



Bayesian Estimation and Comparison of Moment Condition Models

Siddhartha Chib^a, Minchul Shin^b, and Anna Simon^c

^aOlin Business School, Washington University in St. Louis, St. Louis, MO; ^bDepartment of Economics, University of Illinois, Urbana, IL; ^cCNRS - CREST, ENSAE, École Polytechnique, 5, Palaiseau, France

ABSTRACT

In this article, we develop a Bayesian semiparametric analysis of moment condition models by casting the problem within the exponentially tilted empirical likelihood (ETEL) framework. We use this framework to develop a fully Bayesian analysis of correctly and misspecified moment condition models. We show that even under misspecification, the Bayesian ETEL posterior distribution satisfies the Bernstein–von Mises (BvM) theorem. We also develop a unified approach based on marginal likelihoods and Bayes factors for comparing different moment-restricted models and for discarding any misspecified moment restrictions. Computation of the marginal likelihoods is by the method of Chib (1995) as extended to Metropolis–Hastings samplers in Chib and Jeliazkov in 2001. We establish the model selection consistency of the marginal likelihood and show that the marginal likelihood favors the model with the minimum number of parameters and the maximum number of valid moment restrictions. When the models are misspecified, the marginal likelihood model selection procedure selects the model that is closer to the (unknown) true data-generating process in terms of the Kullback–Leibler divergence. The ideas and results in this article broaden the theoretical underpinning and value of the Bayesian ETEL framework with many practical applications. The discussion is illuminated through several examples. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received June 2016
Revised July 2017

KEYWORDS

Bernstein–von Mises theorem; Estimating equations; Exponentially tilted empirical likelihood; Marginal likelihood; Misspecification; Model selection consistency

1. Introduction

Our goal in this article is to develop a Bayesian analysis of moment condition models. By moment condition models, we mean models that are specified only through moment restrictions of the type $E^P[\mathbf{g}(X, \theta)] = \mathbf{0}$, where $\mathbf{g}(X, \theta)$ is a known vector-valued function of a random vector X and an unknown parameter vector θ , and P is the unknown data distribution. Models of this type, which arise frequently in statistics and econometrics, see, for example, Broniatowski and Keziou (2012), can be attractive since full modeling of P is not invoked and inferences about θ are based only on the partial information supplied by the set of moment conditions. For instance, in a regression context, letting $X = (y, x)$ and $y = x\beta + \varepsilon$, where y is the scalar response and x is a scalar predictor, one can learn about the regression parameter β from the orthogonality assumption $E^P[(y - x\beta)x] = 0$ without fully modeling the error distribution or the parameters of the error distribution. More generally, β can be inferred in this setting from the orthogonality conditions $E^P[(y - x\beta)\mathbf{z}] = \mathbf{0}$, given a set of instrumental variables \mathbf{z} . Examples of such moment condition models abound, but for the most part the analysis of such models from the Bayesian perspective has proved elusive since typical parametric and semiparametric Bayesian methods are reliant on a full probability model of P .

On the frequentist side, the recent developments in empirical likelihood (EL) based methods, see, for example, Owen (1988, 1990, 2001), Qin and Lawless (1994), Kitamura and

Stutzer (1997), Imbens (1997), Schennach (2007), Chen and Van Keilegom (2009), and references therein, have opened up a promising approach for dealing with moment condition models. There are emerging cogent arguments for using the EL in Bayesian analysis. For example, Lazar (2003) argued that the EL can be used in a Bayesian framework in place of the data distribution P . In fact, Schennach (2005) showed that it is possible to obtain a nonparametric likelihood closely related to EL, called the exponentially tilted empirical likelihood (ETEL), by marginalizing over P with a nonparametric prior that favors distributions that are close to the empirical distribution function in terms of the Kullback–Leibler (KL) divergence while satisfying the moment restrictions. In addition, Grendar and Judge (2009) showed that the EL is the mode of the posterior of P under a general prior on P . Thus, by combining either the EL or the ETEL functions with a prior $\pi(\theta)$ on θ , moment condition models can in principle be subjected to a Bayesian semiparametric analysis. Applications of this idea are given, for instance, by Lancaster and Jun (2010), Kim and Yang (2011), Yang and He (2012), Xi, Li, and Hu (2016) to handle moment condition models, by Rao and Wu (2010) in complex survey estimation, and by Chaudhuri and Ghosh (2011), Porter, Holan, and Winkle (2015), Chaudhuri, Mondal, and Yin (2017) in small area estimation. On the theory side, Yang and He (2012) showed the asymptotic normality of the Bayesian EL posterior distribution of the quantile regression parameter, and Fang and Mukerjee (2006) and Chang and Mukerjee (2008) studied the higher-order asymptotic and

coverage properties of the Bayesian EL/E TEL posterior distribution for the population mean, while Schennach (2005) and Lancaster and Jun (2010) considered the large-sample behavior of the Bayesian E TEL posterior distribution under the assumption that all moment restrictions are valid. Alternative, non-EL/E TEL-based approaches for moment condition models, which we do not consider in this article, have also been examined, for example, Bornn, Shephard, and Solgi (2015), Florens and Simoni (2016), and Kitamura and Otsu (2011).

The purpose of this article is to establish a number of new results for the Bayesian analysis of moment condition models, within the E TEL framework, complementing and extending the aforementioned papers in important directions. One goal is the Bayesian analysis of moment condition models that are potentially misspecified. For this reason, our analysis is built on the E TEL function which, as shown by Schennach (2007), leads to frequentist estimators of θ that have the same orders of bias and variance (as a function of the sample size) as the EL estimators but, importantly, maintain the root n convergence even under model misspecification (see Schennach 2007, Theorem 1). Within this useful framework, we develop a fully Bayesian treatment of correctly and misspecified moment condition models. We show that even under misspecification, the Bayesian E TEL posterior distribution has desirable properties, and that it satisfies the Bernstein–von Mises (BvM) theorem. Another goal is to develop a Bayesian approach for comparing different moment-restricted models and for discarding any misspecified moment restrictions. For an overview on Bayesian model selection in standard models, we refer to Robert (2007), and references therein. Our proposal is to select the model with the largest marginal likelihood. Since one aim of this model selection comparison is to discard misspecified moment restrictions, we do not consider the model averaging perspective. To operationalize model comparisons in our set-up, in particular when models are defined by different numbers of moment conditions, we show that it is necessary to linearly transform the moment functions $\mathbf{g}(\mathbf{X}, \theta)$ so that all the transformed moments are included in each model. This linear transformation simply consists of adding an extra parameter different from zero to the components of the vector $\mathbf{g}(\mathbf{X}, \theta)$ that correspond to the restrictions not included in a specific model.

We compute the marginal likelihood by the method of Chib (1995), as extended to Metropolis–Hastings samplers in Chib and Jeliazkov (2001). This method renders computation of the marginal likelihood simple and is a key feature of both our numerical and theoretical analysis. Our asymptotic theory covers the following exhaustive possibilities: the case where the models in the comparison set contain only valid moment restrictions, the case where all the models in the set are misspecified, and finally the case where some of the models contain only valid moment restrictions while the others contain at least one invalid moment restriction. Our analysis shows that the marginal likelihood-based selection procedure is consistent in the sense that: (i) it discards misspecified moment restrictions, (ii) it selects the model that is the “less misspecified” when comparing models that are all misspecified, (iii) it selects the model that contains the maximum number of overidentifying valid moment restrictions when comparing correctly specified models, and (iv) when some models are correctly specified and some

are misspecified, it selects the model that is correctly specified and contains the maximum number of overidentifying moment conditions. These important model selection consistency results are based on the asymptotic behavior of the E TEL function, and the validity of the BvM theorem, both under correct specification and misspecification. These results, developed within a formal Bayesian setting, can be viewed as complementary to the less Bayesian formulations described by Variyath, Chen, and Abraham (2010) and Vexler, Deng, and Wilding (2013) where the focus is on quasi-Bayes factors constructed from the EL, and Hong and Preston (2012) where models are compared based on a quasi-marginal likelihood obtained from an approximation to the true P .

The rest of the article is organized as follows. In Section 2, we describe the moment condition model, define the notion of misspecification in this setting, and then discuss the prior-posterior analysis with the E TEL function. We then provide the first pair of major results dealing with the asymptotic behavior of the posterior distribution for both correctly specified and misspecified models. Section 3 introduces our model selection procedure based on marginal likelihoods and the associated large sample results. Throughout the article, for expository purposes, we include numerical examples. Then in Section 4 we discuss the problems of variable selection in a count regression model and instrument validity in an instrumental variable regression. Section 5 concludes. Proofs of our results are collected in the Appendix and in the online Appendix.

2. Setting

Suppose that \mathbf{X} is an \mathbb{R}^{d_x} -valued random vector with (unknown) distribution P . Suppose that the operating assumption is that the distribution P satisfies the d unconditional moment restrictions

$$\mathbf{E}^P[\mathbf{g}(\mathbf{X}, \theta)] = \mathbf{0}, \quad (2.1)$$

where \mathbf{E}^P denotes the expectation taken with respect to P , $\mathbf{g}: \mathbb{R}^{d_x} \times \Theta \mapsto \mathbb{R}^d$ is a vector of known functions with values in \mathbb{R}^d , $\theta := (\theta_1, \dots, \theta_p)' \in \Theta \subset \mathbb{R}^p$ is the parameter vector of interest, and $\mathbf{0}$ is the $d \times 1$ vector of zeros. We assume that $\mathbf{E}^P[\mathbf{g}(\mathbf{X}, \theta)]$ is bounded for every $\theta \in \Theta$. We also suppose that we are given a random sample $\mathbf{x}_{1:n} := (\mathbf{x}_1, \dots, \mathbf{x}_n)$ on \mathbf{X} and that $d \geq p$.

When the number of moment restrictions d exceeds the number of parameters p , the parameter θ in such a setting is said to be overidentified (over restricted). In such a case, there is a possibility that a subset of the moment conditions may be invalid in the sense that the true data-generating process is not contained in the collection of probability measures that satisfy the moment conditions for all $\theta \in \Theta$. That is, there is no parameter θ in Θ that is consistent with the moment restrictions (2.1) under the true data-generating process P . To deal with possibly invalid moment restrictions, we reformulate the moment conditions in terms of an additional nuisance parameter $\mathbf{V} \in \mathfrak{V} \subset \mathbb{R}^d$. For example, if the k th moment condition is not expected to be valid, we subtract $\mathbf{V} = (V_1, \dots, V_d)$ from the moment restrictions where V_k is a free parameter and all other elements of \mathbf{V} are zero. To accommodate this situation, we rewrite the above

conditions as the following augmented moment conditions:

$$E^P[\mathbf{g}^A(X, \boldsymbol{\theta}, \mathbf{V})] = \mathbf{0}, \tag{2.2}$$

where $\mathbf{g}^A(X, \boldsymbol{\theta}, \mathbf{V}) := \mathbf{g}(X, \boldsymbol{\theta}) - \mathbf{V}$. Note that in this formalism, the parameter \mathbf{V} indicates which moment restrictions are active where by “active moment restrictions” we mean the restrictions for which the corresponding components of \mathbf{V} are zero. To guarantee identification of $\boldsymbol{\theta}$, at most $(d - p)$ elements of \mathbf{V} can be different than zero. If all the elements of \mathbf{V} are zero, we recover the restrictions in (2.1).

Let $d_v \leq (d - p)$ be the number of nonzero elements of \mathbf{V} and let $\mathbf{v} \in \mathcal{V} \subset \mathbb{R}^{d_v}$ be the vector that collects all the nonzero components of \mathbf{V} . We call \mathbf{v} the augmented parameter and $\boldsymbol{\theta}$ the parameter of interest. Therefore, the number of active moment restrictions is $d - d_v$. In the following, we write $\mathbf{g}^A(X, \boldsymbol{\theta}, \mathbf{v})$ as a shorthand for $\mathbf{g}^A(X, \boldsymbol{\theta}, \mathbf{V})$, with \mathbf{v} the vector obtained from \mathbf{V} by collecting only its nonzero components.

The central problem of misspecification of the moment conditions, mentioned in the preceding paragraph, can now be formally defined in terms of the augmented moment conditions.

