

QUASI-BAYESIAN ESTIMATION AND INFERENCE WITH CONTROL FUNCTIONS

Ruixuan Liu and Zhengfei Yu

Chinese University of Hong Kong and University of Tsukuba

ABSTRACT. We consider a quasi-Bayesian method that combines a frequentist estimation in the first stage and a Bayesian estimation/inference approach in the second stage. The study is motivated by structural discrete choice models that use the control function methodology to correct for endogeneity bias. In this scenario, the first stage estimates the control function using some frequentist parametric or nonparametric approach. The structural equation in the second stage, associated with certain complicated likelihood functions, can be more conveniently dealt with using a Bayesian approach. This paper studies the asymptotic properties of the quasi-posterior distributions obtained from the second stage. We prove that the corresponding quasi-Bayesian credible set does not have the desired coverage in large samples. Nonetheless, the quasi-Bayesian point estimator remains consistent and is asymptotically equivalent to a frequentist two-stage estimator. We show that one can obtain valid inference by bootstrapping the quasi-posterior that takes into account the first-stage estimation uncertainty.

Key words: Bayesian methods, Bernstein–von Mises theorem, endogenous regressors, control function, semiparametric estimation

JEL classification: C11, C14, C15, C25

Ruixuan Liu, CUHK Business School, Chinese University of Hong Kong, 12 Chak Cheung Street, Cheng Yu Tung Building, Hong Kong, E-mail: ruixuanliu@cuhk.edu.hk.

Zhengfei Yu, Faculty of Humanities and Social Sciences, University of Tsukuba, 1-1-1, Tennodai, Tsukuba, Ibaraki 305-8571, Japan. E-mail: yu.zhengfei.gn@u.tsukuba.ac.jp.

We would like to thank Christoph Breunig, Yanqin Fan, Emmanuel Guerre, Hide Ichimura, Shakeeb Khan, Chang-Jin Kim, Yuan Liao, Zhipeng Liao, Essie Maasoumi, Jing Tao, Jun Yu, and Yichong Zhang for helpful comments and discussions. Liu's research is supported by the GRF Number 14503523 from the Research Grants Council Hong Kong. Yu gratefully acknowledges the support of JSPS KAKENHI Grant Number 21K01419. The usual disclaimer applies.

1. Introduction

The control function approach is commonly used in empirical economic studies to cope with endogenous explanatory variables in the structural equations (Heckman and Robb, 1985; Wooldridge, 2015). This approach typically involves two stages: the first stage projects the endogenous explanatory variables to a set of instruments and other exogenous variables and obtains the residuals; these residuals serve as control functions in the second stage as extra explanatory variables when estimating the structural equations. In this paper, we propose a quasi-Bayesian method that combines a frequentist-type first-stage estimation with some Bayesian approach in the second stage. The study is motivated by structural discrete choice models that apply the control function method to correct for endogeneity bias. In such scenarios, the second-stage estimation remains challenging for frequentist methods. The practical computational convenience of our proposal is easy to digest because the Bayesian paradigm turns the challenging integration or optimization problem in the second stage into a sampling problem of posterior distributions. Therefore, this quasi-Bayesian approach offers researchers great flexibility to combine various non-parametric frequentist methods with state-of-the-art Bayesian algorithms. With a clear understanding of its theoretical foundation, this new procedure can be particularly useful in econometric models with sophisticated likelihood and endogeneity issues.

To formalize the idea, we consider independent and identically distributed (*i.i.d.*) observations $\mathbf{Y}^n = (Y_i, i = 1, \dots, n)^\top$. The component ζ_0 represents a vector of unknown functions that one needs to estimate from the first-stage equation in order to obtain the residuals. Throughout the paper, we assume that the true ζ_0 is identifiable and can be estimated from the subvector $\mathbf{Y}_1^n = (Y_{1,i}, i = 1, \dots, n)^\top$ that consists of the observables from the first stage. We denote the resulting estimator by $\widehat{\zeta}_n$. The second stage postulates a probability density function, $p_\theta^n(\cdot; \zeta)$ for the subvector $\mathbf{Y}_2^n = (Y_{2,i}, i = 1, \dots, n)^\top$ determined by the structural equation with the unknown parameter of interest $\theta \in \Theta \subset \mathbb{R}^p$, for some integer $p \geq 1$. The unique truth is denoted by θ_0 . The Bayesian approach in the second stage posits a prior distribution $\Pi(\theta)$ and updates this prior to the quasi-posterior distribution given the limited-information likelihood¹ in the second stage. Using the Bayes' rule, the quasi-posterior of the finite-dimensional parameter θ is

$$(1.1) \quad \Pi[\theta \in \mathcal{A} | \mathbf{Y}_2^n; \widehat{\zeta}_n] = \frac{\int_{\mathcal{A}} p_\theta^n(\mathbf{Y}_2^n; \widehat{\zeta}_n) d\Pi(\theta)}{\int_{\Theta} p_{\theta'}^n(\mathbf{Y}_2^n; \widehat{\zeta}_n) d\Pi(\theta')}.$$

¹The current approach is based on the “limited information” strategy in the sense that the information contained in both stages (or equations) are not simultaneously considered.

