Published in final edited form as:

J Comput Graph Stat. 2019 ; 28(4): 778–789. doi:10.1080/10618600.2019.1609977.

Beyond prediction: A framework for inference with variational approximations in mixture models

T. Westling,

Center for Causal Inference, University of Pennsylvania

T. H. McCormick

Departments of Statistics & Sociology, University of Washington

Abstract

Variational inference is a popular method for estimating model parameters and conditional distributions in hierarchical and mixed models, which arise frequently in many settings in the health, social, and biological sciences. Variational inference in a frequentist context works by approximating intractable conditional distributions with a tractable family and optimizing the resulting lower bound on the log-likelihood. The variational objective function is typically less computationally intensive to optimize than the true likelihood, enabling scientists to fit rich models even with extremely large datasets. Despite widespread use, little is known about the general theoretical properties of estimators arising from variational approximations to the log-likelihood, which hinders their use in inferential statistics. In this paper we connect such estimators to profile M -estimation, which enables us to provide regularity conditions for consistency and asymptotic normality of variational estimators. Our theory also motivates three methodological improvements to variational inference: estimation of the asymptotic model-robust covariance matrix, a one-step correction that improves estimator efficiency, and an empirical assessment of consistency. We evaluate the proposed results using simulation studies and data on marijuana use from the National Longitudinal Study of Youth.

Keywords

generalized linear mixed models; profile M -estimation

1 Introduction

Thanks to rapid improvements in data availability and user-friendly tools for data storage and manipulation, researchers from an ever-broader set of scientific disciplines now routinely analyze extremely complex, high-dimensional data. However, computing parameter estimates using stalwart statistical techniques such as maximum likelihood and Markov chain Monte Carlo can be a challenge in these settings and is often a bottleneck in

Corresponding author T. Westling tgwest@pennmedicine.upenn.edu.

Supplementary material

Supplementary material includes proofs of Theorems 1 and 2, an intuitive explanation of the over-concentration of the variational Bayes posterior distribution, derivations related to the exponential mixture model, and additional simulation results.

practice. In these situations, researchers often turn to computationally efficient approximations.

Variational approximations are one method of approximating a likelihood function or posterior distribution that are increasingly popular across a range of scientific fields. In public health, for example, Lee and Wand (2016) used a variational approximation to estimate a model for overall and hospital-specific trends in cesarean section rates. In statistical genetics, Raj et al. (2014) used a variational approximation to a multinomial model of allele frequencies across populations of individuals. O'Connor et al. (2010) used a variational approximation to a model of demographics and lexical choice in geo-tagged Twitter data.

Despite their popularity, variational approximations do not typically come with guarantees about the statistical properties of the resulting estimator. This drawback is particularly problematic when a scientist would like to interpret a parameter estimate, in which case estimator consistency is crucial, or report a confidence interval, in which case good coverage rates rely on the ability to accurately estimate the sampling distribution of the estimator. In this paper, we address the problem of inference using variational approximations. We show that, in a wide range of parametric mixture models, well-established theory from profile M -estimation provides an asymptotic lens through which we may understand the large-sample properties of parameter estimates resulting from variational approximations to the log-likelihood. Using the M -estimation framework, we derive conditions for consistency and asymptotic normality of variational estimators.

The theory we establish for variational estimators motivates us to also propose three methodological improvements to these estimators. First, we provide a consistent estimator of the asymptotic covariance matrix of variational estimators. Second, we introduce a one-step correction to the variational estimator that improves large-sample statistical efficiency. Third, we develop an empirical evaluation of estimator consistency for use when the theoretical calculations are intractable. We demonstrate the importance of these methodological advances with two logistic mixture models of marijuana use by age among participants in the National Longitudinal Survey of Youth (NLSY).

The remainder of the paper is organized as follows. This section formally defines the class of models and variational estimators we study. Section 2 connects variational estimation to profile M -estimation and states our theoretical results. Section 3 illustrates the general theoretical results in a few simple models. Section 4 presents our three methodological contributions. Section 5 evaluates our methods using simulated data and demonstrates an application to the NLSY. Section 6 presents a discussion. Technical conditions and proofs of theorems are provided in the supplementary material. Code to replicate all of the empirical analyses in this paper are available at https://github.com/tedwestling/variational_asymptotics.

1.1 Variational estimators

In this paper, we consider inference for a Euclidean parameter θ in a parametric mixture model $p_{\theta}(x) = \int_z p_{\theta}(x, z) d\mu(z)$, where the marginal likelihood $p_{\theta}(x)$ is computationally

expensive to compute. Parametric mixture models have been used in a variety of scientific contexts. For example, mixed-membership models are a type of mixture model that have been used to model text (Blei et al. 2003), social networks (Airoldi et al. 2008), population genetics (Pritchard et al. 2000), and scientific collaborations (Erosheva et al. 2004).

Mixture models have been used in conjunction with both Bayesian and frequentist inferential frameworks. In a frequentist setting, maximum likelihood (ML) estimation comes with guarantees of asymptotic efficiency and methods of conducting inference for many models. These guarantees provide a degree of assurance for scientists that the point estimates and uncertainty intervals will behave in predictable ways.

ML estimation can, however, be computationally burdensome. When the integral in $p_{\theta}(x)$ must be approximated numerically, the cost of this computation increases exponentially with the dimension of the domain $\mathcal{Z}$ of the latent variable since $p_{\theta}(x, z)$ needs to be evaluated at sufficiently many points to accurately approximate the integral. This computational burden is a significant barrier for researchers who want to develop tailored mixture models to flexibly represent the dependencies in their data. As a result, a variety of approximate methods have been developed as alternatives to maximum likelihood.

Variational inference is an approximate method based on optimizing a lower bound for the original objective function. This lower bound is designed to eliminate the need for, or at least reduce the dimension of, any numerical integrals, thereby improving computational efficiency. Variational inference can be used in a frequentist context to approximate the log-likelihood or in a Bayesian context to approximate the posterior distribution. In this paper we focus on the former. We will refer to estimators of θ resulting from optimizing a variational approximation to the log-likelihood as *variational estimators*.

Before providing formal definitions, we distinguish between two key aspects of the variational approximation. First, we can evaluate the properties of the optimizer of the variational lower bound. Second, we could consider the tightness of the variational lower bound to the true objective function. These questions are related. Demonstrating tightness of the lower bound is one way to control the difference between the true and variational optimizers, for example. However, a tight variational lower bound is not a necessary condition for good behavior of the variational estimator, and indeed does not hold in many settings where the variational estimator performs well. In this paper, we address the former of these two components, i.e. the properties of the optimizer of the variational lower bound.

We now move to a formal definition of variational estimation. Blei et al. (2017) presents a thorough introduction to variational inference and many relevant references. Let $X_1, \dots, X_n$ be observed p -variate data generated independently and identically from a distribution P_0 on a sample space $\mathcal{X}$. Let $\mathcal{P} = \{P_{\theta} : \theta \in \Theta\}$ be a statistical model, where Θ is an open subset of $\mathbb{R}^d$ and each P_{θ} has a density $P_{\theta}(x) = \int_{\mathcal{Z}} P_{\theta}(x, z) d\mu$. Here μ is a dominating measure on $\mathcal{Z} \subseteq \mathbb{R}^k$. We can conceptualize this data-generating process as first drawing independent *latent* random variables $Z_1, \dots, Z_n$ from the marginal distribution $p_{\theta, z}(z) = \int_{\mathcal{X}} p_{\theta}(x, z) dx$, then drawing each X_i given Z_i from the conditional distribution $p_{\theta, x|z}(x | Z_i) = p_{\theta}(x, Z_i) / p_{\theta, z}(Z_i)$. Of these, we only observe $X_1, \dots, X_n$.

We are most interested in cases where $p_{\theta}(x)$ cannot be written in closed-form in terms of elementary functions, as in many generalized linear mixed models (McCulloch and Neuhaus 2001) and non-linear hierarchical models (Davidian and Giltinan 1995, Goldstein 2011). In these cases, calculating the log-likelihood of the observed data, $\sum_{i=1}^n \log p_{\theta}(X_i)$, and its derivatives with respect to θ requires numerical integration. When the dimension of the latent variable is large, these numerical integrals are computationally expensive.