Definition 1 (Misspecified model). We say that the augmented moment condition model is misspecified if the set of probability measures implied by the moment restrictions does not contain the true data-generating process P for every $(\boldsymbol{\theta}, \mathbf{v}) \in \Theta \times \mathcal{V}$, that is, $P \notin \mathcal{P}$ where $\mathcal{P} = \bigcup_{(\boldsymbol{\theta}, \mathbf{v}) \in \Theta \times \mathcal{V}} \mathcal{P}_{(\boldsymbol{\theta}, \mathbf{v})}$ and $\mathcal{P}_{(\boldsymbol{\theta}, \mathbf{v})} = \{Q \in \mathbb{M}; E^Q[\mathbf{g}^A(X, \boldsymbol{\theta}, \mathbf{v})] = \mathbf{0}\}$ with \mathbb{M} the set of all probability measures on \mathbb{R}^{d_x} .

In a nutshell, a set of augmented moment conditions is misspecified if there is no pair $(\boldsymbol{\theta}, \mathbf{v})$ in $(\Theta \times \mathcal{V})$ that satisfies $E^P[\mathbf{g}^A(X, \boldsymbol{\theta}, \mathbf{v})] = \mathbf{0}$ where P is the true data-generating process. On the other hand, if such a pair of values $(\boldsymbol{\theta}, \mathbf{v})$ exists then the set of augmented moment conditions is correctly specified.

Throughout the article, we use regression models to understand the various concepts and ideas.

Example 1 (Linear regression model). Suppose that we are interested in estimating the following linear regression model with an intercept and a predictor:

$$y_i = \alpha + \beta z_i + e_i, \quad i = 1, \dots, n, \tag{2.3}$$

where $(z_i, e_i)'$ are independently drawn from some distribution P . Under the assumption that $E^P[e_i|z_i] = 0$, we can use the following moment restrictions to estimate $\boldsymbol{\theta} := (\alpha, \beta)$:

$$E^P[e_i(\boldsymbol{\theta})] = 0, \quad E^P[e_i(\boldsymbol{\theta})z_i] = 0, \quad E^P[(e_i(\boldsymbol{\theta}))^3] = v, \tag{2.4}$$

where $e_i(\boldsymbol{\theta}) := (y_i - \alpha - \beta z_i)$. The first two moment restrictions are derived from the standard orthogonality condition and identify $\boldsymbol{\theta}$. The last restriction potentially serves as additional information. In terms of the notation in (2.1) and (2.2), $\mathbf{x}_i := (y_i, z_i)'$, $\mathbf{g}(\mathbf{x}_i, \boldsymbol{\theta}) = (e_i(\boldsymbol{\theta}), e_i(\boldsymbol{\theta})z_i, e_i(\boldsymbol{\theta})^3)'$, $\mathbf{V} = (0, 0, v)'$, $d_v = 1$, and $\mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{V}) = \mathbf{g}(\mathbf{x}_i, \boldsymbol{\theta}) - (0, 0, v)'$. If one believes that the underlying distribution of e_i is indeed symmetric, then one could use this information by setting v to zero. Otherwise, it is desirable to treat v as an unknown object. If the distribution of e_i is skewed and v is forced to be zero, then the model becomes misspecified because no (α, β) can be consistent with

the three moment restrictions jointly under P . When the augmented parameter \mathbf{v} is treated as a free parameter, the model is correctly specified even under asymmetry.

2.1. Prior-Posterior Analysis

Consider now the question of prior-posterior analysis under the ETEL function. Although our setting is similar to that of Schennach (2005), the presence of the augmented parameter \mathbf{v} and the possibility of misspecification lead to a new analysis and new results.

For any $(\boldsymbol{\theta}, \mathbf{v})$, define the convex hull of $\bigcup_{i=1}^n \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v})$ as the following convex subset of \mathbb{R}^d : $\{\sum_{i=1}^n p_i \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v}); p_i \geq 0, \forall i = 1, \dots, n, \sum_{i=1}^n p_i = 1\}$. Now suppose that (i) $\mathbf{g}^A(\mathbf{x}, \boldsymbol{\theta}, \mathbf{v})$ is continuous in \mathbf{x} for every $(\boldsymbol{\theta}, \mathbf{v}) \in \Theta \times \mathcal{V}$ (or has a finite number of step discontinuities) and (ii) the interior of the convex hull of $\bigcup_{i=1}^n \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v})$ contains the origin. Suppose also that the nonparametric prior on P is the mixture of uniform probability densities described by Schennach (2005), which is capable of approximating any distribution as the number of mixing components increases. Then, adapting the arguments of Schennach (2005), the posterior distribution of $(\boldsymbol{\theta}, \mathbf{v})$ after marginalization over P has the form

$$\pi(\boldsymbol{\theta}, \mathbf{v} | \mathbf{x}_{1:n}) \propto \pi(\boldsymbol{\theta}, \mathbf{v}) p(\mathbf{x}_{1:n} | \boldsymbol{\theta}, \mathbf{v}), \tag{2.5}$$

where $\pi(\boldsymbol{\theta}, \mathbf{v})$ is the prior of $(\boldsymbol{\theta}, \mathbf{v})$ and $p(\mathbf{x}_{1:n} | \boldsymbol{\theta}, \mathbf{v})$ is the ETEL function defined as

$$p(\mathbf{x}_{1:n} | \boldsymbol{\theta}, \mathbf{v}) = \prod_{i=1}^n p_i^*(\boldsymbol{\theta}, \mathbf{v}) \tag{2.6}$$

and $p_i^*(\boldsymbol{\theta}, \mathbf{v})$ are the probabilities that minimize the KL divergence between the probabilities (p_1, \dots, p_n) assigned to each sample observation and the empirical probabilities $(\frac{1}{n}, \dots, \frac{1}{n})$, subject to the conditions that the probabilities (p_1, \dots, p_n) sum to one and that the expectation under these probabilities satisfies the given moment conditions:

$$\max_{p_1, \dots, p_n} \sum_{i=1}^n [-p_i \log(np_i)] \tag{2.7}$$

$$\text{subject to } \sum_{i=1}^n p_i = 1 \quad \text{and} \quad \sum_{i=1}^n p_i \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v}) = \mathbf{0}. \tag{2.8}$$

For numerical and theoretical purposes below, the preceding probabilities are computed more conveniently from the dual (saddlepoint) representation as, for $i = 1, \dots, n$

$$p_i^*(\boldsymbol{\theta}, \mathbf{v}) := \frac{e^{\widehat{\lambda}(\boldsymbol{\theta}, \mathbf{v})' \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v})}}{\sum_{j=1}^n e^{\widehat{\lambda}(\boldsymbol{\theta}, \mathbf{v})' \mathbf{g}^A(\mathbf{x}_j, \boldsymbol{\theta}, \mathbf{v})}},$$

$$\text{where } \widehat{\lambda}(\boldsymbol{\theta}, \mathbf{v}) = \arg \min_{\lambda \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \exp(\lambda' \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v})). \tag{2.9}$$

Therefore, the posterior distribution takes the form

$$\pi(\boldsymbol{\theta}, \mathbf{v} | \mathbf{x}_{1:n}) \propto \pi(\boldsymbol{\theta}, \mathbf{v}) \prod_{i=1}^n \frac{e^{\widehat{\lambda}(\boldsymbol{\theta}, \mathbf{v})' \mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, \mathbf{v})}}{\sum_{j=1}^n e^{\widehat{\lambda}(\boldsymbol{\theta}, \mathbf{v})' \mathbf{g}^A(\mathbf{x}_j, \boldsymbol{\theta}, \mathbf{v})}}, \tag{2.10}$$

which may be called the bayesian exponentially tilted empirical likelihood (BETEL) posterior distribution. It can be efficiently simulated by Markov chain Monte Carlo (MCMC) methods. For example, the one block tailored Metropolis–Hastings (M-H) algorithm (Chib and Greenberg 1995) is applied as follows. Let $q(\boldsymbol{\theta}, \mathbf{v}|\mathbf{x}_{1:n})$ denote a student- t distribution whose location parameter is the mode of the log ETEL function and whose dispersion matrix is the inverse of the negative Hessian matrix of the log ETEL function at the mode. Then, starting from some initial value $(\boldsymbol{\theta}^{(0)}, \mathbf{v}^{(0)})$, we get a sample of draws from the BETEL posterior by repeating the following steps for $s = 1, \dots, S$:

1. Propose $(\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger)$ from $q(\boldsymbol{\theta}, \mathbf{v}|\mathbf{x}_{1:n})$ and solve for $p_i^*(\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger)$, $1 \leq i \leq n$, from the Exponential Tilting saddlepoint problem (2.9).
2. Calculate the M-H probability of move

$$\alpha((\boldsymbol{\theta}^{s-1}, \mathbf{v}^{s-1}), (\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger) | \mathbf{x}_{1:n}) = \min \left\{ 1, \frac{\pi(\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger | \mathbf{x}_{1:n}) q(\boldsymbol{\theta}^{s-1}, \mathbf{v}^{s-1} | \mathbf{x}_{1:n})}{\pi(\boldsymbol{\theta}^{s-1}, \mathbf{v}^{s-1} | \mathbf{x}_{1:n}) q(\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger | \mathbf{x}_{1:n})} \right\}.$$

3. Set $(\boldsymbol{\theta}^s, \mathbf{v}^s) = (\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger)$ with probability $\alpha((\boldsymbol{\theta}^{s-1}, \mathbf{v}^{s-1}), (\boldsymbol{\theta}^\dagger, \mathbf{v}^\dagger) | \mathbf{x}_{1:n})$. Otherwise, set $(\boldsymbol{\theta}^s, \mathbf{v}^s) = (\boldsymbol{\theta}^{s-1}, \mathbf{v}^{s-1})$. Go to step 1.

Note that when the dimension of $(\boldsymbol{\theta}, \mathbf{v})$ is large, the tailored randomized block M-H algorithm of Chib and Ramamurthy (2010) can be used instead for improved simulation efficiency.

Prior specification. In our examples, we focus on two prior distributions. Under the first prior, which we call the default prior, each element θ_k and v_l of $\boldsymbol{\theta}$ and \mathbf{v} , respectively, is given independent student- t distributions with $\nu = 2.5$ degrees of freedom, location zero, and dispersion equal to 5:

$$\theta_k \sim t_{2.5}(0, 5^2) \quad \text{and} \quad v_l \sim t_{2.5}(0, 5^2). \tag{2.11}$$

In the second prior, which we call the training sample prior, an initial portion of the sample (which is not used for subsequent inferences) is used to find the ETEL estimate of the unknown parameters, that is, the maximizer of the ETEL function (2.6) whose definition is recalled in (A.1) in the Appendix. Then, the prior of each element of $(\boldsymbol{\theta}', \mathbf{v}')$ is equal to the default prior except that now the location is set equal to the corresponding ETEL estimate.

To see the different implications of these prior distributions, consider two moment condition models defined by the restrictions:

$$\begin{aligned} M_1: & \mathbf{E}^P[g_1(\mathbf{X}, \boldsymbol{\theta})] = 0, \quad \mathbf{E}^P[g_2(\mathbf{X}, \boldsymbol{\theta})] = 0 \\ M_2: & \mathbf{E}^P[g_1(\mathbf{X}, \boldsymbol{\theta})] = 0, \quad \mathbf{E}^P[g_2(\mathbf{X}, \boldsymbol{\theta})] = v, \end{aligned} \tag{2.12}$$

where both moments restrictions are active under M_1 but only the first is active under M_2 . Then, under the default prior, a prior mean of 0 on v implies the belief that the second moment restriction is likely to hold. On the other hand, in the training sample prior, the prior location of v is determined by the ETEL estimate of v in the training sample. If this is substantially different from zero (relative to the prior dispersion), this prior implies the belief that the second moment restriction is, a priori, less likely to be active.

