimizers from two minimization problems, we invoke Γ -convergence. It is a classical technique in characterizing variational problems. A major reason is that if two functionals Γ -converge, then their minimizer also converge.

We recall the definition of Γ -convergence ([Dal Maso, 2012](#); [Braides, 2006](#)).

Definition 13. Let \mathcal{X} be a metric space and $F_\epsilon : \mathcal{X} \rightarrow \mathbb{R}$ a family of functionals indexed by $\epsilon > 0$. Then the existence of a limiting functional F_0 , the Γ -limit of F_ϵ , as $\epsilon \rightarrow 0$, relies on two conditions:

1. (liminf inequality) for every $x \in \mathcal{X}$ and for every $x_\epsilon \rightarrow x$, we have

$$F_0(x) \leq \liminf_{\epsilon \rightarrow 0} F_\epsilon(x_\epsilon),$$

2. (limsup inequality / existence of a recovery sequence) for every $x \in \mathcal{X}$ we can find a sequence $\tilde{x}_\epsilon \rightarrow x$ such that

$$F_0(x) \geq \limsup_{\epsilon \rightarrow 0} F_\epsilon(\tilde{x}_\epsilon).$$

The first condition says that F_0 is a lower bound for the sequence F_ϵ , in the sense that $F_0(x) \leq F_\epsilon(x_\epsilon) + o(1)$ whenever $x_\epsilon \rightarrow x$. Together with the first condition, the second condition implies that $F_0(x) = \lim_{\epsilon \rightarrow 0} F_\epsilon(\tilde{x}_\epsilon)$, so that the lower bound is sharp.

Γ -convergence is particularly useful for variational problems due to the following fundamental theorem. Before stating the theorem, we first define equi-coerciveness.

Definition 14. (Equi-coerciveness of functionals) A sequence $F_\epsilon : \mathcal{X} \rightarrow \bar{\mathbb{R}}$ is equi-coercive if for all $\epsilon_j \rightarrow 0$ and x_j such that $F_{\epsilon_j}(x_j) \leq t$ there exist a subsequence of j (not relabeled) and a converging sequence x'_j such that $F_{\epsilon_j}(x'_j) \leq F_{\epsilon_j}(x_j) + o(1)$.

Equi-coerciveness of functionals ensures that we can find a precompact minimizing sequence of F_ϵ such that the convergence $x_\epsilon \rightarrow x$ can take place. Now we are ready to state the fundamental theorem.

Theorem 15. (Fundamental theorem of Γ -convergence) Let \mathcal{X} be a metric space. Let (F_ϵ) be an equi-coercive sequence of functions on \mathcal{X} . Let $F = \Gamma\text{-}\lim_{\epsilon \rightarrow 0} F_\epsilon$, then

$$\operatorname{argmin}_{\mathcal{X}} F = \lim_{\epsilon \rightarrow 0} \operatorname{argmin}_{\mathcal{X}} F_\epsilon.$$

The above theorem implies that if all functions F_ϵ admit a minimizer x_ϵ then, up to subsequences, x_ϵ converge to a minimum point of F . We remark that the converse is not true; we may have minimizers of F which are not limits of minimizers of F_ϵ , e.g. $F_\epsilon(t) = \epsilon t^2$ (Braides, 2006).

In this way, Γ -convergence is convenient to use when we would like to study the asymptotic behavior of a family of problem F_ϵ through defining a limiting problem F_0 which is a ‘good approximation’ such that the minimizers converge: $x_\epsilon \rightarrow x_0$, where x_0 is a minimizer of F_0 . Conversely, we can characterize solutions of a difficult F_0 by finding easier approximating F_ϵ (Braides, 2006).

We now prove Lemma 3 for the general mean field family. The family is parametric as in Section 1, so we assume it is indexed by some finite dimensional parameter m . We want to show that the functionals

$$F_n(m) := \text{KL}(q(\theta; m) || \pi^*(\theta | x))$$

Γ -converge to

$$F_0(m) := \text{KL}(q(\theta; m) || \mathcal{N}(\theta; \theta_0 + \delta_n \Delta_{n, \theta_0}, \delta_n V_{\theta_0}^{-1} \delta_n)) - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}$$

in probability as $n \rightarrow 0$. Recall that the mean field family has density

$$q(\theta) = \prod_{i=1}^d \delta_{n, ii}^{-1} q_{h, i}(h),$$

where $h = \delta_n^{-1}(\theta - \theta_0)$.

We need the following mild technical conditions on \mathcal{Q}^d .

Following the change-of-variable step detailed in the beginning of Appendix B, we consider the mean field variational family with densities $q(\theta) = \prod_{i=1}^d \delta_{n, ii}^{-1} q_{h, i}(h)$, where $h = \delta_n^{-1}(\theta - \mu)$ for some $\mu \in \Theta$.

Assumption 2. We assume the following conditions on $q_{h,i}$:

1. $q_{h,i}, i = 1, \dots, d$ have continuous densities.
2. $q_{h,i}, i = 1, \dots, d$ have positive and finite entropies.
3. $\int q'_{h,i}(h) dh < \infty, i = 1, \dots, d$.

The last condition ensures that convergence in finite dimensional parameters implied convergence in TV distance. This is due to a Taylor expansion argument:

$$\text{KL}(q(\theta; m) || q(\theta; m + \delta)) \quad (62)$$

$$= \int q(\theta; m) \cdot (\log q(\theta; m) - \log q(\theta; m + \delta)) d\theta \quad (63)$$

$$= \int q(\theta; m) \cdot (\delta \cdot (\log q(\theta; m))') d\theta + o(1) \quad (64)$$

$$= \delta \int (q(\theta; m))' d\theta \quad (65)$$

$$< \epsilon \quad (66)$$

The last step is true if Assumption 2 is true for q . We also notice that convergence in KL divergence implies convergence in TV distance. Therefore, Assumption 2 implies that convergence in finite dimensional parameter implies convergence in TV distance.

Together with Theorem 15, the Γ -convergence of the two functionals implies $m_n \xrightarrow{P_{\theta_0}} m_0$ where m_n is the minimizer of F_n for each n and m_0 is the minimizer of F_0 . This is due to the last term of $F_0 - \Delta_{n,\theta_0}^\top V_{\theta_0} \Delta_{n,\theta_0}$ is a constant bounded in P_{θ_0} probability and independent of m . The convergence in total variation then follows from Assumption 2 and our argument above.

Lastly, we prove the Γ -convergence of the two functionals for the mean field family.