Variational inference parameter estimates are obtained by maximizing a criterion function motivated as follows. Denote by $\mathcal{Q}_0$ the set of densities dominated by μ , and by $\mathcal{Q}_0^{\mathcal{X}}$ the set of all conditional densities dominated by μ for all $x \in \mathcal{X}$; that is, all $s: \mathcal{Z} \times \mathcal{X} \rightarrow \mathbb{R}$ such that $s(\cdot | x) \in \mathcal{Q}_0$ for all $x \in \mathcal{X}$. Suppose that $P_0 \in \mathcal{P}$, so that $p_0 = P_{\theta_0}$ for some $\theta_0 \in \Theta$. Then θ_0 and the true conditional distribution of the latent variable $\pi_{\theta_0}(z | x) := p_{\theta_0}(x, z) / p_{\theta_0}(x)$ can be represented as

$$(\theta_0, \pi_{\theta_0}) = \arg \max_{\theta \in \Theta, s \in \mathcal{Q}_0^{\mathcal{X}}} E_{P_0} \left[\int_{\mathcal{Z}} \log \left(\frac{p_{\theta}(X, Z)}{s(Z | X)} \right) s(Z | X) d\mu(Z) \right]. \quad (1)$$

To see this, first define

$$f_0(\theta, s) := E_{P_0} \left[-D_{KL}(s(\cdot | X) || \pi_{\theta}(\cdot | X)) \right] = E_{P_0} \left[\int_{\mathcal{Z}} \log \left(\frac{\pi_{\theta}(Z | X)}{s(Z | X)} \right) s(Z | X) d\mu(Z) \right],$$

where D_{KL} denotes the Kullback-Leibler (KL) divergence. Thus, $f_0(\theta, s)$ is the expected KL divergence between $s(\cdot | X)$ and $\pi_{\theta}(\cdot | X)$. By Gibbs' inequality, $f_0(\theta, s) \leq 0 = f_0(\theta_0, \pi_{\theta_0})$ for all $(\theta, s) \in \Theta \times \mathcal{Q}_0^{\mathcal{X}}$. Next, note that θ_0 maximizes

$$\theta \mapsto g_0(\theta) := E_{P_0} \left[\log \frac{p_{\theta}(X)}{p_0(X)} \right] = -D_{KL}(p_0 || p_{\theta}).$$

Therefore, $(\theta_0, \pi_{\theta_0})$ maximizes $(\theta, s) \mapsto$

$f_0(\theta, s) + g_0(\theta)$ over $\Theta \times \mathcal{Q}_0^{\mathcal{X}}$, and after some rearranging, we can see that this is equivalent to the representation in (1).

The expectation-maximization (EM) algorithm can be motivated by (1) by replacing the unknown P_0 with the empirical distribution and alternating between optimization over θ and s . Using similar reasoning to that presented above, this amounts to alternating between computing $\theta_{(t)} := \arg \max_{\theta \in \Theta} \sum_{i=1}^n \int_{\mathcal{Z}} [\log p_{\theta}(X_i, Z)] \pi_{\theta_{(t-1)}}(Z | X_i) d\mu(Z)$, where $\theta_{(t-1)}$ is the previous value of θ , and computing $\pi_{\theta_{(t)}}(\cdot | X_i)$ for each observed X_i . However, if the marginal likelihood $p_{\theta}(x)$ cannot be written in terms of elementary functions, then neither can π_{θ} , and hence the EM algorithm requires numerical integration.

To construct a variational approximation to the log-likelihood, we replace the optimization over $\mathcal{Q}_0^{\mathcal{X}}$ in (1) with an optimization over $\mathcal{Q}^{\mathcal{X}}$, where $\mathcal{Q}$ is a smaller *variational family* of distributions, and as before $\mathcal{Q}^{\mathcal{X}}$ is the set of conditional distributions over $\mathcal{Z}$ such that $s(\cdot | x) \in \mathcal{Q}$ for each $x \in \mathcal{X}$. For example, $\mathcal{Q}$ could consist of all independent products over

each dimension of z (known as mean-field variational inference), all multivariate Gaussian distributions, or all independent Gaussian distributions. For simplicity, we will assume throughout that $\mathcal{Q}$ is indexed by a finite-dimensional Euclidean parameter $\psi \in \Psi$, so that every $s \in \mathcal{Q}^{\mathcal{X}}$ can be identified with a density $s(\cdot | x) = q(\cdot; \psi(x))$. We note that in some cases even when $\mathcal{Q}$ is a semiparametric family, it can be shown that the optimal q lies in a parametric sub-family with a known form, so that our results can still be applied (see, e.g. Section 5.3 of Wainwright and Jordan 2008). For families where this does not apply, our theory could be extended to incorporate semiparametric $\mathcal{Q}$.

Let Ψ^n denote the n -fold Cartesian product $\Psi \times \dots \times \Psi$. For $\psi \in \Psi^n$ and $i \in \{1, \dots, n\}$, we will denote $\psi_i \in \Psi$ the i th element of ψ . Given the observed data $X_1, \dots, X_n$, Ψ^n then parametrizes the set of variational conditional distributions over $X_1, \dots, X_n$, and each ψ_i parametrizes the variational conditional distribution $s(\cdot | X_i) = q(\cdot; \psi_i)$. Given $\mathcal{Q}$ and Ψ , the variational estimator of θ , which we will denote $\hat{\theta}_n$, and the variational conditional estimators ψ_n are the joint maximizers of the following objective function:

$$(\hat{\theta}_n, \psi_n) := \arg \max_{\theta \in \Theta, \psi \in \Psi^n} \sum_{i=1}^n \int \log \left(\frac{p_{\theta}(X_i, Z_i)}{q(Z_i; \psi_i)} \right) q(Z_i; \psi_i) d\mu(Z_i) = \arg \max_{\theta \in \Theta, \psi \in \Psi^n} \mathcal{L}_n(\theta, \psi; \mathbf{X}_n). \quad (2)$$

We note that we are implicitly assuming that the full variational distribution over $(Z_1, \dots, Z_n)$ factors as $\prod_{i=1}^n q(Z_i; \psi_i)$. However, since the true conditional distribution of $(Z_1, \dots, Z_n)$ given $(X_1, \dots, X_n)$ factors as $\prod_{i=1}^n \pi_{\theta_0}(Z_i | X_i)$, the optimal variational distribution will always factor as well, so this assumption comes with no loss of generality.

A crucial piece of motivation for our work is that, since $\mathcal{L}_n$ is typically not proportional to the log-likelihood, it is not clear what the asymptotic properties of the variational estimator $\hat{\theta}_n$ are. In many circumstances, the variational estimator is used for prediction. In such cases, scientists can evaluate the quality of the variational approximation using cross-validation or another held-out data technique. If, however, a scientist would like to go beyond prediction and interpret the point estimator (or, critically, its uncertainty) produced by a variational approximation, not knowing the properties of the estimator is a substantial hindrance. In particular, we would like to know whether $\hat{\theta}_n$ is consistent and, if it is consistent, what the asymptotic distribution of $\sqrt{n}(\hat{\theta}_n - \theta_0)$ is.

Asymptotic properties of variational estimators have been studied in depth for certain specific models, yielding positive results regarding the consistency of variational estimators for Gaussian mixture models (Wang and Titterton 2006), exponential family models with missing values (Wang and Titterton 2004), Poisson mixed models as the cluster size and number of clusters both diverge (Hall, Ormerod and Wand 2011, Hall, Pham, Wand and Wang 2011), Markovian models with missing values (Hall et al. 2002), and stochastic block models for social networks (Bickel et al. 2013). Of particular note are Hall, Ormerod and Wand (2011) and Hall, Pham, Wand and Wang (2011), who derive sharp asymptotics for

Poisson regression with random cluster intercepts as both the number of clusters and observations per cluster diverge. Our work is distinct from these results in two ways. First, we provide results at a general level rather than for a specific model. Second, we focus on the asymptotic regime where the number of clusters is diverging, but the number of observations per cluster is stochastically bounded.

More recently, researchers have begun developing general theoretical results for variational estimators. For example, Pati et al. (2018) studied finite-sample risk bounds for mean-field variational Bayes estimators in a very general setting, and applied their results to derive the rate of convergence of variational Bayes estimators in Latent Dirichlet Allocation and Gaussian mixture models. Wang and Blei (to appear) provided sufficient conditions for a Bernstein-von Mises result for the variational Bayes posterior distribution. We note that both of these recent works are distinct from our goals here, which are to study the asymptotic properties of frequentist variational estimators.

2 Variational approximations and M -Estimation

In this section, we demonstrate the connection between M -estimators and variational inference. The key for this connection is using a profile version of the variational objective function. Viewing variational inference in this way unlocks a deep and broad set of theoretical results developed for M -estimators. We make this connection explicit in this section and then, in Section 4, demonstrate how these theoretical results can be used to develop new methods for scientific practice.