Table 1. Posterior summary for two simulated sample sizes from Example 1 (a regression model with skewed error distribution). The true value of α is 0 and that of β is 1. The summaries are based on 10,000 MCMC draws beyond a burn-in of 1000. The M-H acceptance rate is around 90% in both cases. “Lower” and “upper” refer to the 0.025 and 0.975 quantiles of the simulated draws, respectively, and “ineff” to the inefficiency factor, the ratio of the numerical variance of the mean to the variance of the mean assuming independent draws: an inefficiency factor close to 1 indicates that the MCMC draws, although serially correlated, are essentially independent.

	mean	sd	median	lower	upper	ineff
$n = 250$						
α	-0.03	0.10	-0.03	-0.24	0.16	1.49
β	0.99	0.08	0.98	0.83	1.15	1.70
v	-1.42	0.36	-1.39	-2.20	-0.82	3.21
$n = 2000$						
α	0.00	0.03	0.00	-0.06	0.06	1.30
β	0.99	0.03	0.99	0.93	1.04	1.21
v	-0.97	0.10	-0.97	-1.18	-0.78	1.34

Example 1 (continued). To illustrate the prior-posterior analysis, we generate y_i , $i = 1, \dots, n$ from the regression model in (2.3) with the covariate $z_i \sim \mathcal{N}(0.5, 1)$, intercept $\alpha = 0$, slope $\beta = 1$ and e_i distributed according to the skewed distribution:

$$e_i \sim \begin{cases} \mathcal{N}(0.75, 0.75^2) & \text{with probability 0.5} \\ \mathcal{N}(-0.75, 1.25^2) & \text{with probability 0.5.} \end{cases} \tag{2.13}$$

Our analysis is based on the moment restrictions in (2.4), that is,

$$\mathbf{g}^A(\mathbf{x}_i, \boldsymbol{\theta}, v) = (e_i(\boldsymbol{\theta}), e_i(\boldsymbol{\theta})z_i, e_i(\boldsymbol{\theta})^3 - v)', \quad e_i(\boldsymbol{\theta}) = y_i - \alpha - \beta z_i,$$

with $\boldsymbol{\theta} = (\alpha, \beta)$. These moment conditions are correctly specified because v is free. Under the default independent student- t prior in (2.11), the marginal posterior distributions of α , β , and v are summarized in Table 1 for two different values of n . It can be seen from the 0.025 and 0.975 quantiles (called “lower” and “upper,” respectively) that the marginal posterior distributions of α and β are already concentrated around the true values for $n = 250$ but concentrate even more closely around the true values for $n = 2000$. This example showcases the ease with which such Bayesian inferences are possible.

Notation. In Sections 2.2 and 2.3, and in the online Appendix, we use the following notations. For ease of exposition, we denote $\boldsymbol{\psi} := (\boldsymbol{\theta}, \mathbf{v})$, $\boldsymbol{\psi} \in \Psi$ with $\Psi := \Theta \times \mathcal{V}$. Moreover, $\|\cdot\|_F$ denotes the Frobenius norm and $\|\cdot\|$ the Euclidean norm. The notation \xrightarrow{P} is for convergence in probability with respect to the product measure $P^n = \bigotimes_{i=1}^n P$. The log-likelihood function for one observation is denoted by $l_{n,\boldsymbol{\psi}}$:

$$\begin{aligned} l_{n,\boldsymbol{\psi}}(\mathbf{x}) &:= \log \frac{e^{\widehat{\lambda}(\boldsymbol{\psi})' \mathbf{g}^A(\mathbf{x}, \boldsymbol{\psi})}}{\sum_{j=1}^n e^{\widehat{\lambda}(\boldsymbol{\psi})' \mathbf{g}^A(\mathbf{x}_j, \boldsymbol{\psi})}} \\ &= -\log n + \log \frac{e^{\widehat{\lambda}' \mathbf{g}^A(\mathbf{x}, \boldsymbol{\psi})}}{\frac{1}{n} \sum_{j=1}^n e^{\widehat{\lambda}' \mathbf{g}^A(\mathbf{x}_j, \boldsymbol{\psi})}} \end{aligned}$$

so that the log-ETEL function is $\log p(\mathbf{x}_{1:n} | \boldsymbol{\psi}) = \sum_{i=1}^n l_{n,\boldsymbol{\psi}}(\mathbf{x}_i)$. For a set $\mathcal{A} \subset \mathbb{R}^m$, we denote by $\text{int}(\mathcal{A})$ its interior relative to \mathbb{R}^m . Further notations are introduced as required.

2.2. Asymptotic Properties: Correct Specification

In this section, we first introduce additional notations and assumptions for correctly specified models. Under these assumptions and Assumptions 5–6 in the online Appendix, we establish both the large sample behavior of the BETEL posterior distribution and, in Section 3, the model selection consistency of our marginal likelihood procedure.

Let θ_* be the true value of the parameter of interest θ and \mathbf{v}_* be the true value of the augmented parameter. So, $\boldsymbol{\psi}_* := (\theta_*, \mathbf{v}_*)$. The true value \mathbf{v}_* is equal to zero when the nonaugmented model (2.1) is correctly specified. Moreover, let $\Delta := \mathbf{E}^P[\mathbf{g}^A(X, \boldsymbol{\psi}_*)\mathbf{g}^A(X, \boldsymbol{\psi}_*)']$ and $\Gamma := \mathbf{E}^P[\frac{\partial}{\partial \boldsymbol{\psi}} \mathbf{g}^A(X, \boldsymbol{\psi}_*)]$. Assumption 1 requires that the augmented model is correctly specified in the sense that there is a value of $\boldsymbol{\psi}$ such that (2.2) is satisfied by P , and that this value is unique. A necessary condition for the latter is that $(d - p) \geq d_v \geq 0$.

Assumption 1. Model (2.2) is such that $\boldsymbol{\psi}_* \in \Psi$ is the unique solution to $\mathbf{E}^P[\mathbf{g}^A(X, \boldsymbol{\psi})] = \mathbf{0}$.

The next assumption concerns the prior distribution and is a standard assumption to establish asymptotic properties of Bayesian procedures.

Assumption 2. (a) π is a continuous probability measure that admits a density with respect to the Lebesgue measure; (b) π is positive on a neighborhood of $\boldsymbol{\psi}_*$.

For a correctly specified moment conditions model, the asymptotic normality of the BETEL posterior is established in the following theorem where we denote by $\pi(\sqrt{n}(\boldsymbol{\psi} - \boldsymbol{\psi}_*)|\mathbf{x}_{1:n})$ the posterior distribution of $\sqrt{n}(\boldsymbol{\psi} - \boldsymbol{\psi}_*)$. The result shows that the BETEL posterior distribution has a Gaussian limiting distribution and that it concentrates on a $n^{-1/2}$ -ball centered at the true value of the parameter. An informal discussion of this behavior is given by Schennach (2005) but without the required assumptions. Theorem 1 provides these assumptions. The proof of the result is based on, for example, Lehmann and Casella (1998) and Ghosh and Ramamoorthi (2003) and is given in the online Appendix C.

Theorem 1 (Bernstein–von Mises—correct specification). Under Assumptions 1, 2 and Assumptions 5, 6 in the online Appendix and if in addition, for any $\delta > 0$, there exists an $\epsilon > 0$ such that, as $n \rightarrow \infty$

$$P \left(\sup_{\|\boldsymbol{\psi} - \boldsymbol{\psi}_*\| > \delta} \frac{1}{n} \sum_{i=1}^n (l_{n,\boldsymbol{\psi}}(\mathbf{x}_i) - l_{n,\boldsymbol{\psi}_*}(\mathbf{x}_i)) \leq -\epsilon \right) \rightarrow 1, \quad (2.14)$$

then the posteriors converge in total variation toward a normal distribution, that is,

$$\sup_B \left| \pi(\sqrt{n}(\boldsymbol{\psi} - \boldsymbol{\psi}_*) \in B|\mathbf{x}_{1:n}) - \mathcal{N}_{0,(\Gamma'\Delta^{-1}\Gamma)^{-1}}(B) \right| \xrightarrow{P} 0, \quad (2.15)$$

where $B \subseteq \Psi$ is any Borel set.

According to this result, the posterior distribution $\pi(\boldsymbol{\psi}|\mathbf{x}_{1:n})$ of $\boldsymbol{\psi}$ is asymptotically normal, centered on the true value $\boldsymbol{\psi}_*$ and with variance $n^{-1}(\Gamma'\Delta^{-1}\Gamma)^{-1}$. Thus, the posterior distribution has the same asymptotic variance as the efficient Generalized Method of Moments estimator of Hansen (1982; see also

Chamberlain 1987). Assumption (2.14) in this theorem is a standard identifiability condition (see, e.g., Lehmann and Casella 1998, Assumption 6.B.3) that controls the behavior of the log-ETEL function at a distance from $\boldsymbol{\psi}_*$. Controlling this behavior is important because the posterior involves integration over the whole range of $\boldsymbol{\psi}$. To understand the meaning of this assumption, we remark that asymptotically the log-ETEL function $\boldsymbol{\psi} \mapsto \sum_{i=1}^n l_{n,\boldsymbol{\psi}}(\mathbf{x}_i)$ is maximized at the true value $\boldsymbol{\psi}_*$ because the model is correctly specified. Hence, Assumption (2.14) means that if the parameter $\boldsymbol{\psi}$ is “far” from the true value $\boldsymbol{\psi}_*$, then the log-ETEL function has to be small, that is, has to be far from the maximum value $\sum_{i=1}^n l_{n,\boldsymbol{\psi}_*}(\mathbf{x}_i)$.

2.3. Asymptotic Properties: Misspecification

In this section, we consider the case where the model is misspecified in the sense of Definition 1 and establish that, even in this case, the BETEL posterior distribution has good frequentist asymptotic properties as the sample size n increases. Namely, we show that the BETEL posterior of $\sqrt{n}(\boldsymbol{\psi} - \boldsymbol{\psi}_*)$ is asymptotically normal and the BETEL posterior of $\boldsymbol{\psi}$ concentrates on an $n^{-1/2}$ -ball centered at the pseudo-true value of the parameter. To the best of our knowledge, these properties have not been established yet for misspecified moment condition models.

Because in misspecified models there is no value of $\boldsymbol{\psi}$ for which the true data distribution P satisfies the restriction (2.2), we need to define a pseudo-true value for $\boldsymbol{\psi}$. The latter is defined as the value of $\boldsymbol{\psi}$ that minimizes the KL divergence $K(P||Q^*(\boldsymbol{\psi}))$ between the true data distribution P and a distribution $Q^*(\boldsymbol{\psi})$ defined as $Q^*(\boldsymbol{\psi}) := \operatorname{arginf}_{Q \in \mathcal{P}_\Psi} K(Q||P)$, where $K(Q||P) := \int \log(dQ/dP)dQ$ and \mathcal{P}_Ψ is defined in Definition 1. We remark that these two KL divergences are the population counterparts of the KL divergences used for the definition of the ETEL function in (2.6): the empirical counterpart of $K(Q||P)$ is used to construct the $p_i^*(\boldsymbol{\psi})$ probabilities and is given by (2.7), while the empirical counterpart of $K(P||Q^*(\boldsymbol{\psi}))$ is given by $\log(1/n) - \sum_{i=1}^n l_{n,\boldsymbol{\psi}}(\mathbf{x}_i)/n$, where $\sum_{i=1}^n l_{n,\boldsymbol{\psi}}(\mathbf{x}_i)$ is the log-ETEL function if the dual theorem holds. Roughly speaking, the pseudo-true value is the value of $\boldsymbol{\psi}$ for which the distribution that satisfies the corresponding restrictions (2.2) is the closest to the true P , in the KL sense. By using the dual representation of the KL minimization problem, the P -density $dQ^*(\boldsymbol{\psi})/dP$ admits a closed-form: $dQ^*(\boldsymbol{\psi})/dP = e^{\lambda_\circ(\boldsymbol{\psi})' \mathbf{g}^A(X, \boldsymbol{\psi})} / \mathbf{E}^P[e^{\lambda_\circ(\boldsymbol{\psi})' \mathbf{g}^A(X, \boldsymbol{\psi})}]$ where $\lambda_\circ(\boldsymbol{\psi})$ is the pseudo-true value of the tilting parameter defined as the solution of $\mathbf{E}^P[\exp\{\lambda' \mathbf{g}^A(X, \boldsymbol{\psi})\} \mathbf{g}^A(X, \boldsymbol{\psi})] = \mathbf{0}$ which is unique by the strict convexity of $\mathbf{E}^P[\exp\{\lambda' \mathbf{g}^A(X, \boldsymbol{\psi})\}]$ in λ . Therefore,

$$\begin{aligned} \lambda_\circ(\boldsymbol{\psi}) &:= \operatorname{arg min}_{\lambda \in \mathbb{R}^d} \mathbf{E}^P \left[e^{\lambda' \mathbf{g}^A(X, \boldsymbol{\psi})} \right], \\ \boldsymbol{\psi}_\circ &:= \operatorname{arg max}_{\boldsymbol{\psi} \in \Psi} \mathbf{E}^P \log \left[\frac{e^{\lambda_\circ(\boldsymbol{\psi})' \mathbf{g}^A(X, \boldsymbol{\psi})}}{\mathbf{E}^P [e^{\lambda_\circ(\boldsymbol{\psi})' \mathbf{g}^A(X, \boldsymbol{\psi})}]} \right]. \end{aligned} \quad (2.16)$$

However, in a misspecified model, the dual theorem is not guaranteed to hold and so $\boldsymbol{\psi}_\circ$ defined in (2.16) is not necessarily equal to the pseudo-true value defined as the KL-minimizer. In fact, when the model is misspecified, the probability measures in $\mathcal{P} := \bigcup_{\boldsymbol{\psi} \in \Psi} \mathcal{P}_\boldsymbol{\psi}$, which are implied by the model, might not have a common support with the true P , see Sueishi (2013) for a discussion on this point. Following Sueishi (2013, Theorem 3.1), to

guarantee identification of the pseudo-true value by (2.16) and validity of the dual theorem we introduce the following assumption. This assumption replaces Assumption 1 in misspecified models.