We first rewrite $F_n(m, \mu)$.

$$F_n(m, \mu) := \text{KL}(q(\theta; m, \mu) || \pi^*(\theta | x)) \quad (67)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int q(\theta; m, \mu) \log \pi^*(\theta | x) d\theta \quad (68)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int q(\theta; m, \mu) \log p(\theta) d\theta - \int q(\theta; m, \mu) M_n(\theta; x) d\theta + \log \int p(\theta) \exp(M_n(\theta; x)) d\theta \quad (69)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int q(\theta; m, \mu) \log p(\theta) d\theta - \int q(\theta; m, \mu) M_n(\theta; x) d\theta + \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + \log \det(\delta_n) + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \right] \quad (70)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int q(\theta; m, \mu) \log p(\theta) d\theta - \int q(\theta; m, \mu) M_n(\theta; x) d\theta + \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \right] \quad (71)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int q(\theta; m, \mu) M_n(\theta; x) d\theta + \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + o_P(1) \right] \quad (72)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int q(\theta; m, \mu) [M_n(\theta_0; x) + \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n,\theta_0}] d\theta$$

$$-\frac{1}{2}(\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) + o_P(1)] d\theta - \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + o_P(1) \right] \quad (73)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \int \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n,\theta_0} \cdot q(\theta; m, \mu) d\theta \\ + \int \frac{1}{2} (\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) \cdot q(\theta; m, \mu) d\theta - \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + o_P(1) \right] \quad (74)$$

The first equality is by the definition of KL divergence. The second equality is by the definition of the VB ideal. The third equality is due to the Laplace approximation of the normalizer like we did in Equation (33). The fourth equality is due to the cancellation of the two $\log \det(\delta_n)$ terms. This again exemplifies why we assume a fixed variational family on the rescale variable $\tilde{\theta}$. The fifth equality is due to a similar argument as in Equation (43). The sixth equality is due to the LAN condition of $M_n(\theta; x)$. The seventh equality is due to the computation of each term in the integral as an expectation under the distribution $q(\theta)$. To extend the restriction to some compact set K to the whole space \mathbb{R}^d in the sixth equality, we employ the same argument as in Equation (44).

We notice that when $\mu \neq \theta_0$, we will have $F_n(m) \rightarrow \infty$. On the other hand, we have $\limsup F_n < \infty$. This echoes our consistency result in Lemma 2.

Now we rewrite $F_0(m, \mu)$.

$$\text{KL}(q(\theta; m, \mu) \| \mathcal{N}(\theta; \theta_0 + \delta_n \Delta_{n,\theta_0}, \delta_n V_{\theta_0}^{-1} \delta_n)) \quad (75)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) + \int q(\theta; m, \mu) \log \mathcal{N}(\theta; \theta_0 + \delta_n \Delta_{n,\theta_0}, \delta_n V_{\theta_0}^{-1} \delta_n) d\theta \quad (76)$$

$$= \log |\det(\delta_n)|^{-1} + \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) + \frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + \log \det(\delta_n) \\ + \int q(\theta; m, \mu) \cdot (\theta - \theta_0 - \delta_n \Delta_{n,\theta_0})^\top \delta_n^{-1} V_{\theta_0} \delta_n^{-1} (\theta - \theta_0 - \delta_n \Delta_{n,\theta_0}) d\theta \quad (77)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) + \frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + \Delta_{n,\theta_0}^\top V_{\theta_0} \Delta_{n,\theta_0} \\ - \int \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n,\theta_0} \cdot q(\theta; m, \mu) d\theta + \int \frac{1}{2} (\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) \cdot q(\theta; m, \mu) d\theta. \quad (78)$$

This gives

$$F_0(m, \mu) - \Delta_{n,\theta_0}^\top V_{\theta_0} \Delta_{n,\theta_0} \quad (79)$$

$$= \sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \frac{d}{2} \log(2\pi) + \frac{1}{2} \log \det V_{\theta_0} \\ - \int \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n,\theta_0} \cdot q(\theta; m, \mu) d\theta \\ + \int \frac{1}{2} (\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) \cdot q(\theta; m, \mu) d\theta \quad (80)$$

$$= +\infty \cdot (1 - \mathbb{1}_\mu(\theta_0)) + \left[\sum_{i=1}^d \mathbb{H}(q_{h,i}(h; m)) - \frac{d}{2} \log(2\pi) + \frac{1}{2} \log \det V_{\theta_0} \right. \\ \left. - \int \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n,\theta_0} \cdot q(\theta; m, \mu) d\theta \right. \\ \left. + \int \frac{1}{2} (\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) \cdot q(\theta; m, \mu) d\theta \right] \cdot \mathbb{1}_\mu(\theta_0). \quad (81)$$

The last step is due to our definition of our variational family $q(\theta; m, \mu) = \prod_{i=1}^d \delta_{n,i}^{-1} q_{h,i}(h; m)$, where $h = \delta_n^{-1}(\theta - \mu)$ for some $\mu \in \Theta$. The last step is true as long as the $q_{h,i}$ distributions are not point masses at zero. This is ensured by positive entropy in Assumption 2.

Comparing Equation (74) and Equation (81), we can prove the Γ convergence. Let $m_n \rightarrow m$. When $\mu \neq \theta_0$, $\liminf_{n \rightarrow \infty} F_n(m_n, \mu) = +\infty$. The limsup inequality is automatically satisfied. When $\mu = \theta_0$,

we have $F_n(m, \mu) = F_0(m, \mu) - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0} + o_P(1)$. This implies $F_0(m, \mu) \leq \lim_{n \rightarrow \infty} F_n(m_n, \mu)$ in P_{θ_0} probability by the continuity of F_n ensured by Assumption 2.

We then show the existence of a recovery sequence. When $\mu \neq \theta_0$, $F_0(m, \mu) = +\infty$. The limsup inequality is automatically satisfied. When $\mu = \theta_0$, we can simply choose $m_n = \theta_0$. The limsup inequality is again ensured by $F_0(m, \mu) \leq \lim_{n \rightarrow \infty} F_n(m_n, \mu)$ in P_{θ_0} probability and the continuity of F_n . The Γ -convergence of the F functionals is shown.

We notice that $\Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}$ does not depend on m or μ so that $\operatorname{argmin} F_0 = \operatorname{argmin} F_0 - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}$. The convergence of the KL minimizers is thus proved.