We refer to the estimation/inferential procedure derived from the above posterior as quasi-Bayesian because the term $p_{\theta}^n(\mathbf{Y}_2^n; \widehat{\zeta}_n)$ is not a genuine likelihood, and it depends on some plug-in estimator $\widehat{\zeta}_n$. Two key questions arise: (1). Is the quasi-posterior distribution of θ asymptotically Gaussian? (2). What are the center and dispersion of this quasi-posterior distribution? To our best knowledge, such problems have not been investigated in the literature even when the estimator $\widehat{\zeta}_n$ is entirely parametric. Our paper fills this void, and we allow for any flexible first-stage nonparametric estimation.

Our main objective is to provide a thorough theoretical understanding of the large sample properties of this quasi-Bayesian procedure from the frequentist's perspective. The challenge is that one faces a nonstandard posterior distribution that depends on the first-stage estimation. It is unclear whether the corresponding quasi-Bayesian credible set agrees with the frequentist's confidence set, even asymptotically. We show that this credible set does not have the right coverage in large samples because it ignores the estimation uncertainty from the first stage. On the other hand, the quasi-Bayesian point estimator, such as the quasi-posterior mean, is asymptotically equivalent to a frequentist two-stage estimator; that is, it is root- n asymptotically normal with the same sandwich-type covariance matrix as its frequentist counterpart. In essence, the quasi-posterior distribution is correctly centered but with a wrong dispersion. We also show that a proper bootstrap method will restore the desirable coverage probability and thus is valid for inference. Although the bootstrap requires additional simulation, its implementation can be easily parallelized.

To illustrate our theory, we focus on an endogenous multinomial Probit (MNP) model from Petrin and Train (2010), which we refer to as the Petrin–Train model in the sequel. The framework is popular in modeling consumer choices over different products, say choosing from $j \in \{0, \dots, J\}$, with the 0-th category being the non-purchase option. The endogenous variable $(X_j^e)_{j=0}^J$ can be the price, advertising, or quality in the context. To address the endogeneity problem, researchers often construct the Hausman-type price instruments (Petrin and Train, 2010). The first stage involves the estimation of a vector of $(\zeta_j)_{j=0}^J$, as the conditional mean functions of the endogenous variables given the exogenous covariates and instrument variables. Given the estimated conditional mean functions $\widehat{\zeta}_n = (\widehat{\zeta}_{n,j})_{j=0}^J$, we can extract the residuals as the control functions or variables for the second-stage MNP model. In this scenario, the likelihood function $p_{\theta}^n(\mathbf{Y}_2^n; \widehat{\zeta}_n)$ is the product of the conditional choice probabilities that involve complicated multiple integrals. As reviewed in Train (2009), the Bayesian approach enjoys practical advantages over frequentist methods in estimating the MNP-type models.

Going beyond the Petrin-Train model, our study reflects two generic features of modern econometric models. On the one hand, many models used to describe observed data may be so complex that the likelihoods associated with these models are computationally

intractable. These analytical difficulties can be alleviated by simulation-based procedures, such as the Bayesian method, which has proven successful in many areas. On the other hand, the control function approach provides an effective solution to deal with the endogeneity problem. This approach treats endogeneity as an omitted variable bias problem, in which the inclusion of estimates of the first-stage errors as additional covariates corrects the inconsistency of the second stage. The idea of combining a frequentist first stage and a Bayesian second stage has been practiced in empirical research (Agarwal and Somaini, 2018), though lacking a formal theoretical investigation for such a “hybrid” or quasi-Bayesian procedure. In their empirical analysis of the distribution of students’ preferences for public schools, Agarwal and Somaini (2018) first estimate the believed assignment probabilities at various schools using frequentist methods and then conduct Bayesian estimation in the second stage due to a complicated likelihood similar to the MNP. They also suggest a bootstrap method for inference. Our paper sheds a theoretical light on the asymptotic normality of the two-stage quasi-Bayesian estimator and the bootstrap validity therein.