Assumption 3. For a fixed $\psi \in \Psi$, there exists $Q \in \mathcal{P}_\psi$ such that Q is mutually absolutely continuous with respect to P , where \mathcal{P}_ψ is defined in Definition 1.

This assumption implies that \mathcal{P}_ψ is nonempty. A similar assumption is also made by Kleijn and van der Vaart (2012) to establish the BvM under misspecification. Moreover, because consistency in misspecified models is defined with respect to the pseudo-true value ψ_\circ , we need to replace Assumption 2 (b) by the following assumption which, together with Assumption 2 (a), requires the prior to put enough mass to balls around ψ_\circ .

Assumption 4. The prior distribution π is positive on a neighborhood of ψ_\circ where ψ_\circ is as defined in (2.16).

A first step to establish the BvM theorem is to prove that the misspecified model satisfies a stochastic local asymptotic normality (LAN) expansion around the pseudo-true value ψ_\circ . Namely, that the log-likelihood ratio $l_{n,\psi} - l_{n,\psi_\circ}$, evaluated at a local parameter around the pseudo-true value, is well approximated by a quadratic form. Such a result is established in Theorem C.1 in the online Appendix C. A second key ingredient for establishing the BvM theorem is the requirement that, as $n \rightarrow \infty$, the posterior of ψ concentrates and puts all its mass on $\Psi_n := \{\|\psi - \psi_\circ\| \leq M_n/\sqrt{n}\}$, where M_n is any sequence such that $M_n \rightarrow \infty$. We prove this result in Theorem C.2 in the online Appendix C.

Theorem 2 states that the limit of the posterior distribution of $\sqrt{n}(\psi - \psi_\circ)$ is a Gaussian distribution with mean and variance defined in terms of the population counterpart of $l_{n,\psi}(\mathbf{x})$, which we denote by $\mathcal{L}_{n,\psi}(\mathbf{x}) := \log \frac{\exp(\lambda_\circ(\psi)'g^A(\mathbf{x},\psi))}{E^P[\exp(\lambda_\circ(\psi)'g^A(\mathbf{x},\psi))]} - \log n$ and which involves the pseudo-true value λ_\circ . With this notation, the variance and mean of the Gaussian limiting distribution are $V_{\psi_\circ}^{-1} := -(\mathbf{E}^P[\dot{\mathcal{L}}_{n,\psi_\circ}])^{-1}$ and $\Delta_{n,\psi_\circ} := \frac{1}{\sqrt{n}} \sum_{i=1}^n V_{\psi_\circ}^{-1} \dot{\mathcal{L}}_{n,\psi_\circ}(\mathbf{x}_i)$, respectively, where $\dot{\mathcal{L}}_{n,\psi_\circ}$ and $\ddot{\mathcal{L}}_{n,\psi_\circ}$ denote the first and second derivatives of the function $\psi \mapsto \mathcal{L}_{n,\psi}$ evaluated at ψ_\circ . Let $\pi(\sqrt{n}(\psi - \psi_\circ) | \mathbf{x}_{1:n})$ denote the posterior distribution of $\sqrt{n}(\psi - \psi_\circ)$.

Theorem 2 (Bernstein–von Mises—misspecification). Assume that the matrix V_{ψ_\circ} is nonsingular and that Assumptions 2 (a), 3, 4 and Assumptions 5 (a)–(d), 6 (b), 7, and 8 in the online Appendix hold. If in addition there exists a constant $C > 0$ such that for any sequence $M_n \rightarrow \infty$, as $n \rightarrow \infty$

$$P \left(\sup_{\psi \in \Psi_n^c} \frac{1}{n} \sum_{i=1}^n (l_{n,\psi}(\mathbf{x}_i) - l_{n,\psi_\circ}(\mathbf{x}_i)) \leq -\frac{CM_n^2}{n} \right) \rightarrow 1, \quad (2.17)$$

then the posteriors converge in total variation toward a normal distribution, that is,

$$\sup_B \left| \pi(\sqrt{n}(\psi - \psi_\circ) \in B | \mathbf{x}_{1:n}) - \mathcal{N}_{\Delta_{n,\psi_\circ}, V_{\psi_\circ}^{-1}}(B) \right| \xrightarrow{P} 0, \quad (2.18)$$

where $B \subseteq \Psi$ is any Borel set.

Condition (2.17) involves the log-likelihood ratio $l_{n,\psi}(\mathbf{x}) - l_{n,\psi_\circ}(\mathbf{x})$ and is an identifiability condition, standard in the literature, and with a similar interpretation as condition (2.14). Theorem 2 states that, in misspecified models, the sequence of posterior distributions converges in total variation to a sequence of normal distributions with random mean and fixed covariance matrix $V_{\psi_\circ}^{-1}$. By using the first-order condition for ψ_\circ it can be shown that the random mean Δ_{n,ψ_\circ} has mean zero. We stress that the BvM result of Theorem 2 for the BETEL posterior distribution does not directly follow from the assumptions and results in Kleijn and van der Vaart (2012) because the ETEL function contains random quantities. Therefore, we need to strengthen the assumptions to establish that a stochastic LAN expansion holds for our case.

As the next lemma shows, the quantity Δ_{n,ψ_\circ} relates to the Schennach’s (2007) ETEL frequentist estimator $\hat{\psi}$ (whose definition is recalled in (A.1) in the Appendix for convenience). Because of this connection, it is possible to write the location of the normal limit distribution in a more familiar form in terms of the semiparametric efficient frequentist estimator $\hat{\psi}$.

Lemma 1. Assume that the matrix V_{ψ_\circ} is nonsingular and that Assumption 3 and Assumptions 5 (a)–(d), 6 (b), 7, and 8 in the online Appendix hold. Then, the ETEL estimator $\hat{\psi}$ satisfies

$$\sqrt{n}(\hat{\psi} - \psi_\circ) = \frac{1}{\sqrt{n}} \sum_{i=1}^n V_{\psi_\circ}^{-1} \dot{\mathcal{L}}_{n,\psi_\circ} + o_P(1). \quad (2.19)$$

Therefore, Lemma 1 implies that the BvM Theorem 2 can be reformulated with the sequence $\sqrt{n}(\hat{\psi} - \psi_\circ)$ as the location of the normal limit distribution, that is,

$$\sup_B \left| \pi(\hat{\psi} \in B | \mathbf{x}_{1:n}) - \mathcal{N}_{\hat{\Delta}_{n,\psi_\circ}, V_{\psi_\circ}^{-1}}(B) \right| \xrightarrow{P} 0. \quad (2.20)$$

Two remarks are in order: (I) the limit distribution of $\sqrt{n}(\hat{\psi} - \psi_\circ)$ is centered on zero because $E^P[\dot{\mathcal{L}}_{n,\psi_\circ}] = 0$; (II) the asymptotic covariance matrix of $\sqrt{n}(\hat{\psi} - \psi_\circ)$ is $V_{\psi_\circ}^{-1} E^P[\dot{\mathcal{L}}_{n,\psi_\circ} \dot{\mathcal{L}}_{n,\psi_\circ}'] V_{\psi_\circ}^{-1}$ (which is also derived by Schennach 2007, Theorem 10) and, because of misspecification, it does not coincide with the limiting covariance matrix in the BvM theorem. This consequence of misspecification is also discussed in Kleijn and van der Vaart (2012) and implies that, for $\alpha \in (0, 1)$, the central $(1 - \alpha)$ Bayesian credible sets are not in general $(1 - \alpha)$ confidence sets, even asymptotically. In fact, while credible sets are correctly centered, their width/volume need not be correct since the asymptotic variance matrix in the BvM is not the sandwich asymptotic covariance matrix of the frequentist estimator. See Example 2 in the online Appendix A for an illustration of misspecified models and pseudo-true value.

3. Bayesian Model Selection

3.1. Basic Idea

Now suppose that there are countable candidate models indexed by ℓ . Suppose that model ℓ is characterized by

$$E^P[g^\ell(X, \theta^\ell)] = \mathbf{0}, \quad (3.1)$$

with $\theta^\ell \in \Theta^\ell \subset \mathbb{R}^{p_\ell}$, and $\ell = 1, \dots, J$ for some $J \geq 2$. Different models involve different parameters of interest θ^ℓ and/or different g^ℓ functions. If X contains a dependent variable and covariates it might be that the covariates are not the same for all models, however to lighten the notation we do not explicit this difference across the models.

One or all models may be misspecified. The goal is to compare these models and select the best model. By best model we mean the model that contains the maximum number of over-identifying conditions when all models are correctly specified, and when all models are misspecified, we mean the model that is the closest to the true P . Our purpose in this section is to establish a collection of results on the search for such a best model. We show that this search can be carried out with the help of the marginal likelihoods (defined as the integral of the sampling density over the parameters with respect to the prior density) of the competing models. The model with the largest marginal likelihood satisfies a model selection consistency property in that the model chosen in this way is the best model asymptotically. This property, which has not been established in this context before, is of enormous practical and theoretical importance.

Before getting to the details, it is crucial to understand that there are some subtleties involved in comparing different moment condition models. The central problem is that the marginal likelihood of models with different sets of moment restrictions and different parameters may not be comparable. In fact, when we have different sets of moment restrictions, we need to be careful about dealing with, and interpreting, unused moment restrictions. This can be best explained by an example.

Example 1 (continued). Suppose we do not know if e_i is symmetric. In this case, one might be inclined to compare the following two candidate models:

$$\begin{aligned} \text{Model 1: } & \mathbf{E}^P[e_i(\theta)] = 0, \quad \mathbf{E}^P[e_i(\theta)z_i] = 0. \\ \text{Model 2: } & \mathbf{E}^P[e_i(\theta)] = 0, \quad \mathbf{E}^P[e_i(\theta)z_i] = 0, \quad \mathbf{E}^P[(e_i(\theta))^3] = 0, \end{aligned} \tag{3.2}$$

where $\theta := (\alpha, \beta)$ is the same parameter in the two models and $e_i(\theta) := (y_i - \alpha - \beta z_i)$. As written, these two models are not comparable because the convex hulls associated with the two models do not have the same dimension. More precisely, let $co_1 := \{\sum_{i=1}^n p_i(e_i(\theta), e_i(\theta)z_i)'; p_i \geq 0, \forall i = 1, \dots, n, \sum_{i=1}^n p_i = 1\}$ be the convex hull associated with Model 1, and $co_2 := \{\sum_{i=1}^n p_i(e_i(\theta), e_i(\theta)z_i, e_i(\theta)^3)'; p_i \geq 0, \forall i = 1, \dots, n, \sum_{i=1}^n p_i = 1\}$ be the convex hull associated with Model 2. Because co_1 and co_2 have different dimensions, the $p_i^*(\theta)$ in the two ETEL functions are not comparable because they enforce the zero vector constraint (the second constraint in (2.8)) in different spaces (\mathbb{R}^2 and \mathbb{R}^3).

The foregoing problem can be overcome as follows. We start by defining a grand model that nests all the models that we want to compare. This grand model is constructed such that: (1) it includes all the moment restrictions in the models and, (2) if the same moment restriction is included in two or more models but involves a different parameter in different models, then the grand model includes the moment restriction that involves the parameter of largest dimension. We write the grand model as $\mathbf{E}^P[g^G(X, \theta^G)] = \mathbf{0}$ where g^G has dimension d , and θ^G includes the parameters of all models. Next, each original model

is obtained from this grand model by first subtracting a vector of nuisance parameters V and then restricting θ^G and V appropriately. More precisely, an equivalent version of the original model is obtained by: (I) setting equal to zero the components of θ^G in the shared moment restrictions that are not present in the original model, (II) letting free the components of V that correspond to the over-identifying moment restrictions not present in the original model, and (III) setting equal to zero the components of V that correspond to moment restrictions present in the original model and to moment restrictions that exactly identify the extra parameters arising from the other models.