D Proof of Lemma 4

Notice that the mean field variational families $\mathcal{Q}^d = \{q : q(\theta) = \prod_{i=1}^d \delta_{n, ii}^{-1} q_{h, i}(h), \text{ where } h = \delta_n^{-1}(\theta - \mu) \text{ for some } \mu \in \Theta\}$, or the Gaussian family $\{q : q(\theta) = N(m, \delta_n \Sigma \delta_n)\}$ can be written in the form of

$$q(\theta) = |\det \delta_n|^{-1} q_h(\delta_n^{-1}(\theta - \mu))$$

for some $\mu \in \mathbb{R}^d$, and $\int q_h(h) dh = 0$. This form is due to a change-of-variable step we detailed in the beginning of Appendix B.

We first specify the mild technical conditions on \mathcal{Q}^d .

Assumption 3. *We assume the following conditions on q_h .*

1. *If q_h has zero mean, we assume $\int h^2 \cdot q_h(h) dh < \infty$ and $\sup_{z, x} |(\log p(z, x | \theta))'| \leq C_{11} \cdot q_h(\theta)^{-C_{12}}$ for some $C_{11}, C_{12} > 0$; $|M_n(\theta; x)| \leq C_{13} \cdot q_h(\theta)^{-C_{14}}$ for some $C_{13}, C_{14} > 0$.*
2. *If q_h has nonzero mean, we assume $\int h \cdot q_h(h) dh < \infty$ and $\sup_{z, x} |(\log p(z, x | \theta))'| \leq C_{11} \cdot q_h(\theta)^{-C_{12}}$ for some $C_{11}, C_{12} > 0$; $|M_n(\theta; x)| \leq C_{13} \cdot q_h(\theta)^{-C_{14}}$ for some $C_{13}, C_{14} > 0$.*

The assumption first assumes finite moments for q_h so that we can properly apply a Taylor expansion argument. The second part of this assumption makes sure the derivative of $\log p(z, x | \theta)$ does not increase faster than the tail decrease of $q_h(\cdot)$. For example, if $q_h(\cdot)$ is normal, then the second part writes $\sup_{z, x} |(\log p(z, x | \theta))'| \leq C_{15} \exp(\theta^2)$ for some $C_{15} > 0$, and $|M_n(\theta; x)| \leq C_{13} \exp(\theta^2)$ for some $C_{13} > 0$. The latter is satisfied by the LAN condition. This is in general a rather weak condition. We usually would not expect the derivative of $\log p(z, x | \theta)$ and $M_n(\theta; x)$ to increase this fast as θ increases.

Now we are ready to prove the lemma.

We first approximate the profiled evidence lower bound (ELBO), $\text{ELBO}_p(q(\theta))$.

$$\text{ELBO}_p(q(\theta)) \tag{82}$$

$$:= \sup_{q(z)} \int q(\theta) \left(\log \left[p(\theta) \exp \left\{ \int q(z) \log \frac{p(z, x | \theta)}{q(z)} dz \right\} \right] - \log q(\theta) \right) d\theta, \tag{83}$$

$$= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(\theta) \int q(z) \log \frac{p(z, x | \theta)}{q(z)} dz d\theta \tag{84}$$

$$= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta$$

$$+ \sup_{q(z)} \int q(\theta) \left[\int q(z) \log \frac{p(z, x | \mu)}{q(z)} dz + (\theta - \mu) \left(\int q(z) \log \frac{p(z, x | \mu)}{q(z)} dz \right)' \right]$$

$$= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x | \mu)}{q(z)} dz$$

$$+ \int q(\theta) \frac{1}{2} |\theta - \mu|^2 \left(\int q(z) \log \frac{p(z, x | \mu)}{q(z)} dz \right)'' d\theta \tag{85}$$

$$\begin{aligned}
&\geq \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \\
&\quad - C_{15} \int q(\theta) \frac{1}{2} |\theta - \mu|^2 q_h(\theta)^{-C_{12}} d\theta \tag{86}
\end{aligned}$$

$$\begin{aligned}
&= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \\
&\quad - C_{15} \int q_h(h) \frac{1}{2} |\delta_n h|^2 q_h(\mu + \delta_n h)^{-C_{12}} dh \tag{87}
\end{aligned}$$

$$\begin{aligned}
&\geq \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \\
&\quad - C_{16} \min_i (\delta_{n,ii}^2) \tag{88}
\end{aligned}$$

$$= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz + o(1), \tag{89}$$

for some constant $C_{16} > 0$. The first equality is by the definition of $\text{ELBO}_p(q(\theta))$. The second equality is rewriting the integrand. The third equality is due to mean value theorem where $\tilde{\theta}^\dagger$ is some value between μ and θ . (A very similar argument for q_h with nonzero means can be made starting from here, that is expanding only to the first order term.) The fourth equality is due to $q_h(\cdot)$ having zero mean. The fifth inequality is due to the second part of Assumption 3. The sixth inequality is due to a change of variable $h = \delta_n(\theta - \mu)$. The seventh inequality is due to $q_h(\cdot)$ residing within exponential family with finite second moment (the first part of Assumption 3). The eighth equality is due to $\delta_n \rightarrow 0$ as $n \rightarrow \infty$.

We now approximate $-\text{KL}(q(\theta) \parallel \pi^*(\theta|x))$ in a similar way.

$$-\text{KL}(q(\theta) \parallel \pi^*(\theta|x)) \tag{90}$$

$$= \int q(\theta) \log \frac{p(\theta) \exp(M_n(\theta; x))}{q(\theta)} d\theta \tag{91}$$

$$= \int q(\theta) \left(\log \left[p(\theta) \exp \left\{ \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} \right\} \right] - \log q(\theta) \right) d\theta \tag{92}$$

$$\begin{aligned}
&= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \int q(\theta) \left[\sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \right. \\
&\quad \left. + (\theta - \mu) \left(\sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \right)' + \frac{1}{2} |\theta - \mu|^2 \left(\sup_{q(z)} \int q(z) \log \frac{p(z, x|\tilde{\theta}^\dagger)}{q(z)} dz \right)'' \right] d\theta \tag{93}
\end{aligned}$$

$$\begin{aligned}
&= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \\
&\quad + \int q(\theta) \frac{1}{2} |\theta - \mu|^2 \left(\sup_{q(z)} \int q(z) \log \frac{p(z, x|\tilde{\theta}^\dagger)}{q(z)} dz \right)'' d\theta \tag{94}
\end{aligned}$$

$$\begin{aligned}
&\leq \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \\
&\quad + C_{17} \int q(\theta) \frac{1}{2} |\theta - \mu|^2 q_h(\theta)^{-C_{14}} d\theta \tag{95}
\end{aligned}$$