There is a clear advantage of our proposal over the frequentist two-step method or the Bayesian full information paradigm in the current context, because it separates two tasks: one can correct for the endogeneity bias via robust nonparametric first-stage estimation, while tackling the complicated likelihood in the second stage by a Bayesian approach. This separation of tasks allows for considerable algorithmic flexibility. Regarding the frequentist two-step approach, the second stage requires the simulated method of moments (McFadden, 1989; Hajivassiliou and McFadden, 1998), which calls for Monte Carlo simulation, in addition to solving hard optimization problems for finding the frequentist type point estimator. Referring to the full information Bayesian approach, one would focus on the posterior that is proportional to $\tilde{p}_{\theta,\zeta}^n(\mathbf{Y}^n)d\Pi(\theta)dQ(\zeta)$, where $\tilde{p}_{\theta,\zeta}^n(\mathbf{Y}^n)$ is the density function for the joint distribution of \mathbf{Y}_1^n and \mathbf{Y}_2^n combining two stages together, and $Q(\cdot)$ is a suitable prior on the functional component ζ . The joint likelihood inevitably requires a full model encompassing both stages, for which the applied researcher may be reluctant to hypothesize such a structure (Murphy and Topel, 2002). If the researcher does not want to assume a parametric model for ζ , the full-information Bayes must put a nonparametric prior $Q(\zeta)$ on it. Although the non-parametric Bayesian inference is one of the most vibrant research areas in econometric theory, verification of the related Bernstein–von Mises (BvM) theorem either relies on the Gaussian processes priors or explores particular model structures; see Ghosal and van der Vaart (2017). In addition, designing a feasible algorithm for the posteriors remains a skillful task when the joint likelihood functions from both stages are combined.

1.1. Related Literature

Studying asymptotic properties of quasi-Bayesian procedures forms an important line of work in econometrics. Regarding the discrepancy between quasi-Bayesian credible sets and frequentist confidence sets, we highlight a recurring theme when the generalized information identity fails (Chernozhukov and Hong, 2003). In this case, the local asymptotic normality (LAN) expansion (van der Vaart, 1998) generates a centering term whose asymptotic variance is of the sandwich form that does not match the (minus) second derivative matrix of the (quasi) log-likelihood function. These phenomena appear in studies of the generalized method of moments (GMM) or the limited information approach (Kim, 2002; Chernozhukov and Hong, 2003; Chib, Shin, and Simoni, 2018), as well as in misspecified models (Kleijn and van der Vaart, 2012; Müller, 2013; Kim, 2014). Given the natural link between GMM and two-step estimation (Newey and McFadden, 1994), it is tempting to think that our framework is nested by the works mentioned above. This is not the case. We emphasize that one distinct feature of our problem is that the first-stage estimation is based on a frequentist method and then plugged into the posterior of the second stage. One can view this as approximating the posterior of ζ using an infinitely narrow but tall spike at a given value, like the Dirac delta function. Such feature usually violates the standard assumptions in proving the BvM theorem.

Another line of research closely related to this paper is the profile sampler (Lee, Kosorok, and Fine, 2005; Cheng and Kosorok, 2008). This literature recommends sampling from the profile likelihood $pl_n(\theta) \equiv \sup_{\zeta} p_{\theta}^n(\mathbf{Y}^n; \zeta)$ for semiparametric models indexed by (θ, ζ) . The infinite-dimensional parameter ζ therein must be estimated via the profile maximum likelihood so that there is no additional adjustment term; also, see the last paragraph on page 1357 of Newey (1994) about the profiled out nuisance functions. Such procedures are not required in our setup. The crux of our analysis is to characterize the first-stage estimation effect on the quasi-Bayesian posteriors. On the technical ground, this paper also refines the proof of Lee, Kosorok, and Fine (2005) and Cheng and Kosorok (2008): we show the asymptotic negligibility of the posterior outside a shrinking neighborhood of the truth, where the radius depends on the convergence rate of first-stage estimation. In comparison, the result of Lemma A.1 of Lee, Kosorok, and Fine (2005) is stated for some neighborhood with a *fixed* radius, yet later on, this lemma is cited to cover the case with a shrinking radius of the order $o(n^{-1/3})$. Our proof relaxes this rate to match the well-known requirement at the order of $o(n^{-1/4})$ in general semiparametric models (Newey, 1994).

1.2. Organization

The rest of our paper is organized as follows: Section 2 contains the main theoretical results. This relatively long section is divided into three subsections, in which we introduce necessary assumptions, investigate the asymptotic behavior of the quasi-posterior for the finite-dimensional parameter in the second stage, and show the validity of bootstrap methods for inference. Section 3 applies the general theory to the Petrin–Train model. We present low-level regularity conditions and state the asymptotic results. Section 4 conducts Monte Carlo simulations and provides an empirical illustration. The last section concludes. Proofs of the main results are presented in Appendix A. More technical lemmas and their proofs are collected in Appendix B.