2.1 Variational estimation as M -Estimation

We will study the general properties of the variational estimator $\hat{\theta}_n$ through the lens of M -estimation. An M -estimator of a parameter θ is the maximizer of a data-1 " dependent objective function $M_n(\theta) = \frac{1}{n} \sum_{i=1}^n m(\theta; X_i)$. From (2) we can see that

$$\mathcal{L}_n(\theta, \psi_n; \mathbf{X}_n) = \sum_{i=1}^n v(\theta, \psi_i; X_i) \text{ for } v(\theta, \psi; x) = \int \log \left(\frac{p_\theta(x, z)}{q(z; \psi)} \right) q(z; \psi) d\mu(z).$$

Applying the theory of M -estimation to the vector (θ, ψ_n) with $m(\cdot) = v(\cdot)$ is complicated due to the dependence on ψ_i , which are known as *incidental* parameters specific to each data point. The θ , in contrast, are *structural* parameters shared across all data (Lancaster 2000). Hall, Ormerod and Wand (2011) dealt with this problem for Poisson mixed models by assuming the cluster size was growing with the number of observations, so that the incidental parameters effectively became structural. In our more general setting we could analogously assume that each observed data X_i is composed of replicates $X_{i1}, \dots, X_{im}$ and let m grow with n . However, this would limit the applicability of our results to only cases where clusters are very large. Since, in practice, clusters are often small, we instead apply M -estimation to the *profiled* variational objective.

In order to use the M -estimation framework for the variational estimator $\hat{\theta}_n$, we will express the optimization defined in (2) as a two-stage procedure, where first $\mathcal{L}_n$ is optimized with respect to $\boldsymbol{\psi}_n$ for each fixed θ , then this *profiled* function is optimized with respect to θ . Furthermore, we note that optimizing $\mathcal{L}_n$ with respect to $\boldsymbol{\psi}_n$ for fixed θ is equivalent to optimizing each summand $\nu(\theta, \boldsymbol{\psi}_i; x)$ with respect to $\boldsymbol{\psi}_i$ for fixed θ . Therefore, we will assume that for each $\theta \in \Theta$ and P_0 -a.e. x , the map $\boldsymbol{\psi} \mapsto \nu(\theta, \boldsymbol{\psi}; x)$ possesses a unique point of maximum in Ψ , which we will denote by $\hat{\boldsymbol{\psi}}(\theta; x)$. We then define the profiled single-data objective function

$$m(\theta; x) := \sup_{\boldsymbol{\psi} \in \Psi} \nu(\theta, \boldsymbol{\psi}; x) = \nu(\theta, \hat{\boldsymbol{\psi}}(\theta; x); x).$$

Proposition 1 below asserts that, with this assumption, the variational estimator $\hat{\theta}_n$ of the model parameters from equation (2) is equal to the maximizer of the profiled criterion function $\sum_{i=1}^n m(\theta; X_i)$. This result formally establishes the connection between variational inference and M -estimation that we will use throughout this article.

Proposition 1. *Suppose that, for all $\theta \in \Theta$ and P_0 -a.e. x , $\boldsymbol{\psi} \mapsto \nu(\theta, \boldsymbol{\psi}; X)$ possesses a unique maximizer in Ψ , and that $\theta \mapsto \sum_{i=1}^n m(\theta; X_i)$ possesses at least one maximizer in Θ . Then $\hat{\theta}_n \in \arg \max_{\theta \in \Theta} \sum_{i=1}^n m(\theta; X_i)$.*

The proofs of all results are provided in the supplementary material.

The representation of $\hat{\theta}_n$ provided by Proposition 1 now falls within the M -estimator framework. We can therefore use the existing, well-studied asymptotic theory of M -estimators to better understand the asymptotic properties of variational estimators. In the subsequent sections, we show that, using this representation, the theory for M -estimators yields general results for consistency and asymptotic normality for variational estimators.

2.2 Consistency

We first explore consistency using the M -estimator representation of the variational estimator. An important point that we will return to later is that, depending on the model and approximation, the estimator based on the variational lower bound may not be consistent for the truth. Hence, in what follows, we refer to $\bar{\theta}$ as the limit of $\hat{\theta}_n$, so that $\bar{\theta} = \theta_0$ if and only if $\hat{\theta}_n$ is consistent.

The population objective function $M_0(\theta) = E_{P_0}[m(\theta; X)]$ governs the asymptotic properties of the variational estimator $\hat{\theta}_n$. Under regularity conditions, $\hat{\theta}_n \xrightarrow{P_0} \arg \max_{\theta} M_0(\theta)$, so that if M_0 is uniquely maximized at θ_0 then $\hat{\theta}_n$ is consistent for θ_0 , as we state below.

Theorem 1. *Suppose the function $M_0(\theta) = E_{P_0}[\nu(\theta, \hat{\boldsymbol{\psi}}(\theta; X); X)]$ attains a finite global maximum at $\bar{\theta}$ and conditions (A1)-(A3) hold. Then $\hat{\theta}_n \xrightarrow{P_0} \bar{\theta}$. Regularity conditions (A1)-(A3) justify the application of Theorem 5.14 of van der Vaart (2000) and are provided in the*

supplementary material Condition (A1) requires that $v(\theta, \hat{\psi}(\theta; x); x)$ be upper semi-continuous in θ for a.e. x . This is implied, for instance, if v is upper semi-continuous in θ and ψ and $\hat{\psi}$ is continuous in θ , for a.e. x . Condition (A2) requires that v have a measurable and integrable local envelope function. Condition (A3) requires that $\hat{\theta}_n$ be contained in a compact with probability tending to one. If the parameter space is not compact, (A3) can often be established via a suitable compactification of the parameter space, as in van der Vaart (2000) Example 5.16.

Here and throughout, we define D_ψ and D_θ as the derivative operators with respect to ψ and θ , respectively. If v and $\hat{\psi}$ are sufficiently smooth functions of θ for P_0 -a.e. x (see the supplementary material for additional details), then the Leibniz integral rule implies that $D_\theta M_0 |_{\theta = \bar{\theta}} = E_{P_0}[D_\theta v |_{\theta = \bar{\theta}, \psi = \hat{\psi}(\bar{\theta}; X)}]$, and furthermore since $D_\psi v |_{\psi = \hat{\psi}(\theta; x)} = 0$ by definition of $\hat{\psi}$ as a maximizer, $E_{P_0}[D_\theta v |_{\theta = \bar{\theta}, \psi = \hat{\psi}(\bar{\theta}; X)}] = 0$ as well. Therefore, a preliminary step in assessing whether $\hat{\theta}_n$ is consistent is to determine whether $E_{P_0}[D_\theta v |_{\theta = \bar{\theta}, \psi = \hat{\psi}(\bar{\theta}; x)}] = 0$. If it does not equal zero, then $\hat{\theta}_n$ cannot be consistent. If it does equal zero, and in addition $M_0(\theta)$ is strictly concave and regularity conditions (A1)–(A3) hold, then $\hat{\theta}_n$ is consistent.

In practice, it is often not possible to derive $\hat{\psi}(\theta; x)$ in closed form, which prevents a theoretical assessment of consistency of the variational estimator. This is the situation, for instance, in many generalized linear mixed models. In Section 4.3, we propose an empirical method of assessing consistency that does not require explicit derivation of $\hat{\psi}(\theta; x)$.

2.3 Asymptotic normality

If the variational estimator $\hat{\theta}_n$ is consistent for $\bar{\theta}$ and additional regularity conditions hold then $\sqrt{n}(\hat{\theta}_n - \bar{\theta}) \xrightarrow{d} N(0, V(\bar{\theta}))$ where $v(\theta)$ is the sandwich covariance. Here and throughout, we denote by $D_\bullet^2$ the second derivative operator with respect to $\bullet$.

Theorem 2. Suppose $\hat{\theta}_n \xrightarrow{P_0} \bar{\theta}$, a point of maximum of $M_0(\theta) = E_{P_0}[m(\theta; X)]$, and conditions (B1)–(B4) hold. Then

$$\sqrt{n}(\hat{\theta}_n - \bar{\theta}) \xrightarrow{d} N_d(0, V(\bar{\theta}))$$

where $V(\theta) = A(\theta)^{-1}B(\theta)A(\theta)^{-1}$ for

$$A(\theta) = E_{P_0}\left[D_\theta^2 m(\theta; X)\right] \quad (3)$$

$$B(\theta) = E_{P_0}\left[(D_\theta m(\theta; X))(D_\theta m(\theta; X))^T\right]. \quad (4)$$

In the next section we provide formulas for estimating the matrices A and B regardless of whether $m(\theta, X)$ is known explicitly.