$$\begin{aligned}
&= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz \\
&\quad + C_{17} \int q_h(h) \frac{1}{2} |\delta_n h|^2 q_h(\mu + \delta_n h)^{-C_{14}} dh \tag{96}
\end{aligned}$$

$$\leq \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz + C_{17} \max_i (\delta_{n,ii}^2) \tag{97}$$

$$= \int q(\theta) \log p(\theta) d\theta - \int q(\theta) \log q(\theta) d\theta + \sup_{q(z)} \int q(z) \log \frac{p(z, x|\mu)}{q(z)} dz + o(1) \tag{98}$$

for some constant $C_{17} > 0$. The first equality is by the definition of KL divergence. The second equality is rewriting the integrand. The third equality is due to mean value theorem where $\tilde{\theta}^\dagger$ is some value between μ and θ . (A very similar argument for q_h with nonzero means can be made starting from here, that is expanding only to the first order term.) The fourth equality is due to $q_h(\cdot)$ having zero mean. The fifth inequality is due to the third part of Assumption 3. The sixth inequality is due to a change of variable $h = \delta_n(\theta - \mu)$. The seventh inequality is due to $q_h(\cdot)$ residing within exponential family with finite second moment (the first part of Assumption 3). The eighth equality is due to $\delta_n \rightarrow 0$ as $n \rightarrow \infty$.

Combining the above two approximation, we have $-\text{KL}(q(\theta) \parallel \pi^*(\theta|x)) \leq \text{ELBO}_p(q(\theta)) + o(1)$. On the other hand, we know that $\text{ELBO}_p(q(\theta)) \leq -\text{KL}(q(\theta) \parallel \pi^*(\theta|x))$ by definition. We thus conclude $\text{ELBO}_p(q(\theta)) = -\text{KL}(q(\theta) \parallel \pi^*(\theta|x)) + o_p(1)$.

E Proof of Theorem 5 and Theorem 6

Theorem 5 is a direct consequence of Lemma 2, Lemma 3, and Lemma 4. Lemma 2 and Lemma 3 characterizes the consistency and asymptotic normality of the KL minimizer of the VB ideal. Lemma 4 says the VB posterior shares the same asymptotic properties as the VB ideal. All of them together give the consistency and asymptotic normality of VB posteriors.

Theorem 6 is a consequence of a slight generalization of Theorem 2.3 of [Kleijn et al. \(2012\)](#). The theorem characterizes the consistency and asymptotic normality of the posterior mean estimate under model specification with a common \sqrt{n} -convergence rate. We only need to replace all \sqrt{n} by δ_n^{-1} in their proof to obtain the generalization.

Specifically, we first show that, for any $M_n \rightarrow \infty$,

$$\int_{\|\tilde{\theta}\| > M_n} \|\tilde{\theta}\|^2 q^*(\tilde{\theta}) d\tilde{\theta} \xrightarrow{P_{\theta_0}} 0.$$

This is ensured by Assumption 3.1.

We then consider three stochastic processes: fix some compact set K and for given $M > 0$,

$$t \mapsto Z_{n,M}(t) = \int_{\|\tilde{\theta}\| \leq M} (t - \tilde{\theta})^2 \cdot q_{\tilde{\theta}}^*(\tilde{\theta}) d\tilde{\theta}, \quad (99)$$

$$t \mapsto W_{n,M}(t) = \int_{\|\tilde{\theta}\| \leq M} (t - \tilde{\theta})^2 \cdot \mathcal{N}(\tilde{\theta}; \Delta_{n,\theta_0}, V_{\theta_0}^{-1}) d\tilde{\theta}, \quad (100)$$

$$t \mapsto W_M(t) = \int_{\|\tilde{\theta}\| \leq M} (t - \tilde{\theta})^2 \cdot \mathcal{N}(\tilde{\theta}; X, V_{\theta_0}^{-1}) d\tilde{\theta}. \quad (101)$$

We note that θ_n^* is the minimizer of $t \mapsto Z_{n,\infty}(t)$ and $\int \tilde{\theta} \cdot \arg \min_{q \in \mathcal{Q}^d} \text{KL}(q(\tilde{\theta}) \parallel \mathcal{N}(\tilde{\theta}; X, V_{\theta_0}^{-1})) d\tilde{\theta}$ is the minimizer of $t \mapsto W_\infty(t)$.

By $\sup_{t \in K, \|h\| \leq M} (t - h)^2 < \infty$, we have $Z_{n,M} - W_{n,M} = o_{P_{\theta_0}}(1)$ in $\ell^\infty(K)$ by Theorem 5. Since $\Delta_{n,\theta_0} \xrightarrow{d} X$, the continuous mapping theorem implies that $W_{n,M} - W_M = o_{P_{\theta_0}}(1)$ in $\ell^\infty(K)$. By $\int \theta \cdot q^*(\theta) < \infty$, we have $W_M - W_\infty = o_{P_{\theta_0}}(1)$ as $M \rightarrow \infty$. We conclude that there exists a sequence $M_n \rightarrow \infty$ such that $Z_{n,M_n} - W_\infty = o_{P_{\theta_0}}(1)$. We also have from above that $Z_{n,M_n} - Z_{n,\infty} = o_{P_{\theta_0}}(1)$ in $\ell^\infty(K)$. We conclude that $Z_{n,\infty} - W_\infty = o_{P_{\theta_0}}(1)$ in $\ell^\infty(K)$. By the continuity and convexity of the squared loss, we invoke the argmax theorem and conclude that θ_n^* converges weakly to $\int \tilde{\theta} \cdot \arg \min_{q \in \mathcal{Q}^d} \text{KL}(q(\tilde{\theta}) \parallel \mathcal{N}(\tilde{\theta}; X, V_{\theta_0}^{-1})) d\tilde{\theta}$.

F Proof of Corollary 7

We prove Lemma 3 for the Gaussian family. We want to show that the functionals

$$F_n(m, \Sigma) := \text{KL}(\mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \parallel \pi^*(\theta | x))$$

Γ -converge to

$$F_0(m, \Sigma) := \text{KL}(\mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \parallel \mathcal{N}(\theta; \theta_0 + \delta_n \Delta_{n,\theta_0}, \delta_n V_{\theta_0}^{-1} \delta_n)) - \Delta_{n,\theta_0}^\top V_{\theta_0} \Delta_{n,\theta_0}$$

in probability as $n \rightarrow 0$. We note that this is equivalent to the Γ -convergence of the functionals of q : $\text{KL}(q(\cdot) \parallel \pi_{\tilde{\theta}}^*(\cdot \mid x))$ to $\text{KL}(q(\cdot) \parallel \mathcal{N}(\cdot; \Delta_{n, \theta_0}, V_{\theta_0}^{-1}))$. This is because the second statement is the same as the first up to a change of variable step from θ to $\tilde{\theta}$.