2. Main Theoretical Results

Bayesian methodology is attractive in its own right. Nonetheless, our problem is fundamentally non-Bayesian, as researchers aim to estimate unknown fixed parameters using the control function approach to correct the endogeneity bias. This motivates us to investigate the large sample behavior of the resulting quasi-posteriors under a fixed true probability model $P_0 = P_{\theta_0, \zeta_0}$ that generates the observed data. A thorough understanding will provide solid justification for the quasi-Bayesian methods, which can be attractive to non-Bayesian practitioners who use them because of their convenience.

We establish the asymptotic theory, drawing on two themes. The first is the frequentist two-step estimation and inference in econometrics literature (Newey, 1994; Chen, Linton, and Van Keilegom, 2003; Ichimura and Lee, 2010). The crux therein is how to characterize the influence of the first-step nonparametric estimation on the remaining parametric components and how to make proper adjustments for inference. This analysis is also central to our study, and we formally show how the quasi-posterior ignores the effect from the first-stage estimation. The second theme is the asymptotic analysis of Bayesian methods for semiparametric models. Our theoretical findings have not been reported before; they shed new light on the BvM theorem (van der Vaart, 1998; Ghosal and van der Vaart, 2017). For standard semiparametric models using *full information* Bayesian methods, the BvM theorem states that the marginal posterior for the finite-dimensional parameter is approximately a normal distribution centered at the semiparametric efficient estimator. Hence, the point estimators and credible sets are produced by one stroke therein. However, estimation and inference have to be dealt with separately for the quasi-posterior distribution in our setting.

2.1. Assumptions and Discussions

The (limited-information) log-likelihood function is of fundamental importance in the second-stage Bayesian estimation. Throughout the paper, we focus on the case with *i.i.d.* data, so the likelihood function becomes $p_\theta^n(\mathbf{Y}_2^n; \zeta) = \prod_{i=1}^n p_\theta(Y_{2,i}; \zeta)$ and the log-likelihood is denoted by $\ell_n(\theta; \zeta) \equiv \log p_\theta^n(\mathbf{Y}_2^n; \zeta) = \sum_{i=1}^n \log p_\theta(Y_{2,i}; \zeta)$. We define the following first-order derivative with respect to (w.r.t.) the finite-dimensional parameter θ :

$$(2.1) \quad \dot{\ell}_\theta(Y_2, \theta; \zeta) \equiv \left(\frac{\partial \log p(Y_2, \theta; \zeta)}{\partial \beta^\top}, \frac{\partial \log p(Y_2, \theta; \zeta)}{\partial \eta^\top} \right)^\top,$$

where θ is partitioned into the parameter of interest β and the nuisance parameter η . The metric associated with the parameter θ is $d_\Theta(\cdot, \cdot)$. The second-order derivative is denoted by $\ddot{\ell}_\theta(Y_2, \theta; \zeta)$ in the same vein. Thereafter, the score function is $S_\theta(\theta, \zeta) = P_0 \dot{\ell}_\theta(Y_2, \theta; \zeta)$, with the corresponding empirical version as $S_{\theta,n}(\theta, \zeta) = \mathbb{P}_n \dot{\ell}_\theta(Y_2, \theta; \zeta)$. The negative Hessian matrix $V_0 = -P_0 \ddot{\ell}_\theta(Y_2, \theta_0; \zeta_0)$ is the information matrix for the second-stage likelihood, when ζ_0 is known. The unknown function in the first stage may contain multiple components, so we write $\zeta_0 = (\zeta_{0j})_{j=0}^J$. Consider a map defined by $\Psi(\theta, \zeta) \equiv S_\theta(\theta, \zeta)$. The following pathwise derivative of Ψ plays a key role in determining the effect of the first-stage estimation (Newey, 1994):

$$\dot{\Psi}_\zeta(\theta_0, \zeta_0)[\zeta - \zeta_0] := \sum_{j=0}^J \dot{\Psi}_{\zeta_j}(\theta_0, \zeta_0)[\zeta_j - \zeta_{0j}],$$

where each $\dot{\Psi}_{\zeta_j}$ is a bounded linear functional that maps $\zeta_j - \zeta_{0j}$ to a real number. The proper functional space for ζ is denoted by \mathcal{G} with its pseudo metric $d_G(\cdot, \cdot)$.