The set of models that emerge from this strategy are equivalent to the original collection but, crucially, are now comparable. Also importantly for practice, as our illustrations of this strategy show below, the strategy just described is simple to operationalize. With this formulation, model ℓ , denoted by M_ℓ , is then defined from the grand model as

$$\mathbf{E}^P[g^A(X, \theta^\ell, v^\ell)] = 0, \quad \theta^\ell \in \Theta^\ell \subset \mathbb{R}^{p_\ell}, \tag{3.3}$$

where $g^A(X, \theta^\ell, v^\ell) := g^G(X, \theta^\ell) - V^\ell$ with $V^\ell \in \mathfrak{V} \subset \mathbb{R}^d$ and with $v^\ell \in \mathcal{V}^\ell \subset \mathbb{R}^{d_{v_\ell}}$ being the vector that collects all the nonzero components of V^ℓ . We assume that $0 \leq d_{v_\ell} \leq d - p_\ell$ to guarantee identification of θ^ℓ . The parameter v^ℓ is the augmented parameter and θ^ℓ is the parameter of interest for model ℓ that has been obtained from θ^G by doing the transformation in (I). Hereafter, we use the notation $\psi^\ell := (\theta^\ell, v^\ell) \in \Psi^\ell$ with $\Psi^\ell := \Theta \times \mathcal{V}^\ell$.

Example 1 (continued). To be able to compare Model 1 and Model 2 in (3.2), we construct the grand model as $\mathbf{E}^P[g^G(x_i, \theta)] := \mathbf{E}^P[(e_i(\theta), e_i(\theta)z_i, e_i(\theta)^3)']$. With respect to this grand model, Model 1 and Model 2 are reformulated as M_1 and M_2 , respectively, by applying (II) and (III) above:

$$\begin{aligned} M_1: & \mathbf{E}^P[e_i(\theta)] = 0, \quad \mathbf{E}^P[e_i(\theta)z_i] = 0, \quad \mathbf{E}^P[(e_i(\theta))^3] = v \\ M_2: & \mathbf{E}^P[e_i(\theta)] = 0, \quad \mathbf{E}^P[e_i(\theta)z_i] = 0, \quad \mathbf{E}^P[(e_i(\theta))^3] = 0. \end{aligned} \tag{3.4}$$

So, $\psi^1 = (\theta', v)'$ and $\psi^2 = \theta$. The convex hulls of M_1 and M_2 have both dimension 3 and more importantly, if $co(M_1)$ and $co(M_2)$ denote these two convex hulls, we have $co(M_1) = co(M_2) - V$ where $V = (0, 0, v)'$ so that the two models are comparable. It is important to note how Model 1 in (3.2) and M_1 deal with uncertainty about the third moment restriction: Model 1 in (3.2) ignores its uncertainty completely while M_1 models the degree of uncertainty through the augmented parameter v . This argument is not limited to comparing two models. When we have multiple models, we need to make sure that the grand model encompasses all candidate models through augmented parameters.

We note that this strategy covers both nested and nonnested models. We say that two models are nonnested, in their original formulation, if neither model can be obtained from the other by eliminating some moment restriction, or by setting to zero some parameter, or both. Points (II) and (III) above are important for the treatment of such nonnested models. In fact, points (II) and (III) imply that, if there are moment restrictions not present in the original model that involve parameters that are not in the original model, then a number of these extra moment restrictions equal to the number of the extra parameters has to

be included. This does not alter the original model if these extra moment restrictions exactly identify the extra parameters and so place no restrictions on the data-generating process. Moreover, despite the notation, for nonnested models θ^ℓ in (3.3) might be larger than the parameter in the original model ℓ .

In what follows, we show how to compute the marginal likelihood for a model. Then, in Section 3.3 we formally show that, with probability approaching one as the number of observations increases, the marginal likelihood-based selection procedure favors the model with the minimum number of parameters of interest and the maximum number of valid moment restrictions. We also consider the situation where all models are misspecified. In this case, our model selection procedure selects the model that is closer to the true data-generating process in terms of the KL divergence.

3.2. Marginal Likelihood

For each model M_ℓ , we impose a prior distribution for ψ^ℓ on Ψ^ℓ , and obtain the BETEL posterior distribution based on (2.10). Then, we select the model with the largest marginal likelihood, denoted by $m(\mathbf{x}_{1:n}; M_\ell)$, which we calculate by the method of Chib (1995) as extended to Metropolis–Hastings samplers in Chib and Jeliazkov (2001). This method makes computation of the marginal likelihood simple and is a key feature of our procedure. The main advantage of the Chib (1995) method is that it is calculable from the same inputs and outputs that are used in the MCMC sampling of the posterior distribution. The starting point of this method is the following identity of the log-marginal likelihood introduced in Chib (1995):

$$\log m(\mathbf{x}_{1:n}; M_\ell) = \log \pi(\tilde{\psi}^\ell | M_\ell) + \log p(\mathbf{x}_{1:n} | \tilde{\psi}^\ell, M_\ell) - \log \pi(\tilde{\psi}^\ell | \mathbf{x}_{1:n}, M_\ell), \tag{3.5}$$

where $\tilde{\psi}^\ell$ is any point in the support of the posterior (such as the posterior mean) and the dependence on the model M_ℓ has been made explicit. The first two terms on the right-hand side of this decomposition are available directly whereas the third term can be estimated from the output of the MCMC simulation of the BETEL posterior distribution. For example, in the context of the one block MCMC algorithm given in Section 2.1, from Chib and Jeliazkov (2001), we have that

$$\pi(\tilde{\psi}^\ell | \mathbf{x}_{1:n}, M_\ell) = \frac{\mathbf{E}_1 \left\{ \alpha(\psi^\ell, \tilde{\psi}^\ell | \mathbf{x}_{1:n}, M_\ell) q(\tilde{\psi}^\ell | \mathbf{x}_{1:n}, M_\ell) \right\}}{\mathbf{E}_2 \left\{ \alpha(\tilde{\psi}^\ell, \psi^\ell | \mathbf{x}_{1:n}, M_\ell) \right\}},$$

where \mathbf{E}_1 is the expectation with respect to $\pi(\psi^\ell | \mathbf{x}_{1:n}, M_\ell)$ and \mathbf{E}_2 is the expectation with respect to $q(\psi^\ell | \mathbf{x}_{1:n}, M_\ell)$. These expectations can be easily approximated by simulations.

3.3. Model Selection Consistency Results

In this section, we establish the consistency of our marginal likelihood-based selection procedure for the following exhaustive cases: the case where the models in the comparison set contain only valid moment restrictions, the case where all the models in the set are misspecified, and finally the case where some of the models contain only valid moment restrictions while the others contain at least one invalid moment restriction. Our

proofs of consistency are based on: (I) the results of the BvM theorems for correctly and misspecified models stated in Sections 2.2 and 2.3, and (II) the analysis of the asymptotic behavior of the ETEL function under correct and misspecification which we develop in the online Appendix (see Lemmas D.1 and D.3).

The first theorem states that, if the active moment restrictions are all valid, then the marginal likelihood selects the model that contains the maximum number of overidentifying conditions, that is, the model with the maximum number of active moment restrictions and the smallest number of parameters of interest. This means that the marginal likelihood-based selection procedure enforces parsimony.

For a model M_ℓ , the dimension of the parameter of interest θ^ℓ to be estimated is p_ℓ while the number of active moment restrictions (included in the model for the estimation of θ^ℓ) is $(d - d_{v_\ell})$. Consider two generic models M_1 and M_2 . Then, $d_{v_2} < d_{v_1}$ means that model M_2 contains more active restrictions than model M_1 , and $p_2 < p_1$ means that model M_1 contains more parameters of interest to be estimated than M_2 .

Theorem 3. Let Assumption 2, Assumptions 5, 6 in the online Appendix, and (2.14) hold, and consider $J < \infty$ different models M_ℓ , $\ell = 1, \dots, J$, that satisfy Assumption 1, that is, they are all correctly specified. Then,

$$\lim_{n \rightarrow \infty} P \left(\max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell) < \log m(\mathbf{x}_{1:n}; M_j) \right) = 1$$

if and only if $p_j + d_{v_j} < p_\ell + d_{v_\ell}, \forall \ell \neq j$.

The result of the theorem implies that, with probability approaching 1, the Bayes factor $B_{j\ell} := m(\mathbf{x}_{1:n}; M_j) / m(\mathbf{x}_{1:n}; M_\ell)$ is larger than 1 for every $\ell \neq j$. The result in the theorem is an equivalence result saying that, if we compare models that contain only valid moment restrictions, then the marginal likelihood selects a model M_j if and only if M_j contains the maximum number of overidentifying conditions among all the compared models. An illustration of this theorem is provided in Example 3 in the online Appendix.

Next, we consider the case where all models are wrong in the sense of Definition 1 and establish a major result of enormous practical significance. The result states that if we compare J misspecified models, then the marginal likelihood-based selection procedure selects the model with the smallest KL divergence $K(P || Q^*(\psi^\ell))$ between P and $Q^*(\psi^\ell)$, where $Q^*(\psi^\ell)$ is such that $K(Q^*(\psi^\ell) || P) = \inf_{Q \in \mathcal{P}_{\psi^\ell}} K(Q || P)$ and $dQ^*(\psi) / dP = e^{\lambda_\circ(\psi)' g^A(X, \psi)} / \mathbf{E}^P [e^{\lambda_\circ(\psi)' g^A(X, \psi)}]$ by the dual theorem, as defined in Section 2.3. Because the I-projection $Q^*(\psi^\ell)$ on \mathcal{P}_{ψ^ℓ} is unique (Csiszar 1975), which $Q^*(\psi^\ell)$ is closer to P (in terms of $K(P || Q^*(\psi^\ell))$) depends only on the “amount of misspecification” contained in each model \mathcal{P}_{ψ^ℓ} .

Theorem 4. Let Assumptions 2–4, Assumptions 5–8 in the online Appendix and (2.17) be satisfied. Let us consider the comparison of $J < \infty$ models M_j , $j = 1, \dots, J$ that all use misspecified moments, that is, M_j does not satisfy Assumption 1, $\forall j$. Then,

$$\lim_{n \rightarrow \infty} P \left(\log m(\mathbf{x}_{1:n}; M_j) > \max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell) \right) = 1$$

if and only if $K(P||Q^*(\psi^j)) < \min_{\ell \neq j} K(P||Q^*(\psi^\ell))$, where $K(P||Q) := \int \log(dP/dQ)dP$.

Similarly as in Theorem 3, Theorem 4 establishes the equivalence result that, if we compare models that all use misspecified moments, then the marginal likelihood selects a model M_j if and only if M_j has the smallest Kullback–Leibler divergence $K(P||Q^*(\psi^j))$ between the true data distribution P and $Q^*(\psi^j)$. Remark that the condition $K(P||Q^*(\psi^j)) < K(P||Q^*(\psi^\ell))$, $\forall \ell \neq j$, given in the theorem does not depend on a particular value of ψ^j and ψ^ℓ . Indeed, the result of the theorem hinges on the fact that the marginal likelihood selects the model with the $Q^*(\psi^j)$ closer to P , that is, the model that contains the “less misspecified” moment restrictions for every value of ψ^j .

The result of the theorem also applies to the case where we compare a correctly specified model M_1 to misspecified models. Indeed, if model M_1 is correctly specified then $K(P||Q^*(\psi^1)) = 0$ while if model M_j is misspecified then $K(P||Q^*(\psi^j)) > 0$.