Together with Theorem 15, this implies $m_n, \Sigma_n \xrightarrow{P_{\theta_0}} m_0, \Sigma_0$ where (m_n, Σ_n) is the minimizer of F_n for each n and (m_0, Σ_0) is the minimizer of F_0 . This is due to the last term of $F_0 - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}$ is a constant bounded in P_{θ_0} probability and independent of m, Σ . The convergence in total variation then follows from Lemma 4.9 of Klartag (2007), which gives an upper bound on TV distance between two Gaussian distributions.

Now we prove the Γ -convergence.

We first rewrite $F_n(m, \Sigma)$.

$$F_n(m, \Sigma) \tag{102}$$

$$:= \text{KL}(\mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \parallel \pi^*(\theta \mid x)) \tag{103}$$

$$= \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \log \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) d\theta - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \log \pi^*(\theta \mid x) d\theta \tag{104}$$

$$= -\frac{d}{2} - \frac{d}{2} \log(2\pi) - \log \det(\delta_n) - \frac{1}{2} \log \det(\Sigma) - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \log p(\theta) d\theta \\ - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) M_n(\theta; x) d\theta + \log \int p(\theta) \exp(M_n(\theta; x)) d\theta \tag{105}$$

$$= -\frac{d}{2} - \frac{d}{2} \log(2\pi) - \log \det(\delta_n) - \frac{1}{2} \log \det(\Sigma) - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \log p(\theta) d\theta \\ - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) M_n(\theta; x) d\theta \\ + \left[\frac{d}{2} \log(2\pi) - \frac{1}{2} \log \det V_{\theta_0} + \log \det(\delta_n) + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \right] \tag{106}$$

$$= -\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \log p(\theta) d\theta \\ - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) M_n(\theta; x) d\theta + \left[-\frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + \log p(\theta_0) + o_P(1) \right] \tag{107}$$

$$= -\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) M_n(\theta; x) d\theta \\ + \left[-\frac{1}{2} \log \det V_{\theta_0} + M_n(\theta_0; x) + o_P(1) \right] \tag{108}$$

$$= -\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) - \int \mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) [M_n(\theta_0; x) + \delta_n^{-1}(\theta - \theta_0)^\top V_{\theta_0} \Delta_{n, \theta_0} \\ - \frac{1}{2}(\delta_n^{-1}(\theta - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(\theta - \theta_0) + o_P(1)] d\theta - \left[\frac{1}{2} \log \det V_{\theta_0} - M_n(\theta_0; x) + o_P(1) \right] \tag{109}$$

$$= -\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) - \delta_n^{-1}(m - \theta_0)^\top V_{\theta_0} \Delta_{n, \theta_0} + \frac{1}{2}(\delta_n^{-1}(m - \theta_0))^\top V_{\theta_0} \delta_n^{-1}(m - \theta_0) \\ + \frac{1}{2} \text{Tr}(V_{\theta_0} \cdot \Sigma) - \frac{1}{2} \log \det V_{\theta_0} + o_P(1). \tag{110}$$

The first equality is by the definition of KL divergence. The second equality is calculating the entropy of multivariate Gaussian distribution. The third equality is due to the Laplace approximation of the normalizer like we did in Equation (33). The fourth equality is due to the cancellation of the two $\log \det(\delta_n)$ terms. This again exemplifies why we assume the family to have variance $\delta_n \Sigma \delta_n$. The fifth equality is due to a similar argument as in Equation (43). The intuition is that $\mathcal{N}(m, \delta_n \Sigma \delta_n)$ converges to a point mass as $n \rightarrow \infty$. The sixth equality is due to the LAN condition of $M_n(\theta; x)$. The seventh equality is due to the computation of each term in the integral as an expectation under the Gaussian distribution $\mathcal{N}(m, \delta_n \Sigma \delta_n)$. To extend the restriction to some compact set K to the whole space \mathbb{R}^d in the sixth equality, we employ the same argument as in Equation (44).

We notice that when $m \neq \theta_0$, we will have $F_n(m, \Sigma) \rightarrow \infty$. On the other hand, we have $\limsup F_n < \infty$. This echoes our consistency result in Lemma 2.

Now we rewrite $F_0(m, \Sigma)$.

$$\text{KL}(\mathcal{N}(\theta; m, \delta_n \Sigma \delta_n) \| \mathcal{N}(\theta; \theta_0 + \delta_n \Delta_{n, \theta_0}, \delta_n V_{\theta_0}^{-1} \delta_n)) \quad (111)$$

$$= -\frac{d}{2} + \frac{1}{2} \text{Tr}(V_{\theta_0} \cdot \Sigma) + (m - \theta_0 - \delta_n \Delta_{n, \theta_0})^\top \delta_n^{-1} V_{\theta_0} \delta_n^{-1} (m - \theta_0 - \delta_n \Delta_{n, \theta_0}) \\ - \frac{1}{2} \log \det(\Sigma) - \frac{1}{2} \log \det V_{\theta_0} \quad (112)$$

$$= -\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) - \delta_n^{-1} (m - \theta_0)^\top V_{\theta_0} \Delta_{n, \theta_0} + \frac{1}{2} (\delta_n^{-1} (m - \theta_0))^\top V_{\theta_0} \delta_n^{-1} (m - \theta_0) \\ + \frac{1}{2} \text{Tr}(V_{\theta_0} \cdot \Sigma) - \frac{1}{2} \log \det V_{\theta_0} + \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}. \quad (113)$$

This gives

$$F_0(m, \Sigma) - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0} \quad (114)$$

$$= -\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) - \delta_n^{-1} (m - \theta_0)^\top V_{\theta_0} \Delta_{n, \theta_0} + \frac{1}{2} (\delta_n^{-1} (m - \theta_0))^\top V_{\theta_0} \delta_n^{-1} (m - \theta_0) \\ + \frac{1}{2} \text{Tr}(V_{\theta_0} \cdot \Sigma) - \frac{1}{2} \log \det V_{\theta_0} \quad (115)$$

$$= +\infty \cdot (1 - \mathbb{1}_m(\theta_0)) + [-\frac{d}{2} - \frac{1}{2} \log \det(\Sigma) + \frac{1}{2} \text{Tr}(V_{\theta_0} \cdot \Sigma) - \frac{1}{2} \log \det V_{\theta_0}] \cdot \mathbb{1}_m(\theta_0). \quad (116)$$

This equality is due to the KL divergence between two multivariate Gaussian distributions.