We list all regularity conditions followed by some heuristic discussions in order. They are not necessarily the weakest possible assumptions. For a function $f(\cdot)$ of a random vector Y that follows distribution P , we use the standard empirical process notations: $P_0 f = \int f(y) dP_0(y)$, $\mathbb{P}_n f = n^{-1} \sum_{i=1}^n f(Y_i)$, and $\mathbb{G}_n f = n^{1/2} (\mathbb{P}_n - P_0) f$. We also write \mathbb{E}_0 instead of P_0 in bounding center stochastic terms. Definitions of the P_0 -Glivenko-Cantelli class and the P_0 -Donsker class follow those in van der Vaart and Wellner (1996) and van der Vaart (1998).

Assumption 2.1 (Identification). The function ζ_0 is uniquely identified from the first stage. The expected log-likelihood function $\mathbb{E}_0[\log p_\theta(Y_2; \zeta_0)]$ has a unique maximum at θ_0 over the parameter space Θ in the sense that $\sup_{d_\Theta(\theta, \theta_0) > \delta} \mathbb{E}_0[\log p_\theta(Y_2; \zeta_0)] < \mathbb{E}_0[\log p_{\theta_0}(Y_2; \zeta_0)]$ for any $\delta > 0$. Also, $S_\theta(\theta_0, \zeta_0) = 0$, and the matrix V_0 is positive definite with its smallest eigenvalue $\lambda_1 > 0$.

Assumption 2.2 (First-stage Estimation). The first-stage estimator $\widehat{\zeta}_n \in \mathcal{G}_n$ with probability 1, and $d_G(\widehat{\zeta}_n, \zeta_0) = O_{P_0}(r_n)$, where $r_n = o(n^{-1/4})$ and $nr_n^2 \rightarrow +\infty$.

Assumption 2.3 (Frequentist Two-stage Estimation). There exists a frequentist-type estimator that approximately maximizes the second-stage likelihood, i.e., $\widehat{\theta}_n = \arg \max_{\theta} \ell_n(\theta; \widehat{\zeta}_n) + o_{P_0}(n^{-1/2})$. In addition, it satisfies

$$(2.2) \quad S_{\theta,n}(\widehat{\theta}_n, \widehat{\zeta}_n) = o_{P_0}(n^{-1/2}).$$

Assumption 2.4 (Prior). The prior measure Π on Θ is assumed to be a probability measure with a bounded Lebesgue density π , which is continuous and positive on a neighborhood of the true value θ_0 . In addition, $\int |\theta|^a \pi(\theta) d\theta < +\infty$, for $a \geq 0$.

Assumption 2.5 (Smoothness). We assume the log-likelihood function satisfies the following restrictions for some positive constant terms C_1, \dots, C_4 :

$$\sup_{\zeta \in \mathcal{G}_n} |\mathbb{E}_0 [\log p_{\theta_0}(\cdot; \zeta) - \log p_{\theta_0}(\cdot; \zeta_0)]| \leq C_1 r_n^2,$$

and

$$\mathbb{E}_0 [\log p_{\theta}(\cdot; \zeta) - \log p_{\theta_0}(\cdot; \zeta_0)] + e_n(\theta, \zeta) \leq -C_2 d_{\Theta}^2(\theta, \theta_0) + C_3 d_G^2(\zeta, \zeta_0),$$

for some $e_n(\cdot, \cdot)$ that satisfies

$$\sup_{d_{\Theta}(\theta, \theta_0) \leq \rho, \zeta \in \mathcal{G}_n} |e_n(\theta, \zeta)| \leq C_4 \rho r_n.$$

Assumption 2.6 (Differentiability). The function $\mathbb{E}_0[\log p_{\theta}(Y_{2i}; \zeta)]$ is continuous with respect to ζ uniformly in $\theta \in \Theta$. The function $\Psi(\theta, \zeta)$ is Fréchet differentiable at (θ_0, ζ_0) , i.e. there exists a continuous and non-singular matrix $\dot{\Psi}_{\theta}(\theta_0, \zeta_0)$ and a continuous linear functional $\dot{\Psi}_{\zeta}(\theta_0, \zeta_0)$ such that

$$\|\Psi(\theta, \zeta) - \Psi(\theta_0, \zeta_0) - \dot{\Psi}_{\theta}(\theta_0, \zeta_0)(\theta - \theta_0) - \dot{\Psi}_{\zeta}(\theta_0, \zeta_0)[\zeta - \zeta_0]\| \leq o(d_{\Theta}(\theta, \theta_0)) + O(d_G^2(\zeta, \zeta_0)).$$