Conditions (B1)–(B4), stated in the supplementary material, guarantee that $m(\theta, X)$ satisfies the conditions of van der Vaart (2000) Theorem 5.23. Condition (B1) states that $\hat{\psi}(\theta; x)$ exists for all θ and a.e. x , and (B2) states that it is twice continuously differentiable in θ in a neighborhood of $\bar{\theta}$ for a.e. x . If for θ in a neighborhood of $\bar{\theta}$ and a.e. x , (i) v is twice continuously differentiable in ψ , (ii) $D_{\psi} v|_{\theta, \hat{\psi}(\theta; x)} = 0$, (iii) $D_{\psi}^2 v$ is invertible, and (iv) $D_{\psi} v$ is twice continuously differentiable in θ , then the implicit function theorem implies (B1) and (B2).

Condition (B3) requires that v be twice continuously differentiable in θ in a neighborhood of $\bar{\theta}$ and $\hat{\psi}(\bar{\theta}; x)$ for a.e. x . The differentiability of v required by this condition and the implicit function theorem from the previous paragraph depend on the smoothness of $p_{\theta}(x, z)$ and the variational density $q(z, \psi)$. For instance, by the Leibniz integral rule, if P_{θ} is twice continuously differentiable in θ at x and for $q(\cdot; \psi)$ -a.e. z , and its second derivative is dominated by a $q(\cdot; \psi)$ -integrable function, then v is twice continuously differentiable in θ at ψ and x .

Finally, condition (B4) requires that v and $\hat{\psi}$ be Lipschitz functions in neighborhoods of $\bar{\theta}$ and $\hat{\psi}(\bar{\theta}; x)$ for every x , and that their Lipschitz constant be bounded by a square-integrable function of x . The Lipschitz property of v and $\hat{\psi}$ for fixed x is implied by the differentiability required by (B2) and (B3). Square-integrability of the Lipschitz constant as a function of x is not guaranteed, but is a relatively mild requirement since the neighborhoods around $\bar{\theta}$ and $\hat{\psi}(\bar{\theta}; x)$ may be arbitrarily small.

3 Illustrations of the general theory

In this section, we illustrate the use of our theoretical results for assessing the consistency and asymptotic efficiency of variational estimators in two mixture models. For each model, we highlight the main features necessary to apply our general results, and leave detailed derivations for the supplementary material.

3.1 Consistent and efficient variational estimation

As our first illustration of our general theoretical results, we demonstrate that a variational estimator is consistent and efficient in an exponential mixture model. Suppose that each data unit i consists of a vector of observations $X_i = (X_{i1}, \dots, X_{ip})$. Conditional on independent latent random variables $Z_1, \dots, Z_n$ each distributed as $\text{Exp}(\beta)$, these observations are generated independently as $X_{ij} \sim \text{Exp}(Z_j)$. The parameter vector is $\theta = \beta \in \mathbb{R}^+$. This could serve, for instance, as a model of the lifetimes of clusters of memoryless units.

The marginal density of X_i is now $p_{\theta}(x) = \Gamma(p+1)\beta\left(\beta + \sum_{j=1}^p x_j\right)^{-(p+1)}$, and hence the true conditional distribution of Z_i given X_i is $\text{Gamma}(p+1, \beta + \sum_{j=1}^p x_j)$. Therefore, any variational family of conditional distributions that includes the gamma family as a sub-class

will yield a variational estimator $\hat{\theta}_n$ that is equal to the MLE. However, for the purpose of demonstrating our theoretical method of assessing consistency, it is illustrative to consider a variational class that does not include the true conditional distribution. We will show that in this example, using the mis-specified variational class of log-normal distributions still yields a consistent, and even efficient, variational estimator.

Suppose the variational class is taken to be all log-normal distributions, parametrized by $\psi = (\mu, \sigma^2) \in \mathbb{R} \times \mathbb{R}^+ = \Psi$. Straightforward computation then gives

$$\nu(\theta, \psi; x) \propto \log \beta + (p+1)\mu - \left(\beta + \sum_{j=1}^d x_j \right) e^{\mu + \sigma^2 / 2} + \log \sigma.$$

This is a smooth function, and by composition laws for concave functions, we can see that $\nu(\theta, \psi; x)$ is strictly concave in ψ for each fixed θ and x (Boyd and Vandenberghe 2004). Therefore, the unique zero of the gradient of ν with respect to ψ is the unique ψ maximizing ν for fixed θ and x . This gives $\hat{\mu}(\theta; x) = \log \frac{p+1}{\beta + \sum_{j=1}^p x_j} - (p+1)^{-1} / 2$ and

$\hat{\sigma}(\theta; x) = (p+1)^{-1/2}$. Thus, the profile objective function $m(\theta; x) = \nu(\theta, \hat{\psi}(\theta; x); x)$ can be written explicitly up to a constant as $\log \frac{\beta}{(\beta + \sum_{j=1}^p x_j)^{p+1}}$. Condition (A1) is satisfied

because m smooth in θ .

Condition (A2) is satisfied because $\sup_{\theta} m(\theta; x) = c - p \log(\sum_{j=1}^p x_j)$ for some $c < \infty$, and the expectation of this expression is finite. Condition (A3), which requires tightness of $\hat{\theta}_n$, can be established either by restricting the parameter space to a compact, or by extending the parameter space to $[0, \infty]$ equipped with the metric $d(\beta_1, \beta_2) = |\arctan \beta_1 - \arctan \beta_2|$, as in van der Vaart (2000) Example 5.16.

Since conditions (A1)–(A3) hold, Theorem 1 implies that $\hat{\theta}_n \xrightarrow{P} \bar{\theta}$, the point of maximum of $\theta \mapsto M_0(\theta) = E_{P_0}[m(\theta, X)]$. In this case, since $m(\theta, x)$ is equal up to a constant to the log-likelihood of a single observation, by a standard argument involving Jensen's inequality, $M_0(\theta)$ is uniquely maximized at θ_0 . Therefore, $\hat{\theta}_n$ is consistent even though the variational class *does not include* the true conditional distribution.

Conditions (B1)–(B3) are satisfied because both ν and $\hat{\psi}$ are smooth in θ , and the second derivative of m is bounded up to a constant in a neighborhood of θ_0 by $(\sum_{j=1}^p x_j)^{-1}$, which is P_0 -integrable. The Lipschitz condition (B4) is also satisfied because ν and $\hat{\psi}$ are differentiable with bounded derivatives in a neighborhood of $(\theta_0, \hat{\psi})$ and θ_0 , respectively. The asymptotic variance of $\sqrt{n}(\hat{\theta}_n - \theta_0)$, as implied by Theorem 2, is equal to $(1 + 2/p)\beta_0^2$.

3.2 Inconsistent variational estimation

We now consider an extension of the previous model in which a variational estimator is inconsistent. We keep an identical setup from the previous model, but now, we model the latent variable as $Z_i \sim \text{Gamma}(\alpha, \beta)$ rather than $E \propto \beta$. This is a more flexible model indexed by the parameter $\theta = (\alpha, \beta) \in \mathbb{R}^+ \times \mathbb{R}^+$.

The marginal density of X_p is now $p_\theta(x) = \Gamma(p + \alpha) \Gamma(\alpha)^{-1} \beta^\alpha (\beta + \sum_{j=1}^p x_j)^{-(p + \alpha)}$, and hence the true conditional distribution of Z_i given X_i is $\text{Gamma}(p + \alpha, \beta + \sum_{j=1}^p X_{ij})$. As before, for illustrative purposes we take the variational class to be all log-normal distributions, parametrized by $\psi = (\mu, \sigma^2) \in \mathbb{R} \times \mathbb{R}^+ = \Psi$. We now have

$$v(\theta, \psi; x) \propto \alpha \log \beta - \log \Gamma(\alpha) + (p + \alpha) \mu - \left(\beta + \sum_{j=1}^p x_j \right) e^{\mu + \sigma^2 / 2} + \log \sigma.$$

Once again, v is a smooth function, and is strictly concave in ψ for each fixed θ and x . Setting its derivative with respect to μ and σ to zero and solving gives

$$\hat{\mu}(\theta; x) = \log \frac{p + \alpha}{\beta + \sum_{j=1}^p x_j} - (p + \alpha)^{-1} / 2 \text{ and } \hat{\sigma}(\theta) = (p + \alpha)^{-1} / 2. \text{ Thus,}$$

$$\begin{aligned} m(\theta; x) &\propto \alpha \log \beta - \log \Gamma(\alpha) + (p + \alpha) \log \frac{p + \alpha}{\beta + \sum_{j=1}^p x_j} - (p + \alpha) - \frac{1}{2} \log(p + \alpha) \\ &= \log p_\theta(x) - \log \Gamma(p + \alpha) - (\alpha + p) + (p + \alpha) \log(p + \alpha) - \frac{1}{2} \log(p + \alpha). \end{aligned}$$

Conditions (A1)-(A3) can be checked for this example much as in the previous example. Therefore, Theorem 1 again implies that $\hat{\theta}_n$ tends in probability to the point of maximum of $M_0(\theta)$. As before, M_0 is not available in closed form in terms of elementary functions. However, we have $M_0(\theta) = E_{P_0} [\log p_\theta(X)] + f(\alpha)$, where $f'(\alpha) > 0$ for all α . Since $E_{P_0} [\log p_\theta(X)]$ is smooth and maximized at θ_0 , this implies that $D_\theta M_0|_{\theta=\theta_0} = 0$, so that θ_0 cannot be the point of maximum of M_0 . This shows that the variational estimator using a mis-specified log-normal conditional distribution is inconsistent in this example.