Comparing Equation (110) and Equation (116), we can prove the Γ convergence. Let $m_n \rightarrow m$ and $\Sigma_n \rightarrow \Sigma$. When $m \neq \theta_0$, $\liminf_{n \rightarrow \infty} F_n(m_n, \Sigma_n) = +\infty$. The limsup inequality is automatically satisfied. When $m = \theta_0$, we have $F_n(m, \Sigma) = F_0(m, \Sigma) - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0} + o_P(1)$. This implies $F_0(m, \Sigma) \leq \lim_{n \rightarrow \infty} F_n(m_n, \Sigma_n)$ in P_{θ_0} probability by the continuity of F_n .

We then show the existence of a recovery sequence. When $m \neq \theta_0$, $F_0(m, \Sigma) = +\infty$. The limsup inequality is automatically satisfied. When $m = \theta_0$, we can simply choose $\Sigma_m = \Sigma$ and $m_n = \theta_0$. The limsup inequality is again ensured by $F_0(m, \Sigma) \leq \lim_{n \rightarrow \infty} F_n(m_n, \Sigma_n)$ in P_{θ_0} probability and the continuity of F_n . The Γ -convergence of the F functionals is shown.

We notice that $\Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}$ does not depend on any of m, Σ so that $\text{argmin} F_0 = \text{argmin} F_0 - \Delta_{n, \theta_0}^\top V_{\theta_0} \Delta_{n, \theta_0}$. The convergence of the KL minimizers is thus proved.

G Proof of Lemma 8

We first note that

$$\text{KL}(\mathcal{N}(\cdot; \mu_0, \Sigma_0) \| \mathcal{N}(\cdot; \mu_1, \Sigma_1)) \quad (117)$$

$$= \frac{1}{2} [\text{Tr}(\Sigma_1^{-1} \Sigma_0) + (\mu_1 - \mu_0)^\top \Sigma_1^{-1} (\mu_1 - \mu_0) - d + \log \frac{\det(\Sigma_1)}{\det(\Sigma_0)}]. \quad (118)$$

Clearly, the optimal choice of μ_0 is $\hat{\mu}_0 = \mu_1$. Next, we write $\Sigma_0 = \text{diag}(\lambda_1, \dots, \lambda_d)$. The KL divergence minimization objective thus becomes

$$\frac{1}{2} [\sum_{i=1}^d (\Sigma_1^{-1})_{ii} \lambda_i + \log \det(\Sigma_1) - \sum_{i=1}^d \log \lambda_i].$$

Taking its derivative with respect to each λ_i and setting it to zero, we have

$$(\Sigma_1^{-1})_{ii} = \lambda_i^{-1}.$$

The optimal Σ_0 thus should be diagonal with $\hat{\Sigma}_{0,ii} = ((\Sigma_1^{-1})_{ii})^{-1}$ for $i = 1, 2, \dots, d$. In this sense, mean field (factorizable) approximation matches the precision matrix at the mode.

Moreover, by the inequality (Amir-Moez and Johnston, 1969; Beckenbach and Bellman, 2012)

$$\det(\Sigma_1^{-1}) \leq \prod_i (\Sigma_1^{-1})_{ii} = \det(\hat{\Sigma}_0^{-1}),$$

we have

$$\mathbb{H}(\hat{\Sigma}_0^{-1}) = \frac{1}{2} \log((2\pi e)^d \cdot \det(\hat{\Sigma}_0)) \leq \frac{1}{2} \log((2\pi e)^d \cdot \det(\Sigma_1)) = \mathbb{H}(\Sigma_1).$$

H Proof of Corollary 10

We only need to verify the local asymptotic normality of $L_n(\mu; x)$ here. By Equation (2) of [Westling and McCormick \(2015\)](#), we know the variational log likelihood writes $L_n(\mu; x) = \sum_i m(\mu; x_i)$. We Taylor-expand it around the true value μ_0 :

$$L_n(\mu_0 + \frac{s}{\sqrt{n}}; x) \tag{119}$$

$$= \sum_i m(\mu_0 + \frac{s}{\sqrt{n}}; x_i) \tag{120}$$

$$= \sum_i m(\mu_0; x_i) + \frac{s}{\sqrt{n}} \sum_i D_\mu m(\mu_0; x_i) + \frac{1}{n} s^\top [\sum_i D_\mu^2 m(\mu_0; x_i)] s. \tag{121}$$

Due to X_i 's being independent and identically distributed, we have

$$\frac{1}{\sqrt{n}} \sum_i D_\mu m(\mu_0; x_i) = \sqrt{n} \cdot \frac{\sum_i D_\mu m(\mu_0; x_i)}{n} \xrightarrow{d} \mathcal{N}(0, B(\mu))$$

under P_{μ_0} , where $B(\mu) = \mathbb{E}_{P_{\mu_0}} [D_\mu m(\theta; x) D_\mu m(\mu; x)^\top]$. The convergence in distribution is due to central limit theorem. The mean zero here is due to conditions B2, B4, and B5 of [Westling and McCormick \(2015\)](#) (See point 4 in the first paragraph of Proof of Theorem 2 in [Westling and McCormick \(2015\)](#) for details.)

By strong law of large numbers, we also have

$$\frac{1}{n} \sum_i D_\mu^2 m(\mu_0; x_i) \xrightarrow{P_{\theta_0}} \mathbb{E}_{P_{\theta_0}} [D_\mu^2 m(\mu_0; x_i)].$$

This gives the local asymptotic normality, for s in a compact set,

$$L_n(\mu_0 + \frac{s}{\sqrt{n}}; x) = L_n(\mu_0; x) + s^\top \Sigma Y - \frac{1}{2} s^\top \Sigma s + o_P(1),$$

where

$$Y \sim \mathcal{N}(0, A(\mu_0)^{-1} B(\mu_0) A(\mu_0)^{-1}),$$

and

$$\Sigma = A(\mu_0) = \mathbb{E}_{P_{\theta_0}} [D_\mu^2 m(\mu_0; x_i)]$$

is a positive definite matrix.

The consistent testability assumption is satisfied by the existence of consistent estimators. This is due to Theorem 1 of [Westling and McCormick \(2015\)](#).

The corollary then follows from Theorem 5 and Theorem 6 in Appendix E.