Assumption 2.7 (Complexity). (i) The functional classes $\{\log p_{\theta}(Y_{2i}; \zeta) : \theta \in \Theta, \zeta \in \mathcal{G}_n\}$ and $\{\ddot{l}_{\theta}(\cdot, \theta; \zeta) : \theta \in \Theta, \zeta \in \mathcal{G}_n\}$ belong to P_0 -Glivenko-Cantelli classes. (ii) The functional class $\{\dot{l}_{\theta}(Y_2, \theta_0; \zeta) : \zeta \in \mathcal{G}_n\}$ belongs to a P_0 -Donsker class. The following functional class

$$\{\log p_{\theta}(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta) - (\theta - \theta_0)^{\top} \dot{l}_{\theta_0}(\cdot; \zeta) : d_{\Theta}(\theta, \theta_0) \leq Cr_n, \zeta \in \mathcal{G}_n\}$$

divided by each ordinates of $(\theta - \theta_0)$ belongs to a P_0 -Donsker class with its second moment going to zero w.p.1. (iii) There exists some function $\phi_n(\cdot)$ such that

$$\mathbb{E}_0 \left[\sup_{d_{\Theta}(\theta, \theta_0) \leq \rho, \zeta \in \mathcal{G}_n} \mathbb{G}_n |\log p_{\theta}(\cdot, \theta; \zeta) - \log p_{\theta}(\cdot, \theta_0; \zeta)| \right] \leq \phi_n(\rho),$$

for $\rho \geq r_n$. $\phi_n(\cdot)$ is a sequence of functions defined on $(0, \infty)$ that satisfies: $\rho \mapsto \phi_n(\rho)/\rho^\gamma$ is decreasing for some $\gamma < 2$; $\rho^{-2}\phi_n(\rho) \leq \sqrt{n}$ for every n .

Assumption 2.8 (Normality). Let $\Delta_{n,0} \equiv \sqrt{n} \left(S_{\theta,n}(\theta_0, \zeta_0) + \dot{\Psi}_\zeta(\theta_0, \zeta_0)[\widehat{\zeta}_n - \zeta_0] \right)$. Assume that the following linear representation holds:

$$\Delta_{n,0} = \frac{1}{\sqrt{n}} \sum_{i=1}^n [\Gamma_2(Y_{2i}) + \Gamma_1(Y_{1i})] + o_{P_0}(1),$$

with the influence function $\Gamma_2(\cdot) + \Gamma_1(\cdot)$. Also, the following asymptotic normality holds:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n [\Gamma_2(Y_{2i}) + \Gamma_1(Y_{1i})] \Rightarrow \mathbb{N}(0, \Omega_0),$$

for some finite covariance matrix Ω_0 .

Assumptions 2.1 to 2.3 are generic, concerning the identifiability of the model parameters, the asymptotics of the first stage nonparametric estimation, as well as the existence of a frequentist two-stage estimator. The frequentist two-stage estimator $\widehat{\theta}_n$ merely serves as a theoretical device, as it will become the centering point in the local asymptotic normal (LAN) expansion (Ghosh and Ramamoorthi, 2002; Chernozhukov and Hong, 2003; Lee, Kosorok, and Fine, 2005). In our context of analyzing structural discrete choice models, $\widehat{\theta}_n$ is difficult to compute, which motivates the quasi-Bayesian procedure. The Assumption 2.4 on the prior density is also standard (Ghosh and Ramamoorthi, 2002; Chernozhukov and Hong, 2003).

Compared with the standard conditions for the BvM theorem (van der Vaart, 1998), our setting involves two complications due to the presence of the estimated control function. The first task is to ensure that the posterior probability concentrates in any small neighborhood of the true parameter uniformly over a set \mathcal{G}_n to which the estimated $\widehat{\zeta}_n$ belongs. The second distinction is that one needs an adjustment term in the LAN expansion, which accounts for the first-stage estimation error. Assumptions 2.5 and 2.7 (iii) are needed to ensure that the quasi-likelihood ratio decays exponentially fast for $\{\theta : d_\Theta(\theta, \theta_0) > \delta\}$ with a fixed radius $\delta > 0$, as well as to strengthen the exponential type decay for $\{\theta : Cr_n \leq d_\Theta(\theta, \theta_0) \leq \delta\}$, with the rate r_n stated in Assumption 2.2. The differentiability condition in Assumption 2.6 w.r.t. θ and ζ are needed to account for the two stages in our quasi-Bayesian method. Finally, Assumption 2.8 is used in the LAN expansion of the log-likelihood ratio.