While the limit $\bar{\theta}$ is not available explicitly, we can approximate it using numerical integration and optimization. Figure 1 shows the limits of the variational estimators as a function of the true parameter value for $p = 5$. The bias is small when α_0 and β_0 are small, but increases as α_0 and β_0 grow.

4 Practical tools for inference with variational estimators

We now propose three methodological innovations based on the asymptotic results from Section 2. First, we demonstrate how to leverage asymptotic normality to enhance uncertainty estimators. Second, we show that a one-step correction can be applied to improve the efficiency of the variational estimator. Finally, we address the difficulty of

theoretical assessment of consistency mentioned in Section 2, providing a way to test the consistency of a variational estimator when theoretical calculations are intractable.

4.1 Sandwich covariance estimation

We now discuss computation of consistent covariance estimators. Recall that in practice, $m(\theta; X)$ is often not available in closed form. Fortunately, the derivatives of $m(\theta; X)$ can be expressed in terms of the derivatives of $\nu(\theta, \psi; X)$, which are always available, because $\nu(\theta, \psi; X)$ is a result of the model and variational family used. Thus, using the chain rule, the asymptotic variance can be estimated whether or not $m(\theta; X)$ is available explicitly. We denote by $D_\theta \nu$ and $D_\psi \nu$ the first partial derivatives of ν , and $D_{\theta\theta}^2 \nu$, $D_{\theta\psi}^2 \nu$, and $D_{\psi\psi}^2 \nu$ the second derivatives of ν .

Concerning $D_\theta m(\theta; x)$, which appears in equation (4), since $\hat{\psi}(\theta; x)$ maximizes ν for fixed θ, x , $D_\theta m(\theta; x) = D_\theta \nu(\theta, \psi; x) |_{\psi = \hat{\psi}(\theta; x)}$. For $D_\theta^2 m(\theta; x)$ in equation (3),

$$D_\theta^2 m(\theta; x) = \left[D_{\theta\theta}^2 \nu - D_{\theta\psi}^2 \nu \left(D_{\psi\psi}^2 \nu \right)^{-1} D_{\theta\psi}^2 \nu^T \right]_{\psi = \hat{\psi}(\theta; x)},$$

as we show in the supplementary material, where we abbreviate $\nu(\theta, \psi; X)$ as ν for presentation. Replacing the appropriate derivatives in the definition of $V(\theta)$ with the above expressions and the population expectations with empirical ones gives a way to calculate the asymptotic covariance only knowing $\nu(\theta, \psi; X)$ and its derivatives (which can be calculated numerically), as opposed to $m(\theta; x)$, the computation of which involves optimization.

We now have $\hat{A}_n(\hat{\theta}_n)^{-1} \hat{B}_n(\hat{\theta}_n) \hat{A}_n(\hat{\theta}_n)^{-1} \xrightarrow{P} V(\bar{\theta})$ where

$$\hat{A}_n(\theta) = \frac{1}{n} \sum_{i=1}^n \left[D_{\theta\theta}^2 \nu - D_{\theta\psi}^2 \nu \left(D_{\psi\psi}^2 \nu \right)^{-1} D_{\theta\psi}^2 \nu^T \right]_{\psi = \hat{\psi}_i, x = X_i}, \quad (5)$$

$$\hat{B}_n(\theta) = \frac{1}{n} \sum_{i=1}^n \left[(D_\theta \nu)(D_\theta \nu)^T \right]_{\psi = \hat{\psi}_i, x = X_i}. \quad (6)$$

Equations (5) and (6) provide a formula for constructing an asymptotic covariance matrix for the variational estimator $\hat{\theta}_n$. This covariance can be used to construct asymptotically calibrated Wald intervals, regions, and hypothesis tests about θ_0 if $\bar{\theta} = \theta_0$. Furthermore, the sandwich covariance is model-robust in the sense that it is valid even if $P_0 \notin \mathcal{P}$.

For an MLE under correct model specification, $A(\theta) = B(\theta)$ and the asymptotic covariance reduces to $A(\theta)^{-1}$, the inverse Fisher information matrix. In this case the sandwich covariance is only needed for model-robust uncertainty estimation. However, when m is not proportional to the log-likelihood, as is often true with variational inference, A and B are not

necessarily equal even under correct model specification. Therefore the sandwich covariance is necessary even if $P_0 \in \mathcal{P}$.

4.2 One-step correction

The variational estimator $\hat{\theta}_n$ is not guaranteed to be asymptotically efficient since the variational objective function need not be proportional to the log-likelihood. Hence while Wald-type intervals, regions, and tests using the sandwich estimator proposed in the last section will be asymptotically valid, they may be suboptimal since $\hat{\theta}_n$ may have larger asymptotic variance than the MLE. In these cases, a one-step correction to the variational estimator yields a more efficient estimator.

The one-step estimator is $\hat{\theta}_n^{(1)} = \hat{\theta}_n - I_n(\hat{\theta}_n)^{-1} S_n(\hat{\theta}_n)$, where $l(\theta, x) = \log p_{\theta}(x)$ and

$$S_n(\theta) = \frac{1}{n} \sum_{i=1}^n D_{\theta} l(\theta; X_i), \quad I_n(\theta) = \frac{1}{n} \sum_{i=1}^n (D_{\theta} l(\theta; X_i))(D_{\theta} l(\theta; X_i))^T$$

are the score and observed information at θ . Under regularity conditions $\sqrt{n}(\hat{\theta}_n^{(1)} - \theta_0) \xrightarrow{P} N(0, I(\theta_0)^{-1})$ for $I(\theta_0)$ the Fisher information matrix, which is the same asymptotic distribution as the maximum likelihood estimator or posterior mean.

Computing S_n and I_n require numerical integration in the same way that computing the MLE would. Indeed, the one-step correction is a single step of a Newton-Raphson algorithm for finding the MLE starting at $\hat{\theta}_n$. However, unlike finding the MLE, this one-step procedure only requires a single calculation of these quantities, so requires less computation than finding the exact MLE. Nevertheless, in some cases the one-step correction may not be computationally feasible for the same reasons that computing the MLE is not.

4.3 An empirical test of the consistency of variational estimators

In many cases, including generalized linear mixed models, neither $\hat{w}(\theta; x)$, $m(\theta; x)$, nor $M_0(\theta)$ are available analytically. This presents a challenge not present in the classical M -estimation scenario and seriously undermines the goal of theoretically evaluating the consistency of variational estimators. Simulation studies could be used to assess consistency for any particular fixed, known truth, but would be computationally burdensome.

Here, we propose a method for evaluating the consistency of a variational estimator at a single fixed parameter value θ^* when $m(\theta; x)$ is not available explicitly. Suppose that the data were generated from $P_0 = P_{\theta^*}$. Then a crucial condition for consistency of the variational estimator at θ^* , as stated in Theorem 1, is that $M^*(\theta) := E_{\theta^*}[m(\theta; X)]$ be maximized at θ^* . If M^* is smooth and θ^* is in the interior of the parameter space, then M^* being maximized at θ^* implies that $D_{\theta} M^*(\theta^*) = 0$. Furthermore, as long as $|D_{\theta} m(\theta; x)| \leq h(x)$ for all θ in a neighborhood of θ^* and P_{θ^*} -a.e. x for a P_{θ^*} -integrable function h , then by the dominated convergence theorem,

$D_{\theta} M^*(\theta^*) = E_{\theta^*}[D_{\theta} m(\theta^*; X)] = E_{\theta^*}[D_{\theta} v(\theta^*, \hat{\psi}(\theta^*; X); X)]$. Our proposed method for numerically evaluating consistency of the variational estimator under P_{θ^*} is motivated by numerically assessing whether $E_{\theta^*}[D_{\theta} v(\theta^*, \hat{\psi}(\theta^*; X); X)] = 0$. Our method unfolds in the following steps.