I Proof of Corollary 11

We first verify the local asymptotic normality of the variational log likelihood:

$$\underline{\ell}(\beta, \sigma^2) := \sup_{\mu, \lambda} \ell(\beta, \sigma^2, \mu, \lambda) \tag{122}$$

$$\begin{aligned} &= \sup_{\mu, \lambda} \sum_{i=1}^m \sum_{j=1}^n \{Y_{ij}(\beta_0 + \beta_1 X_{ij} + \mu_i) - \exp(\beta_0 + \beta_1 X_{ij} + \mu_i + \lambda_i/2) - \log(Y_{ij}!)\} \\ &\quad - \frac{m}{2} \log(\sigma^2) + \frac{m}{2} - \frac{1}{2\sigma^2} \sum_{i=1}^n (\mu_i^2 + \lambda_i) + \frac{1}{2} \sum_{i=1}^m \log(\lambda_i). \end{aligned} \tag{123}$$

We take the Taylor expansion of the variational log likelihood at the true parameter values $\beta_0^0, \beta_1^0, (\sigma^2)^0$:

$$\underline{\ell}(\beta_0^0 + \frac{s}{\sqrt{m}}, \beta_1^0 + \frac{t}{\sqrt{mn}}, (\sigma^2)^0 + \frac{r}{\sqrt{m}}) \tag{124}$$

$$\begin{aligned}
&= \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) + \frac{s}{\sqrt{m}} \frac{\partial}{\partial \beta_0} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) + \frac{t}{\sqrt{mn}} \frac{\partial}{\partial \beta_1} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) \\
&\quad + \frac{r}{\sqrt{m}} \frac{\partial}{\partial \sigma^2} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) + \frac{s^2}{m} \frac{\partial^2}{\partial \beta_0^2} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) + \frac{t^2}{mn} \frac{\partial^2}{\partial \beta_1^2} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) \\
&\quad + \frac{r^2}{m} \frac{\partial^2}{\partial (\sigma^2)^2} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) + o_P(1). \tag{125}
\end{aligned}$$

Denote $\hat{\mu}_i$ and $\hat{\lambda}_i$ as the optimal μ and λ at the true values $(\beta_0^0, \beta_1^0, (\sigma^2)^0)$. Also write $Y_{i\cdot} = \sum_{j=1}^n Y_{ij}$ and $B_i = \sum_{j=1}^n \exp(\beta_0 + \beta_1 X_{ij})$. Let $\hat{\beta}_1, \hat{\beta}_0, \hat{\sigma}^2$ be the maximizers of the $\underline{\ell}$. Hence, we write $\hat{B}_i = \sum_{j=1}^n \exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij})$. Finally, the moment generating function of X writes $\phi(t) = \mathbb{E}\{\exp(tX)\}$.

The cross terms are zero due to Equation (5.21), Equation (5.29), Equation (5.37), and Equation (5.50) of Hall et al. (2011).

We next compute each of the six derivatives terms.

For the first term $\frac{\partial}{\partial \beta_0} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0)$, we have

$$\frac{1}{\sqrt{m}} \frac{\partial}{\partial \beta_0} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) \tag{126}$$

$$= \frac{1}{\sqrt{m}} \sum_{i=1}^m [Y_{i\cdot} - B_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i)] \tag{127}$$

$$\begin{aligned}
&= \frac{1}{\sqrt{m}} \sum_{i=1}^m [Y_{i\cdot} - \hat{B}_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i)] - \frac{1}{\sqrt{m}} (\beta_0 - \hat{\beta}_0) \sum_{i=1}^m \hat{B}_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) \\
&\quad - \frac{1}{\sqrt{m}} (\beta_1 - \hat{\beta}_1) \sum_{i=1}^m \hat{B}_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) + o_P(1) \tag{128}
\end{aligned}$$

$$= \sqrt{m} \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^n ((\beta_0 - \hat{\beta}_0) + (\beta_1 - \hat{\beta}_1) X_{ij}) \exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij}) \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) + o_P(1) \tag{129}$$

$$= \sqrt{m} \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^n ((\beta_1 - \hat{\beta}_1)(X_{ij} - \gamma(\beta_1^0)) + \bar{U}) \exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij}) \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) + o_P(1) \tag{130}$$

$$= \sqrt{m} \bar{U} \frac{1}{m} \sum_{i=1}^m \hat{B}_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) + o_P(1) \tag{131}$$

$$= \sqrt{m} \bar{U} \frac{1}{m} \sum_{i=1}^m \hat{B}_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) + o_P(1) \tag{132}$$

$$\stackrel{d}{\rightarrow} \mathcal{N}(0, (\sigma^2)^0) [\exp(\beta_0^0) \mathbb{E} \exp(\beta_1 X_{ij} + U_i)] + o_P(1) \tag{133}$$

$$= \mathcal{N}(0, (\sigma^2)^0) [\exp(\beta_0^0 - \frac{1}{2} (\sigma^2)^0) \phi(\beta_1^0)]. \tag{134}$$

The first equality is due to differentiation with respect to β_0 . The second equality is due to Taylor expansion around the VFE. The third equality is due to Equation (3.5) of Hall et al. (2011). The fourth equation is due to Equation (5.21) of Hall et al. (2011). The fifth equation is due to Equation (5.1) of Hall et al. (2011). The sixth equation is due to the weak law of law numbers and Slutsky's theorem together with Equation (3.4), Equation (5.16), and Equation (5.18) of Hall et al. (2011). The seventh equality is due to the equation below Equation (5.80) of Hall et al. (2011).

We then compute the fourth term.

$$\frac{1}{m} \frac{\partial^2}{\partial \beta_0^2} \underline{\ell}(\beta_0^0, \beta_1^0, (\sigma^2)^0) \tag{135}$$

$$= \frac{1}{m} \sum_{i=1}^m \hat{B}_i \exp(\hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i) + o_P(1) \tag{136}$$

$$\stackrel{P}{\rightarrow} \exp(\beta_0 - \frac{1}{2} (\sigma^2)^0) \phi(\beta_1^0). \tag{137}$$

This step is due to a similar computation to above. The last step is due to weak law of large numbers and the equation below Equation (5.80) of Hall et al. (2011).