2.2. Large Sample Behaviors of Quasi-Posteriors

Our main use of the quasi-posterior is to interpret it as a frequentist's device, similar to that of a sampling distribution. Also, we are interested in whether the mean of this quasi-posterior is a legitimate point estimator for θ_0 . Define the posterior density function of $h \equiv \sqrt{n}(\theta - \hat{\theta}_n)$ by

$$\tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n) \equiv \frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}^n(\mathbf{Y}_2^n|\hat{\zeta}_n)\pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})}{\int p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}^n(\mathbf{Y}_2^n|\hat{\zeta}_n)\pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})dh}.$$

Furthermore, the limiting normal density is

$$\pi_\infty(h) \equiv (2\pi)^{-p/2} |\det V_0|^{1/2} \exp\left\{-\frac{h^\top V_0 h}{2}\right\}.$$

In order to metrize the weak convergence, we consider the following ‘‘total variation of moments’’ norm (Chernozhukov and Hong, 2003) for a real-valued function f on \mathcal{S} as

$$\|f\|_{TVM(a)} \equiv \int_{\mathcal{S}} (1 + |h|^a) |f(h)| dh,$$

for some $a \geq 0$.

Theorem 2.1 (Posterior Measures). *Suppose Assumptions 2.1 to 2.8 hold. Then, we have*

$$(2.3) \quad \|\tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n) - \pi_\infty(h)\|_{TVM(a)} = o_{P_0}(1).$$

Consequently, the rescaled sequence of quasi-posteriors for θ converges in total variation to a standard normal measure, that is,

$$(2.4) \quad \sup_{\xi \in \mathbb{R}^p} |\Pi(\sqrt{n}V_0^{1/2}(\theta - \hat{\theta}_n) \leq \xi | \mathbf{Y}_2^n; \hat{\zeta}_n) - \Phi_p(\xi)| \rightarrow_{P_0} 0.$$

where $\Phi_p(\cdot)$ denotes the p -dimensional standard normal distribution.

Let $\mathcal{C}_n(\alpha)$ be the quasi-Bayesian credible set constructed from the quasi-posterior, that is, $\mathcal{C}_n(\alpha)$ satisfies $\Pi(\theta \in \mathcal{C}_n(\alpha) | \mathbf{Y}_2^n; \hat{\zeta}_n) = 1 - \alpha$ for a given nominal level $\alpha \in (0, 1)$. An important implication of Theorem 2.1 is that $\mathcal{C}_n(\alpha)$ does not have the correct asymptotic coverage probability in general, as stated in Corollary 2.1.

Corollary 2.1. *Let B_n be the set that satisfies $(2\pi)^{-p/2} \int_{B_n} e^{-h^\top h/2} dh \rightarrow 1 - \alpha$ as $n \rightarrow \infty$, and recall the definition of $\Delta_{n,0}$ in Assumption 2.8. Under the same assumptions for Theorem 2.1, we have*

$$(2.5) \quad P_0\{\theta_0 \in \mathcal{C}_n(\alpha)\} = P_0\{-V_0^{-1/2}\Delta_{n,0} \in B_n\}.$$

In addition, if $V_0 = \Omega_0$, then we have

$$(2.6) \quad P_0\{\theta_0 \in \mathcal{C}_n(\alpha)\} \rightarrow 1 - \alpha.$$

Theorem 2.1 may look similar to the standard parametric BvM theorem at first glance. However, the departures are the presence of the first-stage estimation $\widehat{\zeta}_n$ and the centering point $\widehat{\theta}_n$, which is a two-stage frequentist estimator whose asymptotic covariance matrix takes the sandwich form $V_0^{-1}\Omega_0V_0^{-1}$, where Ω_0 corresponds to the variance of $\Gamma_2(Y_{2i}) + \Gamma_1(Y_{1i})$ in Assumption 2.8. On the other hand, as Corollary 2.2 below shows, the variance of the quasi-posterior of θ is governed by the information matrix V_0 , which only captures the variance of $\Gamma_2(Y_{2i})$. The discrepancy between Ω_0 and V_0 prevents the right hand side of (2.5) from converging to the desired coverage probability $1 - \alpha$ in general. With the additional restriction $V_0 = \Omega_0$, which corresponds to the case where the generalized information equality holds (Chernozhukov and Hong, 2003), $\mathcal{C}_n(\alpha)$ will have the correct coverage probability coincidentally, as (2.6) shows. The general coverage failure of the quasi-Bayesian credible set is similar to what Kleijn and van der Vaart (2012) found when they studied the Bayesian procedure for misspecified models. Considering the formal decision theory of the interval estimation problem, Müller (2013) showed that the asymptotic risk of a vanilla posterior is worse than that of a modified posterior using the sandwich covariance matrix; see Section 2.4 of Müller (2013). Below, we show that the quasi-posterior mean is asymptotically equivalent to the frequentist's two-stage estimator $\widehat{\theta}_n$. However, the rescaled posterior variance is not the same as the asymptotic variance of the frequentist two-stage estimator.