Example 4 (Model selection when all models are misspecified). For $i = 1, \dots, n$, let $y_i = \alpha + \beta z_i + e_i$. Here, we generate $z_i \sim \mathcal{N}(0.5, 1)$ and e_i from the skewed distribution in (2.13) with mean zero and variance 1.625, independently of z_i . Let $\theta := (\alpha, \beta)'$, $e_i(\theta) := (y_i - \alpha - \beta z_i)$, and the true value of θ be $(0, 1)'$. We compare the following models. Model 4: $E^P[(e_i(\theta), e_i(\theta)z_i, e_i(\theta)^3, e_i(\theta)^2 - 2)'] = 0$, Model 5: $E^P[(e_i(\theta), e_i(\theta)z_i, e_i(\theta)^2 - 2)'] = 0$, and Model 6: $E^P[e_i(\theta), e_i(\theta)^2 - 2] = 0$ which, written in terms of an encompassing grand model, become, respectively:

$$\begin{aligned} M_4: E^P[e_i(\theta)] &= 0, & E^P[e_i(\theta)z_i] &= 0, & E^P[(e_i(\theta))^3] &= 0, \\ & & E^P[(e_i(\theta))^2 - 2] &= 0 & & \\ M_5: E^P[e_i(\theta)] &= 0, & E^P[e_i(\theta)z_i] &= 0, & E^P[(e_i(\theta))^3] &= v_1, \\ & & E^P[(e_i(\theta))^2 - 2] &= 0 & & \\ M_6: E^P[e_i(\theta)] &= 0, & E^P[e_i(\theta)z_i] &= v_2, & E^P[(e_i(\theta))^3] &= v_1, \\ & & E^P[(e_i(\theta))^2 - 2] &= 0 & & \end{aligned} \tag{3.6}$$

with $\psi^4 = \theta$, $\psi^5 = (\theta, v_1)'$, and $\psi^6 = (\theta, v_1, v_2)'$. Thus, compared to Example 3 in the online Appendix, here we change the moment restriction that involves the variance of e_i . When the underlying distribution has variance different from 2, all models M_4 , M_5 , and M_6 are misspecified due to the new moment restriction: $E^P[(e_i(\theta))^2 - 2] = 0$. In Table 2, we report the percentage of times the marginal likelihood selects each model out of 500 trials, by sample size, under the default and training sample prior (based on 50 prior observations).

Because we know the true data-generating process, we can compute, for each model, the KL divergence between the true

Table 2. Model selection when all models are misspecified. Frequency (%) of times each of the three models in Example 5 are selected by the marginal likelihood criterion in 500 trials, by sample size, for two different prior distributions.

Model	Default prior			Training sample prior		
	M_4	M_5	M_6	M_4	M_5	M_6
$n = 250$	2.6	56.8	40.5	3.0	62.0	35.0
$n = 500$	0.2	28.1	71.6	1.0	31.0	68.0
$n = 1000$	0.0	4.2	95.8	0.0	4.0	96.0
$n = 2000$	0.0	0.0	100.0	0.0	0.0	100.0

model P and $Q^*(\psi^j)$ at the pseudo-true parameter ψ^j for model M_j based on (2.16). Using 10,000,000 simulated draws from P , our calculations show that $K(P||Q^*(\psi^j))$ is equal to 0.0283 for M_4 , 0.0096 for M_5 , and 1.4901×10^{-13} for M_6 . Intuitively, M_6 is the closest to the true model since it imposes fewer restrictions (only two moment restrictions are active). This means that the set of probability distributions that satisfy M_6 is larger than (and contains) the sets of probabilities that conform with M_4 and M_5 . This flexibility ensures that the divergence between the set of probabilities that satisfy M_6 and P (as measured by the KL) will be at least as small as for M_4 and M_5 . As the empirical results show, under each prior, the best model M_6 picked out by our marginal likelihood ranking is also the model that is the closest to the true model, consistent with the prediction of our theory.

Finally, suppose that some of the models that we consider are correctly specified and others are misspecified in the sense of Definition 1. This means that, for the latter, one or more of the active moment restrictions are invalid, or in other words, that one or more components of V are incorrectly set equal to zero. Indeed, all the models for which the active moment restrictions are valid are not misspecified even if some invalid moment restrictions are included among the inactive moment restrictions. This is because there always exists a value $v \in \mathbb{R}^{d_{v_i}}$ that equates the invalid moment restriction. In this case, the true v_* for this model will be different from the zero vector: $v_* \neq \mathbf{0}$ and the true value of the corresponding tilting parameter λ will be zero.

For this situation, Theorems 3 and 4 together imply an interesting corollary: the marginal likelihood selects the correctly specified model that contains the maximum number of overidentifying moment conditions. Without loss of generality, denote this model by M_1 . Then we have the following result.

Corollary 1. Let Assumptions 2–4, and Assumptions 5–8 in the online Appendix hold, and let either (2.14) or (2.17) be satisfied, depending on the model. Let us consider the comparison of J different models M_j , $j = 1, 2, \dots, J$ where M_1 satisfies Assumption 1 whereas M_j , $j \neq 1$ can either satisfy Assumption 1 or not. Then,

$$\lim_{n \rightarrow \infty} P \left(\log m(\mathbf{x}_{1:n}; M_1) > \max_{j \neq 1} \log m(\mathbf{x}_{1:n}; M_j) \right) = 1$$

if and only if $(p_1 + d_{v_1}) < (p_j + d_{v_j}), \forall j \neq 1$ such that M_j satisfies Assumption 1.

This corollary says that, if we compare a set of models, some of them are correctly specified and the others are misspecified, then the marginal likelihood selects model M_1 if and only if M_1 is correctly specified and contains the maximum number of overidentifying moment conditions among the correctly specified models.

4. Applications

The techniques discussed in the previous sections have wide-ranging applications to various statistical settings, such as generalized linear models, and to many different fields, such as biostatistics and economics. In fact, the methods discussed

above can be applied to virtually any problem that, in the frequentist setting, would be approached by generalized method of moments or estimating equation techniques. To illustrate some of the possibilities, we consider in this section two important problems: one in the context of count regression, and the second in the setting of instrumental variable (IV) regression.

4.1. Count Regression: Variable Selection

Suppose that $y_i, i = 1, \dots, n$ arise from the negative binomial (NB) regression model

$$y_i | \boldsymbol{\beta}, \mathbf{x}_i \sim NB\left(\frac{p}{1-p} \mu_i, p\right), \quad \mu_i > 0, \quad p \in (0, 1)$$

$$\log(\mu_i) = \mathbf{x}'_i \boldsymbol{\beta}, \tag{4.1}$$

where μ_i is the size parameter, $\mathbf{x}_i = (x_{1,i}, x_{2,i}, x_{3,i})'$, and $\boldsymbol{\beta} = (\beta_1 = 1, \beta_2 = 1, \beta_3 = 0)$. Thus, $x_{3,i}$ is a redundant regressor. Each explanatory variable $x_{j,i}$ is generated iid from an $\mathcal{N}(0.4, 1/9)$ distribution and p is set equal to 1/2. In this setting, suppose we wish to learn about $\boldsymbol{\beta}$ under the moment conditions

$$\mathbf{E}^P [(y_i - \exp(\beta_1 x_{1,i} + \beta_2 x_{2,i} + \beta_3 x_{3,i})) \mathbf{x}_i] = \mathbf{0}$$

$$\mathbf{E}^P \left[\left(\frac{y_i - \exp(\mathbf{x}'_i \boldsymbol{\beta})}{\sqrt{\exp(\mathbf{x}'_i \boldsymbol{\beta})}} \right)^2 - 1 \right] = v. \tag{4.2}$$

The first type of moment restriction (one for each $x_{j,i}$ for $j = 1, 2, 3$) is derived from the fact that the conditional expectation of y_i is $\exp(\mathbf{x}'_i \boldsymbol{\beta})$ and this identifies $\boldsymbol{\beta}$. The second type of restriction is suggested by a Poisson model (which is misspecified when the data arise from an NB). More specifically, if $v = 0$ that moment condition asserts that the conditional variance of y_i is equal to the conditional mean.

Suppose that we are interested in determining if x_3 is a redundant regressor and if the conditional mean and variance are equal. To solve this problem, we can create the following four models based on the grand model (4.2) with the following restrictions:

- M_1 : β_1 and β_2 are free parameters, $\beta_3 = 0$ and $v = 0$.
- M_2 : $\beta_1, \beta_2, \beta_3$ are free parameters and $v = 0$.
- M_3 : β_1, β_2 and v are free parameters, and $\beta_3 = 0$.
- M_4 : $\beta_1, \beta_2, \beta_3$ and v are free parameters. (4.3)

As required, each model has the same moment restrictions. The different models arise from the different restrictions on β_3 and v . In this set-up, models M_3 and M_4 are the correctly specified models but M_3 has more overidentifying moment restrictions than M_4 . We conduct our MCMC analysis and compute the marginal likelihoods of the four models by the Chib (1995) method under the default student- t prior distribution on $\boldsymbol{\beta}$ given in (2.11). The results are given in Table 3. The results show that the frequency of selecting M_3 is 94% for $n = 250$ and this percentage increases with sample size, in accordance with our theory. In addition, neither model M_1 nor M_2 (which state equality of the conditional mean and variance) is picked for any sample size.

Table 3. Frequency (%) of times each of the four models in (4.3) are selected by the marginal likelihood criterion in 500 trials. Model choice with Negative Binomial DGP. Model M_3 , defined by $\beta_3 = 0$ and v free, is the true model. The other models are defined in the text.

Model	M_1	M_2	M_3	M_4
$n = 250$	0.00	0.00	0.94	0.06
$n = 500$	0.00	0.00	0.95	0.05
$n = 1000$	0.00	0.00	0.96	0.04

In the online Appendix A, we report a similar analysis where the data are generated from a Poisson model. We emphasize that our analysis of these data, and the comparison across models, was light in terms of assumptions. The Poisson and negative binomial distributions are used to simply obtain a sample. These distributional forms are not featured in the estimation or the model comparison. A reader of this article wondered how a parametric Poisson model would have performed for these data. Since the data were generated under either a Poisson model or a model close to a Poisson model, the marginal likelihood of the Poisson model (correctly) is higher than that of the moment model. But this performance suffers dramatically if the data are generated from a count process that is quite different from the Poisson. For instance, suppose that the data are generated under the assumption that the first three moment conditions hold. We have developed a way of generating such a sample which works as follows. We first generate a large population of count data from an arbitrary count process (say $y_i = \lfloor \exp\{\beta_1 x_{1,i} + \beta_2 x_{2,i} + 20\mathcal{N}(0, 1)\} \rfloor$), setting any negative observations to zero and where $\lfloor a \rfloor$ denotes the largest integer less than or equal to a). We then find the ETEL probabilities p_i^* consistent with the given moment conditions. Finally, we sample the population of observations according to these probabilities. The resulting sample satisfies the moment conditions but has no connection to the Poisson or negative binomial distributional forms. For such a design, in 500 replications, in the parametric Poisson models (one with $\beta_3 = 0$ and one with β_3 free), the Poisson model with $\beta_3 = 0$ is selected 42% when $n = 250$, 46% when $n = 500$, and 45% when $n = 1000$. Thus, the Poisson assumption is not capable of selecting the correct case. In addition, when these two Poisson models are compared along with the four moment models in (4.3), for which the marginal likelihoods are computed with our method, M_3 is decisively preferred over the Poisson models in terms of marginal likelihood and the frequency of times it is selected is similar to that reported above in Table 3.

4.2. IV Regression

Consider now the commonly occurring situation with observational data where one is interested in learning about the causal effect parameter β in the model

$$y = \alpha + x\beta + w\delta + \varepsilon, \quad \mathbf{E}^P[\varepsilon] = 0$$

but the covariate x is correlated with the error, due to say unmeasured or uncontrolled factors, apart from w , that are correlated with x and that reside in ε . Also suppose that one has two valid instrumental variables z_1 and z_2 that (by definition) are correlated with x but uncorrelated with ε . In this setting, we

Table 4. Posterior summary for two simulated sample sizes from IV regression model with skewed error. The true value of α is 1, of β is 0.5, and of δ is 0.7. The summaries are based on 10,000 MCMC draws beyond a burn-in of 1000. The M-H acceptance rate is around 90% in both cases.

	mean	sd	median	lower	upper	ineff
$n = 250$						
α	1.26	0.11	1.26	1.04	1.48	1.26
β	0.55	0.03	0.55	0.48	0.61	1.41
δ	0.28	0.17	0.28	-0.06	0.60	1.27
$n = 2000$						
α	1.01	0.05	1.01	0.92	1.10	1.30
β	0.51	0.02	0.51	0.48	0.54	1.33
δ	0.66	0.07	0.66	0.52	0.81	1.38

can learn about $\theta := (\alpha, \beta, \delta)$ from the overidentified moment restrictions

$$E^P [(y_i - \alpha - x_i\beta - w_i\delta)] = 0 \tag{4.4}$$

$$E^P [(y_i - \alpha - x_i\beta - w_i\delta) z_{1i}] = 0 \tag{4.5}$$

$$E^P [(y_i - \alpha - x_i\beta - w_i\delta) z_{2i}] = 0 \tag{4.6}$$

$$E^P [(y_i - \alpha - x_i\beta - w_i\delta) w_i] = 0, \quad (i \leq n) \tag{4.7}$$

without having to model the distribution of ε or the model connecting z to x .