We now compute the third term.

$$\frac{1}{\sqrt{m}} \frac{\partial}{\partial \sigma^2} \ell(\beta_0^0, \beta_1^0, (\sigma^2)^0) \quad (138)$$

$$= \frac{1}{\sqrt{m}} \left[-\frac{m}{2(\sigma^2)^0} + \frac{1}{2\{(\sigma^2)^0\}^2} \sum_{i=1}^m (\hat{\mu}_i^2 + \hat{\lambda}_i) \right] \quad (139)$$

$$= -\sqrt{m} \left(\frac{1}{2(\sigma^2)^0} \right) \left(1 - \frac{\hat{\sigma}^2}{(\sigma^2)^0} \right) \quad (140)$$

$$= \frac{1}{2\{(\sigma^2)^0\}^2} \sqrt{m} (\hat{\sigma}^2 - (\sigma^2)^0) \quad (141)$$

$$\xrightarrow{d} \mathcal{N}(0, 2\{(\sigma^2)^0\}^2) \frac{1}{2\{(\sigma^2)^0\}^2}. \quad (142)$$

The first equality is due to differentiation with respect to σ^2 . The second equality is due to Equation (5.3) of Hall et al. (2011). The third equality is rearranging the terms. The fourth equation is due to Equation (3.6) of Hall et al. (2011).

We then compute the sixth term.

$$\frac{1}{m} \frac{\partial^2}{\partial (\sigma^2)^2} \ell(\beta_0^0, \beta_1^0, (\sigma^2)^0) \quad (143)$$

$$= \left[-\frac{1}{2} \frac{1}{\{(\sigma^2)^0\}^2} + \frac{1}{\{(\sigma^2)^0\}^3} \frac{1}{m} \sum_{i=1}^m (\hat{\mu}_i^2 + \hat{\lambda}_i) \right] \quad (144)$$

$$= \left[-\frac{1}{2} \frac{1}{\{(\sigma^2)^0\}^2} + \frac{\hat{\sigma}^2}{\{(\sigma^2)^0\}^3} \right] + o_P(1) \quad (145)$$

$$\xrightarrow{P} \frac{1}{2\{(\sigma^2)^0\}^2}. \quad (146)$$

This is due to a similar computation to above. The last step is due to Equation (3.6) of Hall et al. (2011) and the weak law of large numbers.

We next compute the second term.

$$\frac{1}{\sqrt{mn}} \frac{\partial}{\partial \beta_1} \ell(\beta_0^0, \beta_1^0, (\sigma^2)^0) \quad (147)$$

$$= \frac{1}{\sqrt{mn}} \sum_{i=1}^m \sum_{j=1}^n X_{ij} (Y_{ij} - \exp(\beta_0^0 + \beta_1^0 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i)) \quad (148)$$

$$= \frac{1}{\sqrt{mn}} \sum_{i=1}^m \sum_{j=1}^n X_{ij} (Y_{ij} - \exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i))$$

$$+ (\beta_0^0 - \hat{\beta}_0) \frac{1}{\sqrt{mn}} \sum_{i=1}^m \sum_{j=1}^n X_{ij} (-\exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i))$$

$$+ (\beta_1^0 - \hat{\beta}_1) \frac{1}{\sqrt{mn}} \sum_{i=1}^m \sum_{j=1}^n X_{ij}^2 (-\exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i)) + o_P(1) \quad (149)$$

$$= \sqrt{mn} \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n ((\beta_0^0 - \hat{\beta}_0) + (\beta_1^0 - \hat{\beta}_1) X_{ij}) X_{ij} (-\exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i))$$

$$+ o_P(1) \quad (150)$$

$$= \sqrt{mn} \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n ((\beta_1^0 - \hat{\beta}_1) (X_{ij} - \gamma(\beta_1^0)) + \bar{U}) X_{ij} (-\exp(\hat{\beta}_0 + \hat{\beta}_1 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i))$$

$$+ o_P(1) \quad (151)$$

$$\xrightarrow{d} \mathcal{N}(0, \tau^2) \exp(\beta_0^0 - \frac{1}{2} \sigma^2) \phi''(\beta_1^0). \quad (152)$$

The first equality is due to differentiation with respect to β_1 . The second equality is due to Taylor expansion around VFE. The third equality is due to Equation (3.4) of Hall et al. (2011). The fourth equality is due to Equation (5.16), Equation (5.18), and Equation (5.21) of Hall et al. (2011). The fifth equation is due to the weak law of law numbers and Slutsky's theorem together with Equation (3.5) and the equation below Equation (5.80) of Hall et al. (2011).

We lastly compute the fifth term.

$$\frac{1}{mn} \frac{\partial^2}{\partial \beta_1^2} \ell(\beta_0^0, \beta_1^0, (\sigma^2)^0) \quad (153)$$

$$= \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n X_{ij}^2 (\exp(\beta_0^0 + \beta_1^0 X_{ij} + \hat{\mu}_i + \frac{1}{2} \hat{\lambda}_i)) \quad (154)$$

$$= \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n X_{ij}^2 (\exp(\beta_0^0 + \beta_1^0 X_{ij} + U_i)) \quad (155)$$

$$\xrightarrow{P} \exp(\beta_0^0 - \frac{1}{2} \sigma^2) \phi''(\beta_1^0). \quad (156)$$

This is due to a similar computation to above. The last step is due to the weak law of large numbers.

The calculation above gives the full local asymptotic expansion of $\ell(\beta_0, \beta_1, \sigma^2)$.

The consistent testability assumption is satisfied by the existence of consistent estimators. This is due to Theorem 3.1 of Hall et al. (2011). The corollary then follows directly from Theorem 5 and Theorem 6 in Appendix E.

References

- Amir-Moez, A. and Johnston, G. (1969). On the product of diagonal elements of a positive matrix. *Mathematics Magazine*, 42(1):24–26.
- Beckenbach, E. F. and Bellman, R. (2012). *Inequalities*, volume 30. Springer Science & Business Media.
- Braides, A. (2006). A handbook of Γ -convergence. *Handbook of Differential Equations: stationary partial differential equations*, 3:101–213.
- Dal Maso, G. (2012). *An introduction to Γ -convergence*, volume 8. Springer Science & Business Media.
- Hall, P., Pham, T., Wand, M. P., Wang, S. S., et al. (2011). Asymptotic normality and valid inference for Gaussian variational approximation. *The Annals of Statistics*, 39(5):2502–2532.
- Klartag, B. (2007). A central limit theorem for convex sets. *Inventiones mathematicae*, 168(1):91–131.
- Kleijn, B., Van der Vaart, A., et al. (2012). The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381.
- Lu, Y., Stuart, A. M., and Weber, H. (2016). Gaussian approximations for probability measures on \mathbf{R}^d . *arXiv preprint arXiv:1611.08642*.
- Piera, F. J. and Parada, P. (2009). On convergence properties of Shannon entropy. *Problems of Information Transmission*, 45(2):75–94.
- Westling, T. and McCormick, T. H. (2015). Establishing consistency and improving uncertainty estimates of variational inference through M-estimation. *arXiv preprint arXiv:1510.08151*.