1. Fix θ^* and b very large (for instance 10^4 or 10^5).
2. For $j = 1, \dots, b$:
 - a. Simulate $X_j^* \sim P_{\theta^*}$.
 - b. Find $\psi_j^* = \hat{\psi}(\theta^*; X_j^*)$ by numerically optimizing $\psi \mapsto v(\theta^*, \psi; X_j^*)$.
 - c. Evaluate $G_j^* = D_{\theta} v|_{\theta^*, \psi_j^*, X_j^*}$.
3. Test the null hypothesis that $E_{\theta^*}[G_j^*] = 0$ either using independent t -tests on each component or Hotelling's T^2 test on the entire vector.

If the test rejects the null hypothesis then the variational estimator cannot be consistent; if not then one can be arbitrarily certain (with large enough b) that the mean score is zero at θ^* . If a weakly significant p -value is found and it is unclear what to conclude, the experiment could be repeated with a larger b .

This method is a necessary, but not sufficient test of consistency. As we explain more below, asymptotically we expect our method to have few false negatives (indication that the estimator is inconsistent when it is actually consistent) but possibly false positives (indications that the estimator is consistent when it is actually inconsistent). The first reason for potential false positives is that $E_{\theta^*}[G_j^*] = 0$ is a necessary but not sufficient condition for consistency. Even if its gradient is zero, θ^* it need not be a global maxima of the objective function. The second reason for potential false positives is that the method can only assess consistency at a single parameter value θ^* rather than on the entirety of the parameter space. Typically one will first use the variational algorithm to estimate $\hat{\theta}_n$, then use this method to assess consistency at $\theta^* = \hat{\theta}_n$. If the estimator is consistent for every θ in a neighborhood of θ_0 then for n large enough $\hat{\theta}_n$ will be in that neighborhood and the method will not indicate inconsistency. On the other hand if $\hat{\theta}_n \xrightarrow{P_0} \bar{\theta} \neq \theta_0$ then this method is approximately assessing whether the algorithm is consistent near $\bar{\theta}$. If the variational algorithm is consistent at $\bar{\theta}$ but not at θ_0 then the method would indicate that the estimator is consistent when in fact it is not. Despite the possibility of false positives, we do not know of any other practical ways to assess consistency of variational estimators when the limit objective is not available in closed form.

5 Numerical studies

In this section, we empirically evaluate the variational estimator, the sandwich covariance, and the one-step correction in mixed effects logistic regression models. In these models, theoretical assessment of the consistency and efficiency of variational inference is challenging because the profiled criterion function is not available in closed form. Hence, we turn to our empirical assessment of consistency and numerical studies to assess the properties of variational estimators.

We consider mixed effects logistic regression models – first with random intercepts, then with random intercepts, slopes, and quadratic terms – using data on marijuana use in adolescents in the United States from the National Longitudinal Survey of Youth 1997 (Bureau of Labor Statistics, U.S. Department of Labor 2013). The data consist of approximately yearly interviews of $n = 8660$ youth from 1997 to 2012, with the number of interviews per youth ranging from four to sixteen. For youth i 's j th interview, we consider the binary outcome Y_{ij} of whether the youth used marijuana in the thirty days preceding the interview. We focus on understanding the relationship between marijuana use, age, and sex. Since our goal is to understand the properties of variational estimators, we use the data, along with “known” parameter values, to simulate outcomes. This way we can assess the accuracy of parameter estimates and coverage of uncertainty intervals. We also use our methods to conduct an analysis of the real NLSY data.

The results indicate that variational estimators are not always consistent: in the first example the estimator is consistent for some parameters and not for others, and in the second example it is not consistent for any parameters. The first example also demonstrates that even when the variational estimator is consistent, it is not necessarily efficient. In either case the sandwich covariance matrix provides good confidence interval coverage rates and the one-step correction improves efficiency. The empirical evaluation of consistency correctly identifies inconsistency of the parameter vector as a whole, but not always inconsistency of individual parameters.

5.1 Logistic regression with random intercepts

First we consider logistic regression with random intercepts. Let Z_i be a random intercept controlling each youth's overall propensity for marijuana use, SEX_i be an indicator that the youth is male, and AGE_{ij} be youth i 's age at interview j . Denote $p_{ij} = P(Y_{ij} = 1 \mid Z_i, SEX_i, AGE_{ij})$. Our first model for marijuana usage is then

$$\text{logit}(p_{ij}) = \begin{cases} Z_i + \beta_0 + \beta_1(AGE_{ij} / 35) + \beta_2(AGE_{ij} / 35)^2, & SEX_i = 0 \\ Z_i + \beta_3 + \beta_4(AGE_{ij} / 35) + \beta_5(AGE_{ij} / 35)^2, & SEX_i = 1. \end{cases}$$

Each youth's outcomes $Y_{i1}, \dots, Y_{in_i}$ are assumed conditionally independent given Z_i and we model Z_i as IID $N(0, \sigma^2)$. The parameter vector is $\theta = (\beta, \log(\sigma^2))$. The inclusion of the quadratic effect of age is important because we expect that marijuana usage peaks some time in young adulthood and decreases thereafter. This model form is similar to that used in the analysis of age-crime curves (Fabio et al. 2011).

To estimate θ we consider a variational class of conditional distributions over $\mathbf{z}_n$ consisting of all independent Gaussian distributions. This is known as a Gaussian variational approximation (GVA). The variational parameters are $\psi_i = (m_i, \log s_i)$, m_i being the mean and s_i the standard deviation of the variational conditional distribution of γ_i . The variational objective involves one-dimensional numerical integrals. To optimize the variational objective function we use a variational EM algorithm using the statistical software R (R Core Team 2018). We used the R package fastGHQuad (Blocker 2018) for numerical integration. In this case it is not possible to express the profiled objective function explicitly.

To evaluate our methods, we conducted a simulation study based on the NLSY data. For each of 1000 simulations, we draw a bootstrap sample of youth. Conditional on these youth's age and sex we simulated $Y_1, \dots, Y_n$ from the model, treating the variational estimate $\hat{\theta}_n$ for the data as the true parameter value θ_0 . We then estimated the model parameters and asymptotic covariance matrix using maximum likelihood with the R package lme4 (Bates et al. 2015), the Gaussian variational approximation, and the one-step correction to the variational estimator. Finally, we used our proposed method to assess consistency of the variational algorithm at the estimated parameter value.

We first examine the accuracy of point estimates for regression fixed effects. All three estimators concentrate on the true values of the fixed effects β_0 through β_5 (box plots are provided in the supplementary material). Table 1 shows the variance of the estimators for each of the seven model parameters. The variational estimator has slightly larger variance than the MLE, but the one-step correction nearly matches the variance of the MLE. Thus, as we asserted theoretically, the one-step correction is efficient as long as the variational estimator is consistent, even when the variational estimator is inefficient.

Moving now to the random intercept variance, the MLE concentrates on the true variance component, $\log(\sigma^2)$, while the variational estimator and one-step correction do not. Our empirical assessment of consistency described in Section 4.3 correctly identifies this inconsistency. The multivariate Hotelling test rejected in every simulation with $p < 10^{-16}$, correctly indicating that the population mean gradient of the entire parameter vector was significantly different from zero. Additionally, no more than 2.5% of the marginal t -tests rejected at the 0.01 level for each of the fixed effects, in line with their apparent consistency, while every one of the 1000 simulations rejected the marginal t -test for the variance parameter with $p < 10^{-16}$. These results are better than what is guaranteed theoretically, since the theory does not guarantee that the marginal t -tests will accurately reflect the consistency or inconsistency of individual parameters.

We now move to examining uncertainty intervals. Table 2 shows the estimated coverage of marginal 95% Wald-type confidence intervals of the model parameters for each the three estimators. The coverage of the confidence intervals of the linear and quadratic age fixed effects using the sandwich covariance for the variational estimator and the inverse Fisher information for the one-step correction are within the Monte Carlo error of the nominal 95%. The variational confidence intervals for the sex-specific intercepts β_0 and β_3 are too small at 90%, likely because of the underestimation of σ^2 , the variance of the random intercept. The coverage of the confidence intervals of $\log \sigma^2$ is close to zero for the

variational estimator and one-step correction, which is not surprising given that the estimator is inconsistent. The coverage of $\log \sigma^2$ is not shown for the MLE because lme4 does not provide an interval for this parameter.