To demonstrate the performance of our Bayesian prior-posterior analysis in this setting, we generate data on $(y_i, x_i, z_{1i}, z_{2i})$ from a design that incorporates a skewed marginal distribution of ε and substantial correlation between x and ε . In our DGP, we assume that $y = 1 + 0.5x + 0.7w + \varepsilon$, $x = z_1 + z_2 + w + u$, and generate z_j from $\mathcal{N}(0.5, 1)$ and w from $\text{Uniform}(0, 1)$. The errors (ε, u) are generated from a Gaussian copula whose covariance matrix has 1 on the diagonal, and 0.7 on the off-diagonal, such that the ε marginal distribution is the skewed bivariate mixture $0.5\mathcal{N}(0.5, 0.5^2) + 0.5\mathcal{N}(-0.5, 1.118^2)$ and the u marginal distribution is $\mathcal{N}(0, 1)$. We generate $n = 250$ and $n = 2000$ observations from this design and use moment conditions (4.4)–(4.7) and our default student- t prior given in (2.11) to learn about θ . The results shown in Table 4 and Figure 1 demonstrate clearly the ability of our method to concentrate on the true values of the parameters, under minimal assumptions.

Now suppose that we are unsure that z_2 is an appropriate instrument. We can address this concern by estimating a new model M_2 in which the moment condition (4.6) is not active. The marginal likelihood of this model can be compared with

the marginal likelihood of the previous model M_1 . The results show that for $n = 250$, the log-marginal likelihood of M_1 is -1395.807 and that of M_2 is -1398.092 , while for $n = 2000$, the corresponding log-marginal likelihoods are -15217.78 and -15222.65 , respectively, thus correctly indicating for both sample sizes that z_2 is an appropriate instrument.

5. Conclusion

In this article, we have developed a fully Bayesian framework for estimation and model comparisons in statistical models that are defined by moment restrictions. The Bayesian analysis of such models has always been viewed as a challenge because traditional Bayesian semiparametric methods, such as those based on Dirichlet process mixtures and variants thereof, are not suitable for such models. What we have shown in this article is that the exponentially tilted empirical likelihood setting is an immensely useful organizing framework within which a fully Bayesian treatment of such models can be developed. We have established a number of new, powerful results surrounding the Bayesian ETEL framework including the treatment of models that are possibly misspecified. We show how the moment conditions can be reexpressed in terms of additional nuisance parameters and that the Bayesian ETEL posterior distribution satisfies a Bernstein–von Mises theorem. We have also developed a framework for comparing moment condition models based on marginal likelihoods and Bayes factors and provided a suitable large sample theory for model selection consistency. Our results show that the marginal likelihood favors the model with the minimum number of parameters and the maximum number of valid moment restrictions. When the models are misspecified, the marginal likelihood-based selection procedure selects the model that is closer to the (unknown) true data-generating process in terms of the Kullback–Leibler divergence. The ideas and results illumined in this article now provide the means for analyzing a whole array of models from the Bayesian viewpoint. This broadening of the scope of Bayesian techniques to previously intractable problems is likely to have far-reaching practical consequences.

Appendix: Proofs for Section 3

In this Appendix, we provide the proof of Theorems 3 and 4. The proofs of all the other results are in the online Appendix.




Figure 1. Posterior densities of β in the IV regression with skewed error. Posterior densities are based on 10,000 draws beyond a burn-in of 1000. The M-H acceptance rate is about 90% for each sample size.

Notation: Let $\widehat{\lambda}(\psi) := \arg \min_{\lambda \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n [\exp\{\lambda' g^A(x_i, \psi)\}]$. We recall the Schennach (2007) ETEL estimator of ψ , denoted by $\widehat{\psi} := (\widehat{\theta}, \widehat{v})$:

$$\widehat{\psi} := \arg \max_{\psi \in \Psi} \frac{1}{n} \sum_{i=1}^n \left[\widehat{\lambda}(\psi)' g^A(x_i, \psi) - \log \frac{1}{n} \sum_{j=1}^n \exp\{\widehat{\lambda}(\psi)' g^A(x_j, \psi)\} \right]. \tag{A.1}$$

The following notation is used hereafter. The ETEL estimator of ψ^ℓ in model M_ℓ is

$$\widehat{\psi}^\ell := \arg \max_{\psi^\ell \in \Psi^\ell} \frac{1}{n} \sum_{i=1}^n \left[\widehat{\lambda}(\psi^\ell)' g^A(x_i, \psi^\ell) - \log \frac{1}{n} \sum_{j=1}^n \exp\{\widehat{\lambda}(\psi^\ell)' g^A(x_j, \psi^\ell)\} \right], \tag{A.2}$$

where $\widehat{\lambda}(\psi^\ell) = \arg \min_{\lambda \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n [\exp\{\lambda' g^A(x_i, \psi^\ell)\}]$. Denote $\widehat{g}^A(\psi^\ell) := \frac{1}{n} \sum_{i=1}^n g^A(x_i, \psi^\ell)$, $\widehat{g}_\ell^A := \widehat{g}^A(\psi^\ell)$, $\widehat{L}(\psi^\ell) := \exp\{\widehat{\lambda}(\psi^\ell)' \widehat{g}^A(\psi^\ell)\} [n^{-1} \sum_{i=1}^n \exp\{\widehat{\lambda}(\psi^\ell)' g^A(x_i, \psi^\ell)\}]^{-1}$, and $L(\psi^\ell) := \exp\{\lambda_{\circ}(\psi^\ell)' \mathbb{E}^P[g^A(x, \psi^\ell)]\} (\mathbb{E}^P[\exp\{\lambda_{\circ}(\psi^\ell)' g^A(x, \psi^\ell)\}])^{-1}$. Moreover, we use the notation $\Sigma_\ell := (\Gamma_\ell' \Delta_\ell^{-1} \Gamma_\ell)^{-1}$ where $\Gamma_\ell := \mathbb{E}^P[\frac{\partial}{\partial \psi^\ell} g^A(X, \psi_\ast^\ell)]$ and $\Delta_\ell := \mathbb{E}^P[g^A(X, \psi_\ast^\ell) g^A(X, \psi_\ast^\ell)']$. In the proofs, we omit measurability issues which can be dealt with in the usual manner by replacing probabilities with outer probabilities.

Proof of Theorem 3. By (3.5) and Lemmas D.1 and D.2 in the online Appendix, we obtain

$$\begin{aligned} & P\left(\max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell) < \log m(\mathbf{x}_{1:n}; M_j)\right) \\ &= P\left(\max_{\ell \neq j} \left[-\frac{n \widehat{g}_\ell^A{}' \Delta^{-1} \widehat{g}_\ell^A}{2} + \log \pi(\widehat{\psi}^\ell | M_\ell) - \frac{(p_\ell + d_{v_\ell})}{2} (\log n - \log(2\pi)) + \frac{1}{2} \log |\Sigma_\ell| \right] \right. \\ &\quad \left. + \frac{n \widehat{g}_j^A{}' \Delta^{-1} \widehat{g}_j^A}{2} + o_p(1) < \log \pi(\widehat{\psi}^j | M_j) - \frac{(p_j + d_{v_j})}{2} (\log n - \log(2\pi)) + \frac{1}{2} \log |\Sigma_j| \right). \end{aligned} \tag{A.3}$$

Remark that $n \widehat{g}_j^A{}' \Delta^{-1} \widehat{g}_j^A \xrightarrow{d} \chi_{d-(p_j+d_{v_j})}^2, \forall j$, so that $n \widehat{g}_j^A{}' \Delta^{-1} \widehat{g}_j^A = O_p(1)$. Suppose first that $(p_\ell + d_{v_\ell}) > (p_j + d_{v_j}), \forall \ell \neq j$. Since $-n \widehat{g}_\ell^A{}' \Delta^{-1} \widehat{g}_\ell^A < 0$ for every ℓ , we lower bound (A.3) as

$$\begin{aligned} & P\left(\max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell) < \log m(\mathbf{x}_{1:n}; M_j)\right) \\ &\geq P\left(\frac{n \widehat{g}_j^A{}' \Delta^{-1} \widehat{g}_j^A}{2} + o_p(1) < \log n \left[\frac{\min_{\ell \neq j} (p_\ell + d_{v_\ell}) - p_j - d_{v_j}}{2} \right. \right. \\ &\quad \left. \left. - \frac{\min_{\ell \neq j} (p_\ell + d_{v_\ell}) - p_j - d_{v_j}}{2 \log n} \log(2\pi) - \frac{\log[\max_{\ell \neq j} \pi(\widehat{\psi}^\ell | M_\ell) / \pi(\widehat{\psi}^j | M_j)]}{\log n} \right] \right) \\ &= P\left(\underbrace{\frac{n \widehat{g}_j^A{}' \Delta^{-1} \widehat{g}_j^A}{2} + o_p(1)}_{=: \mathcal{I}_n} < \log n \left[\frac{\min_{\ell \neq j} (p_\ell + d_{v_\ell}) - p_j - d_{v_j}}{2} \right. \right. \\ &\quad \left. \left. - \frac{\min_{\ell \neq j} (p_\ell + d_{v_\ell}) - p_j - d_{v_j}}{2 \log n} \log(2\pi) - \frac{\log[\max_{\ell \neq j} \pi(\widehat{\psi}^\ell | M_\ell) / \pi(\widehat{\psi}^j | M_j)]}{\log n} \right] \right) \end{aligned}$$

$$< \log n \left[\underbrace{\frac{\min_{\ell \neq j} (p_\ell + d_{v_\ell}) - p_j - d_{v_j}}{2}}_{=: \mathcal{I}_n} + \mathcal{O}_p((\log n)^{-1}) \right]. \tag{A.4}$$

Because $\mathcal{I}_n = \mathcal{O}_p(1)$ (and is asymptotically positive) and \mathcal{I}_n is strictly positive as $n \rightarrow \infty$ (since $(p_\ell + d_{v_\ell}) > (p_j + d_{v_j}), \forall \ell \neq j$) and converges to $+\infty$, then the probability converges to 1. This proves one direction of the statement.

To prove the second direction of the statement, suppose that $\lim_{n \rightarrow \infty} P(\max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell) < \log m(\mathbf{x}_{1:n}; M_j)) = 1$ and consider the following upper bound (which follows from (A.3) and the fact that $n \widehat{g}_j^A{}' \Delta^{-1} \widehat{g}_j^A > 0, \forall n$):

$$\begin{aligned} & P\left(\max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell) < \log m(\mathbf{x}_{1:n}; M_j)\right) \\ &\leq P\left(\log m(\mathbf{x}_{1:n}; M_\ell) < \log m(\mathbf{x}_{1:n}; M_j)\right), \quad \forall \ell \neq j \\ &\leq P\left(-\frac{n \widehat{g}_\ell^A{}' \Delta^{-1} \widehat{g}_\ell^A}{2} + o_p(1) + \log n \left[\frac{(p_j + d_{v_j}) - (p_\ell + d_{v_\ell})}{2} \right. \right. \\ &\quad \left. \left. + \mathcal{O}_p\left(\frac{1}{\log n}\right) \right] < 0\right), \quad \forall \ell \neq j. \end{aligned} \tag{A.5}$$

Because the probability in the first line of (A.5) converges to 1 as $n \rightarrow \infty$ then, necessarily, the probability in the last line of (A.5) converges to 1 which is possible only if $(p_j + d_{v_j}) < (p_\ell + d_{v_\ell})$ because $\log n \left[\frac{(p_j + d_{v_j}) - (p_\ell + d_{v_\ell})}{2} \right]$ is the dominating term since $-\frac{n \widehat{g}_\ell^A{}' \Delta^{-1} \widehat{g}_\ell^A}{2} < 0$ and it remains bounded as $n \rightarrow \infty$. Since the first inequality in (A.5) holds $\forall \ell \neq j$ then convergence to 1 of the probability in the last line of (A.5) is possible only if $(p_j + d_{v_j}) < (p_\ell + d_{v_\ell}), \forall \ell \neq j$. \square