5.2 Logistic regression with random quadratics

We now alter the model presented above to include random slopes and quadratic terms for each youth. The random intercept model may not accurately capture the dependence structure of a single subject's marijuana use over time, since the random intercepts model implies an exchangeable marginal correlation structure, which is unrealistic given the longitudinal nature of the data. A more realistic model allows random slopes and quadratic terms as well, so that the latent variable Z_i now has three components. For the conditional probability $p_{ij} = P(Y_{ij} = 1 \mid Z_i, SEX_i, AGE_{ij})$ we now have

$$\text{logit}(p_{ij}) = \begin{cases} (Z_{i0} + \beta_0) + (Z_{i1} + \beta_1)(AGE_{ij} / 35) + (Z_{i2} + \beta_2)(AGE_{ij} / 35)^2, & SEX_i = 0 \\ (Z_{i0} + \beta_3) + (Z_{i1} + \beta_4)(AGE_{ij} / 35) + (Z_{i2} + \beta_5)(AGE_{ij} / 35)^2, & SEX_i = 1. \end{cases}$$

Thus β_0, β_1 , and β_2 are the coefficients of the quadratic curve for the average female, and analogously for males. We model the random effects $\mathbf{z}_n$ as IID mean zero multivariate Gaussian with covariance matrix Σ . Once again we use MLE, a Gaussian variational approximation, and a one-step correction to the Gaussian variational approximation to estimate the average random effects and covariance matrix.

The marginal likelihood for this model involves an intractable integral over $\mathbb{R}^3$. GVA only requires numerical integration over a one-dimensional integral, and is hence less computationally intensive than ML estimation. To compare the computational burden of these methods, we simulated 100 data sets each at sample sizes $n = 100, 250$, and 500 , and used the lme4 package, which uses a Laplace approximation to the log-likelihood, GVA, the one-step correction, and ML estimation to obtain estimates of the parameters in the logistic regression with random quadratics model. We used the implementation of the L-BFGS-B algorithm (Byrd et al. 1995) in the optim function in R (R Core Team 2018) to compute the GVA and MLE.

Figure 3 shows box plots of the computation time in minutes of these four algorithms. The MLE was the most computationally expensive – at sample size $n = 500$, the average computation time was already 70 minutes. GVA and GVA+OS required an average of 6.3 and 9.7 minutes respectively to compute with $n = 500$ observations. lme4 was the most computationally efficient, requiring an average of 2.4 minutes.

We conducted a simulation study with the same structure as the study in the last section to compare the point estimates and CI coverage of lme4 (Bates et al. 2015), GVA, and GVA+OS using all 8660 observations. Box plots of the three estimators of the mean random effects are shown in Figure 2. The pattern is very different from the random intercept model. The lme4 estimates are slightly inconsistent (for random effects with dimension larger than one the lme4 package uses a Laplace approximation to the likelihood). The GVA estimates are even more biased than the lme4 estimates. Despite this, it appears that our proposed one-

step corrected fixed effects are roughly centered around the true values. This is surprising since our theory does not guarantee that the one-step correction will be consistent when the variational estimate is not. All three estimators performed quite poorly in terms of estimating the covariance matrix of the random effects.

Table 3 shows the estimated coverage of 95% CIs for the mean random effects for the three estimators. The variational sandwich coverage was close to 0 in every case due to the bias in the parameter estimate seen in Figure 2. The lme4 CIs also do not perform well, with substantially lower than desired coverage. The one-step correction coverage is closest to the desired 95%. These intervals are conservative, containing the true value more than 95% of the time.

The multivariate Hotelling test of consistency soundly rejected for every simulation, correctly indicating that the variational parameter estimator is consistent. In practice, therefore, while we would not be able to perform the same empirical evaluation as we have here (since we do not know the true model parameters to compute accuracy and coverage), we would have information that calls into question the viability of the GVA procedure for this model. The marginal t -tests of consistency for β did not reject in the majority of simulations. Hence, while the marginal t -tests were an accurate diagnostic tool in the random intercept setting, they were not in the random quadratic setting.

5.3 Analysis of marijuana use in NLSY97

We used our one-step corrected estimator to assess the likelihood of marijuana usage by age and sex in the NLSY. We performed our analysis in R (R Core Team 2018) using the packages numDeriv (Gilbert and Varadhan 2016), parallel (R Core Team 2018), fastGHQuad (Blocker 2018), ggplot2 (Wickham 2016), reshape (Wickham 2007), plyr (Wickham 2011), MASS (Venables and Ripley 2002), and DescTools (Signorell 2019). Figure 4 shows the estimated mean curves as a function of age for both the random intercepts and random quadratics models and for both females and males. The average male and female from the random quadratics model have slightly faster increases, peak at younger ages, and decrease earlier than the average male and female from the random intercepts model. In both models the average male has higher overall probability and slightly later peak usage: in the random intercepts model, the estimated peak female usage probability occurs at 21.3 years (95% CI: [18.1, 24.5]), and peak male usage at 22.2 years (95% CI: [19.9, 24.6]). In the random quadratics model, the estimated peak female usage probability occurs at 17.9 years (95% CI: [15.9, 20.0]), and peak male usage at 19.1 years (95% CI: [17.3, 20.8]).

6 Discussion

We have presented a general framework for understanding the properties of variational estimation for parametric mixture models. The key insight of our work comes from representing the profiled variational objective function as an M -estimator. Once we make this connection, we can leverage a rich toolkit of asymptotic and methodological results available for this context.

The theory does not guarantee that variational estimators are consistent, and it is often difficult to derive the profile objective function necessary to assess consistency. We proposed an empirical test of consistency based on estimating the gradient of the profile objective at a single parameter value. This proposed method worked well in practice, correctly indicating whether variational estimator is inconsistent in two generalized linear mixed models.

We also used the asymptotic theory to propose a sandwich covariance estimator to provide calibrated confidence regions of variational estimators and a one-step correction to the variational estimator. Both of these methods work well when the variational estimator is consistent, and in fact the one-step correction exceeded our expectations by correcting some of the bias in fixed-effect variational parameter estimators in a logistic regression model with random quadratics.

Our theory is limited to models which are IID at some level. While this includes many hierarchical and longitudinal models, it excludes models for fully dependent time series, spatial data, and dyadic data. Extending the theory to cover those cases could be a fruitful next step. Additionally, we made the simplifying assumption that the variational class is of fixed and finite dimension, but our theory could be extended to other variational classes.

Our theory only provides results regarding the asymptotic behavior of the variational estimator $\hat{\theta}_n$ of the structural parameters θ_0 . It does not cover the behavior of the variational parameters $\hat{\psi}_i$, which govern the unit-specific variational conditional distributions of the latent variable Z_i given the observed data X_i . The behavior of these parameters is important in settings where confidence regions with valid coverage are desired for the Z_i . This is the case, for instance, when these Z_i correspond to specific fixtures of the real world, such as counties or schools. However, we note that, since the variational family typically does not contain the true conditional distribution, it may be difficult to provide regions with good coverage of Z_i using variational inference. By contrast, as we have demonstrated, $\hat{\theta}_n$ may be consistent even when the variational family does not contain the true conditional. This was one reason we chose to focus on the asymptotic behavior of $\hat{\theta}_n$. Nevertheless, we would conjecture that $\hat{\psi}_i$, given X_i converges in distribution to

$$\arg \max_{\psi \in \Psi} \int \left(\log \frac{p_{\bar{\theta}}(X_i, z)}{q(z; \psi)} \right) q(z; \psi) dz$$

for each i , where $\bar{\theta}$ is the limit in probability of the variational estimator $\hat{\theta}_n$. We leave further discussion along these lines to future work.

We developed our theory in the context of a fixed variational family of conditional distributions. A natural question is whether our theory provides insight about which variational families yield consistency. Unfortunately, it appears to be difficult to address this question in a general manner. As a simple example, it would intuitively seem that if a given class of variational distributions yields a consistent estimator, then any enlargement of the

class should also yield a consistent estimator. However, it is not clear whether this is true based on our theory. This would be an important topic of future research.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank two referees, an associate editor, and the editor for providing constructive feedback that helped them improve this manuscript. The authors also gratefully acknowledge grant 62389-CS-YIP from the United States Army Research Office, grants SES-1559778 and DMS-1737673 from the National Science Foundation, and grant number K01 HD078452 from the National Institute of Child Health and Human Development (NICHD).

References

- Airoldi EM, Blei DM, Fienberg SE and Xing EP (2008), 'Mixed membership stochastic blockmodels', *Journal of Machine Learning Research* 9, 1981–2014. [PubMed: 21701698]
- Bates D, Mächler M, Bolker B and Walker S (2015), 'Fitting Linear Mixed-Effects Models Using lme4', *Journal of Statistical Software* 67(1), 1–48.
- Bickel P, Choi D, Chang X and Zhang H (2013), 'Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels', *The Annals of Statistics* 41 (4), 1922–1943.
- Blei DM, Kucukelbir A and McAuliffe JD (2017), 'Variational inference: A review for statisticians', *Journal of the American Statistical Association* 112(518), 859–877.
- Blei DM, Ng AY and Jordan MI (2003), 'Latent dirichlet allocation', *Journal of Machine Learning Research* 3, 993–1022.
- Blocker AW (2018), fastGHQuad: Fast 'Rcpp' Implementation of Gauss-Hermite Quadrature. R package version 1.0. **URL:** <https://CRAN.R-project.org/package=fastGHQuad>**URL:**
- Boyd S and Vandenberghe L (2004), *Convex optimization*, Cambridge University Press.
- Bureau of Labor Statistics, U.S. Department of Labor (2013), 'National Longitudinal Survey of Youth 1997 cohort, 1997-2011 (rounds 1-15)', Produced by the National Opinion Research Center, the University of Chicago and distributed by the Center for Human Resource Research, The Ohio State University, Columbus, OH.
- Byrd R, Lu P, Nocedal J and Zhu C (1995), 'A limited memory algorithm for bound constrained optimization', *SIAM Journal on Scientific Computing* 16(5), 1190–1208.
- Davidian M and Giltinan DM (1995), *Nonlinear Models for Repeated Measurement Data*, Vol. 62, CRC press.
- Erosheva E, Fienberg S and Lafferty J (2004), 'Mixed-membership models of scientific publications', *Proceedings of the National Academy of Sciences* 101, 5220–5227.
- Fabio A, Tu L-C, Loeber R and Cohen J (2011), 'Neighborhood socioeconomic disadvantage and the shape of the age-crime curve', *American Journal of Public Health* 101, S325–S332. [PubMed: 21778512]
- Gilbert P and Varadhan R (2016), numDeriv: Accurate Numerical Derivatives. R package version 20168–1. **URL:** <https://CRAN.R-project.org/package=numDeriv>**URL:**
- Goldstein H (2011), *Multilevel Statistical Models*, Vol. 922, John Wiley & Sons.
- Hall P, Humphreys K and Titterton D (2002), 'On the adequacy of variational lower bound functions for likelihood-based inference in Markovian models with missing values', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(3), 549–564.
- Hall P, Ormerod JT and Wand MP (2011), 'Theory of Gaussian variational approximation for a Poisson mixed model', *Statistica Sinica* 21(1), 369–389.
- Hall P, Pham T, Wand MP and Wang SS (2011), 'Asymptotic normality and valid inference for Gaussian variational approximation', *The Annals of Statistics* 39(5), 2502–2532.

- Lancaster T (2000), 'The incidental parameter problem since 1948', *Journal of Econometrics* 95(2), 391–413.
- Lee CYY and Wand MP (2016), 'Variational methods for fitting complex Bayesian mixed effects models to health data', *Statistics in Medicine* 35, 165–188. [PubMed: 26415742]
- McCulloch CE and Neuhaus JM (2001), *Generalized Linear Mixed Models*, Wiley Online Library.
- O'Connor B, Eisenstein J, Xing EP and Smith NA (2010), 'Discovering demographic language variation', Technical report.
- Pati D, Bhattacharya A and Yang Y (2018), On Statistical Optimality of Variational Bayes, in Storkey A and Perez-Cruz F, eds, 'Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics', Vol. 84 of *Proceedings of Machine Learning Research*, PMLR, Playa Blanca, Lanzarote, Canary Islands, pp. 1579–1588.
- Pritchard JK, Stephens M and Donnelly P (2000), 'Inference of population structure using multilocus genotype data', *Genetics* 155(2), 945–959. [PubMed: 10835412]
- R Core Team (2018), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- Raj A, Stephens M and Pritchard JK (2014), 'fastSTRUCTURE: Variational Inference of Population Structure in Large SNP Data Sets', *Genetics* 197(2), 573–589. [PubMed: 24700103]
- Signorell A (2019), DescTools: Tools for Descriptive Statistics. R package version 0.99.27. **URL:** <https://cran.r-project.org/package=DescTools>**URL:**
- van der Vaart AW (2000), *Asymptotic Statistics*, Vol. 3, Cambridge University Press.
- Venables WN and Ripley BD (2002), *Modern Applied Statistics with S*, fourth edn, Springer, New York ISBN 0-387-95457-0. **URL:** <http://www.stats.ox.ac.uk/pub/MASS4>**URL:**
- Wainwright MJ and Jordan MI (2008), 'Graphical models, exponential families, and variational inference', *Foundations and Trends in Machine Learning* 1(1-2), 1–305.
- Wang B and Titterton DM (2004), Convergence and asymptotic normality of variational Bayesian approximations for exponential family models with missing values, in 'Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence', UAI '04, AUAI Press, Arlington, Virginia, United States, pp. 577–584.
- Wang B and Titterton DM (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 and Blei DM (to appear), 'Frequentist Consistency of Variational Bayes', *Journal of the American Statistical Association*.
- Wickham H (2007), 'Reshaping data with the reshape package', *Journal of Statistical Software* 21 (12). **URL:** <http://www.jstatsoft.org/v21/i12/paper>**URL:**
- Wickham H (2011), 'The Split-Apply-Combine Strategy for Data Analysis', *Journal of Statistical Software* 40(1), 1–29. **URL:** <http://www.jstatsoft.org/v40/i01/>**URL:**
- Wickham H (2016), *ggplot2: Elegant Graphics for Data Analysis*, Springer-Verlag New York **URL:** <http://ggplot2.org>**URL:**

Fig. 1.

Limits of the variational parameter estimates in the Exponential-Gamma mixture model using a mis-specified variational conditional distribution. The left display shows the limit $\bar{\alpha}$ as a function of the true α_0 used to generate the data. Note that $\bar{\alpha}$ does not depend on β_0 . The right display shows the limit $\bar{\beta}$ as a function of the true β_0 for four values of α_0 . The identity function is shown as a solid black line.

Fig. 2.

Estimates of mean random effects from the logistic regression with random quadratics simulation study. The dotted line indicates the true parameter value. “lme4” corresponds to estimate from the lme4 package, which uses a Laplace approximation. “GVA” stands for Gaussian variational approximation, and “GVA+OS” refers to the one-step correction.

Fig. 3.

Box plots of the computation time of four estimation methods of the logistic regression with random quadratics model. Three sample sizes are shown. “lme4” refers to the lme4 package, which uses a Laplace approximation to the marginal likelihood, “GVA” stands for Gaussian variational approximation, “GVA+OS” refers to the one-step correction, and “MLE” stands for maximum likelihood estimation, performed using L-BFGS-B optimization.

Fig. 4. One-step correction point estimates and pointwise 95% confidence intervals of probability of having used marijuana in the past month. Curves on the left are for the average female, right are average male. Both the logistic regression with random intercepts and random quadratics are shown.

Table 1

Estimator variances in the logistic regression with random intercepts simulation

	β_0	β_1	β_2	β_3	β_4	β_5	$\log(\sigma^2)$
MLE	0.16	1.68	1.05	0.12	1.25	0.77	8.6×10^{-3}
GVA	0.23	1.90	1.19	0.19	1.46	0.91	0.61
One-step correction to GVA	0.17	1.67	1.04	0.13	1.26	0.78	0.60

Table 2

Coverage of 95% confidence intervals in the logistic regression with random intercepts simulation.

	β_0	β_1	β_2	β_3	β_4	β_5	$\log(\sigma^2)$
Maximum likelihood	0.94	0.94	0.94	0.95	0.95	0.95	–
GVA + sandwich covariance	0.90	0.94	0.94	0.90	0.94	0.94	0.09
One-step correction to GVA	0.93	0.94	0.94	0.91	0.95	0.95	0.02

Table 3

Coverage of 95% CIs in the logistic regression with random quadratics simulation.

	β_0	β_1	β_2	β_3	β_4	β_5
Laplace approximation via lme4	0.72	0.68	0.60	0.67	0.61	0.52
GVA + sandwich	0.01	0.01	0.00	0.013	0.01	0.00
One-step correction to GVA	0.98	0.98	0.98	0.98	0.98	0.98