Proof of Theorem 4. We can write $\log p(\mathbf{x}_{1:n} | \psi^\ell; M_\ell) = -n \log n + n \log \widehat{L}(\psi^\ell)$. Then, we have

$$\begin{aligned} & P\left(\log m(\mathbf{x}_{1:n}; M_j) > \max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell)\right) \\ &= P\left(n \log \widehat{L}(\psi_\circ^j) + \log \pi(\psi_\circ^j | M_j) - \log \pi(\psi_\circ^j | \mathbf{x}_{1:n}, M_j) \right. \\ &\quad \left. > \max_{\ell \neq j} [n \log \widehat{L}(\psi_\circ^\ell) + \log \pi(\psi_\circ^\ell | M_\ell) - \log \pi(\psi_\circ^\ell | \mathbf{x}_{1:n}, M_\ell)]\right) \\ &= P\left(n \log L(\psi_\circ^j) + n \log \frac{\widehat{L}(\psi_\circ^j)}{L(\psi_\circ^j)} + \mathcal{B}_j \right. \\ &\quad \left. > \max_{\ell \neq j} \left[n \log L(\psi_\circ^\ell) + \mathcal{B}_\ell + n \log \frac{\widehat{L}(\psi_\circ^\ell)}{L(\psi_\circ^\ell)} \right] \right), \end{aligned} \tag{A.6}$$

where $\forall \ell, \mathcal{B}_\ell := \log \pi(\psi_\circ^\ell | M_\ell) - \log \pi(\psi_\circ^\ell | \mathbf{x}_{1:n}, M_\ell)$ and $\mathcal{B}_\ell = \mathcal{O}_p(1)$ under the assumptions of Theorem 2. By definition of $dQ^*(\psi)$ in Section 2.3 we have that: $\log L(\psi_\circ^\ell) = \mathbb{E}^P[\log dQ^*(\psi_\circ^\ell) / dP] = -\mathbb{E}^P[\log dP / dQ^*(\psi_\circ^\ell)] = -K(P || Q^*(\psi_\circ^\ell))$. Remark that $\mathbb{E}^P[\log(dP/dQ^*(\psi_\circ^\ell))] > \mathbb{E}^P[\log(dP/dQ^*(\psi_\circ^j))]$ means that the KL divergence between P and $Q^*(\psi_\circ^\ell)$, is smaller for model M_1 than for model M_2 , where $Q^*(\psi_\circ^\ell)$ minimizes the KL divergence between $Q \in \mathcal{P}_{\psi_\circ^\ell}$ and P for $\ell \in \{1, 2\}$ (notice the inversion of the two probabilities).

First, suppose that $\min_{\ell \neq j} \mathbb{E}^P[\log(dP/dQ^*(\psi_\circ^\ell))] > \mathbb{E}^P[\log(dP/dQ^*(\psi_\circ^j))]$. By (A.6):

$$\begin{aligned} & P\left(\log m(\mathbf{x}_{1:n}; M_j) > \max_{\ell \neq j} \log m(\mathbf{x}_{1:n}; M_\ell)\right) \\ &\geq P\left(\log \frac{\widehat{L}(\psi_\circ^j)}{L(\psi_\circ^j)} - \max_{\ell \neq j} \log \frac{\widehat{L}(\psi_\circ^\ell)}{L(\psi_\circ^\ell)} + \frac{1}{n} (\mathcal{B}_j - \max_{\ell \neq j} \mathcal{B}_\ell) \right. \\ &\quad \left. > \underbrace{\max_{\ell \neq j} \log L(\psi_\circ^\ell) - \log L(\psi_\circ^j)}_{=: \mathcal{I}_n} \right). \end{aligned} \tag{A.7}$$

This probability converges to 1 because $\mathcal{I}_n = K(P||Q^*(\psi^j)) - \min_{\ell \neq j} K(P||Q^*(\psi^\ell)) < 0$ by assumption, and $[\log \widehat{L}(\psi^\ell) - \log L(\psi^\ell)] \xrightarrow{P} 0$, for every $\psi^\ell \in \Psi^\ell$ and every $\ell \in \{1, 2\}$ by Lemma D.3 in the online Appendix.

To prove the second direction of the statement, suppose that $\lim_{n \rightarrow \infty} P(\log m(x_{1:n}; M_j) > \max_{\ell \neq j} \log m(x_{1:n}; M_\ell)) = 1$. By (A.6) it holds, $\forall \ell \neq j$

$$\begin{aligned} & P\left(\log m(x_{1:n}; M_j) > \max_{\ell \neq j} \log m(x_{1:n}; M_\ell)\right) \\ & \leq P\left(\log \frac{\widehat{L}(\psi^j)}{L(\psi^j)} - \log \frac{\widehat{L}(\psi^\ell)}{L(\psi^\ell)} + \frac{1}{n}(\mathcal{B}_j - \mathcal{B}_\ell) > \log \frac{L(\psi^\ell)}{L(\psi^j)}\right). \end{aligned} \quad (\text{A.8})$$

Convergence to 1 of the left-hand side implies convergence to 1 of the right-hand side which is possible only if $\log L(\psi^\ell) - \log L(\psi^j) < 0$. Since this is true for every model ℓ , then this implies that $K(P||Q^*(\psi^j)) < \min_{\ell \neq j} K(P||Q^*(\psi^\ell))$ which concludes the proof. \square

Supplementary Material

The supplementary material in the online Appendix contains further examples, assumptions, and the technical proofs of the results in the article.

Acknowledgments

The authors gratefully thank the co-editor, an associate editor, and four anonymous referees for their many constructive comments on the previous version of the article.

Funding

Anna Simoni gratefully acknowledges financial support from ANR-13-BSH1-0004 and ANR-11-LABEX-0047.

References

- Bornn, L., Shephard, N., and Solgi, R. (2015), "Moment Conditions and Bayesian Nonparametrics," Technical Report. arXiv:1507.08645. [1657]
- Broniatowski, M., and Keziou, A. (2012), "Divergences and Duality for Estimation and Test Under Moment Condition Models," *Journal of Statistical Planning and Inference*, 142, 2554–2573. [1656]
- Chamberlain, G. (1987), "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions," *Journal of Econometrics*, 34, 305–334. [1660]
- Chang, I. H., and Mukerjee, R. (2008), "Bayesian and Frequentist Confidence Intervals Arising from Empirical-type Likelihoods," *Biometrika*, 95, 139–147. [1656]
- Chaudhuri, S., and Ghosh, M. (2011), "Empirical Likelihood for Small Area Estimation," *Biometrika*, 98, 473–480. [1656]
- Chaudhuri, S., Mondal, D., and Yin, T. (2017), "Hamiltonian Monte Carlo Sampling in Bayesian Empirical Likelihood Computation," *Journal of the Royal Statistical Society, Series B*, 79, 293–320. [1656]
- Chen, S. X., and Van Keilegom, I. (2009), "A Review on Empirical Likelihood Methods for Regression," *TEST*, 18, 415–447. [1656]
- Chib, S. (1995), "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association*, 90, 1313–1321. [1657, 1663, 1665]
- Chib, S., and Greenberg, E. (1995), "Understanding the Metropolis-Hastings Algorithm," *The American Statistician*, 49, 327–335. [1659]
- Chib, S., and Jeliazkov, I. (2001), "Marginal Likelihood From the Metropolis-Hastings Output," *Journal of the American Statistical Association*, 96, 270–281. [1657, 1663]
- Chib, S., and Ramamurthy, S. (2010), "Tailored Randomized Block MCMC methods with Application to DSGE Models," *Journal of Econometrics*, 155, 19–38. [1659]
- Csiszar, I. (1975), "I-Divergence Geometry of Probability Distributions and Minimization Problems," *Annals of Probability*, 3, 146–158. [1663]
- Fang, K.-T., and Mukerjee, R. (2006), "Empirical-Type Likelihoods Allowing Posterior Credible Sets with Frequentist Validity: Higher-Order Asymptotics," *Biometrika*, 93, 723–733. [1656]
- Florens, J.-P., and Simoni, A. (2016), "Gaussian Processes and Bayesian Moment Estimation," Technical Report. arXiv:1607.07343. [1657]
- Ghosh, J., and Ramamoorthi, R. (2003), *Bayesian Nonparametrics*, New York: Springer-Verlag. [1660]
- Grendar, M., and Judge, G. (2009), "Asymptotic Equivalence of Empirical Likelihood and Bayesian MAP," *The Annals of Statistics*, 37, 2445–2457. [1656]
- Hansen, P. (1982), "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054. [1660]
- Hong, H., and Preston, B. (2012), "Bayesian Averaging, Prediction and Nonnested Model Selection," *Journal of Econometrics*, 167, 358–369. [1657]
- Imbens, G. W. (1997), "One-Step Estimators for Over-Identified Generalized Method of Moments Models," *The Review of Economic Studies*, 64, 359–383. [1656]
- Kim, M.-O., and Yang, Y. (2011), "Semiparametric Approach to a Random Effects Quantile Regression Model," *Journal of the American Statistical Association*, 106, 1405–1417. [1656]
- Kitamura, Y., and Otsu, T. (2011), "Bayesian Analysis of Moment Condition Models Using Nonparametric Priors," Technical Report. [1657]
- Kitamura, Y., and Stutzer, M. (1997), "An Information-Theoretic Alternative to Generalized Method of Moments Estimation," *Econometrica*, 65, 861–874. [1656]
- Kleijn, B., and van der Vaart, A. (2012), "The Bernstein-Von-Mises Theorem Under Misspecification," *Electronic Journal Statistics*, 6, 354–381. [1661]
- Lancaster, T., and Jun, S. J. (2010), "Bayesian Quantile Regression Methods," *Journal of Applied Econometrics*, 25, 287–307. [1656]
- Lazar, N. A. (2003), "Bayesian Empirical Likelihood," *Biometrika*, 90, 319–326. [1656]
- Lehmann, E. L., and Casella, G. (1998), *Theory of Point Estimation (Springer Texts in Statistics)* (2nd ed.), New York: Springer. [1660]
- Owen, A. B. (1988), "Empirical Likelihood Ratio Confidence Intervals for a Single Functional," *Biometrika*, 75, 237–249. [1656]
- (1990), "Empirical Likelihood Ratio Confidence Regions," *Annals of Statistics*, 18, 90–120. [1656]
- (2001), *Empirical Likelihood (Chapman & Hall/CRC Monographs on Statistics & Applied Probability)*, Boca Raton, FL: Chapman & Hall/CRC. [1656]
- Porter, A. T., Holan, S. H., and Wikle, C. K. (2015), "Bayesian Semiparametric Hierarchical Empirical Likelihood Spatial Models," *Journal of Statistical Planning and Inference*, 165, 78–90. [1656]
- Qin, J., and Lawless, J. (1994), "Empirical Likelihood and General Estimating Equations," *Annals of Statistics*, 22, 300–325. [1656]
- Rao, J., and Wu, C. (2010), "Bayesian Pseudo-Empirical-Likelihood Intervals for Complex Surveys," *Journal of the Royal Statistical Society, Series B*, 72, 533–544. [1656]
- Robert, C. (2007), *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*, New York: Springer. [1657]
- Schennach, S. M. (2005), "Bayesian Exponentially Tilted Empirical Likelihood," *Biometrika*, 92, 31–46. [1656, 1658, 1660]
- (2007), "Point Estimation with Exponentially Tilted Empirical Likelihood," *Annals of Statistics*, 35, 634–672. [1656, 1657, 1661, 1667]
- Sueishi, N. (2013), "Identification Problem of the Exponential Tilting Estimator under Misspecification," *Economics Letters*, 118, 509–511. [1660]
- Variyath, A. M., Chen, J., and Abraham, B. (2010), "Empirical Likelihood Based Variable Selection," *Journal of Statistical Planning and Inference*, 140, 971–981. [1657]
- Vexler, A., Deng, W., and Wilding, G. E. (2013), "Nonparametric Bayes Factors Based on Empirical Likelihood Ratios," *Journal of Statistical Planning and Inference*, 143, 611–620. [1657]
- Xi, R., Li, Y., and Hu, Y. (2016), "Bayesian Quantile Regression Based on the Empirical Likelihood with Spike and Slab Priors," *Bayesian Analysis*, 11, 821–855. [1656]
- Yang, Y., and He, X. (2012), "Bayesian Empirical Likelihood for Quantile Regression," *The Annals of Statistics*, 40, 1102–1131. [1656]