Corollary 2.2. *Denote the mean and variance of the quasi-posterior by $\widetilde{\theta}_n$ and $\widetilde{\text{Var}}_n(\theta)$. Assume that the prior for θ satisfies $\int |\theta|^2 d\Pi(\theta) < +\infty$. Then we have*

$$(2.7) \quad \widetilde{\theta}_n = \widehat{\theta}_n + o_{P_0}(n^{-1/2}), \text{ and } n^{-1}\widetilde{\text{Var}}_n(\theta)^{-1} = V_0^{-1} + o_{P_0}(1).$$

As a consequence of the first equality in (2.7), the quasi-posterior mean is asymptotically normal:

$$\sqrt{n}(\widetilde{\theta}_n - \theta_0) \Rightarrow \text{N}(0, V_0^{-1}\Omega_0V_0^{-1}).$$

Remark 2.1. *In the context of GMM, Chernozhukov and Hong (2003) constructed a quasi-likelihood by exponentiating the quadratic criterion function. Chernozhukov and Hong (2003) demonstrated that the weighting matrix can be properly chosen so that the generalized information identity holds, and the resulting posterior distribution can be utilized for asymptotically valid confidence sets. When the generalized information equality fails, Chernozhukov and Hong (2003) also suggested that the posterior still contains useful information for inference, as the posterior variance is related to the Hessian matrix. If Ω_0 is easier to*

obtain, one can form the normal confidence interval by plugging in the sandwich-type covariance matrix estimator, which combines the posterior variance and some external estimand for Ω_0 . Our motivating example does not belong to this case, as it is not computationally easy to obtain a consistent estimator for Ω_0 .

2.3. Validity of Bootstrap Inference

Now, we present the validity of a proper bootstrap inferential procedure that takes into account the first-stage estimation error. We show that the bootstrap point estimator can mimic the asymptotic behavior of the quasi-Bayesian point estimator. The essence of this analysis lies in the bootstrap likelihood with multinomial weights. The non-trivial part is to deal with the infinite-dimensional ζ from the first stage. Our theory also connects the bootstrap procedure suggested by Agarwal and Somaini (2018) to a more general context with two-stage semiparametric estimation. In principle, the theory works for other exchangeable weights (van der Vaart and Wellner, 1996). However, the most convenient approach is Efron's nonparametric bootstrap (Efron, 1979) because it only requires generating a sequence of new data sets by sampling the rows of the original data with replacement and then updating the posterior with the same prior. It is straightforward to compute all bootstrap calculations in parallel.

Let the bootstrap weights $\mathbf{M}_n = (M_{n1}, \dots, M_{nn})^\top$ follow the multinomial distribution $Mult(n, (n^{-1}, \dots, n^{-1}))$ (Efron, 1979). We also write $\mathbf{Y}_2^{n*} = (\mathbf{Y}_2^n, \mathbf{M}_n)^\top$. The resulting marginal posterior distribution on \mathcal{B} is given by the Bayes formula:

$$(2.8) \quad \Pi^*[\theta \in A | \mathbf{Y}_2^{n*}; \hat{\zeta}_n^*] = \frac{\int_A p_\theta^{n*}(\mathbf{Y}_2^n; \hat{\zeta}_n^*) d\Pi(\theta)}{\int_{\Theta} p_\theta^{n*}(\mathbf{Y}_2^n; \hat{\zeta}_n^*) d\Pi(\theta)},$$

where $p_\theta^{n*}(\mathbf{Y}_2^n; \zeta) \equiv \prod_{i=1}^n p_\theta^{M_{ni}}(Y_{2,i}; \zeta)$ stands for the bootstrap likelihood function and $\hat{\zeta}_n^*$ stands for the first-stage bootstrap estimator. The bootstrap quasi-Bayesian point estimator is

$$(2.9) \quad \tilde{\theta}_n^* = \int_{\Theta} \theta d\Pi^*(\theta | \mathbf{Y}_2^{n*}; \hat{\zeta}_n^*).$$

Also, denote the frequentist-type two-stage bootstrap estimator by $\hat{\theta}_n^*$. We denote the bootstrap empirical measure as \mathbb{P}_n^* and the bootstrap empirical process as $\mathbb{G}_n^* = \sqrt{n}(\mathbb{P}_n^* - \mathbb{P}_n)$. For a sequence of random variables Z_n^* , we write $Z_n^* = o_{P^*}(1)$ if the law of Z_n^* is governed by the bootstrap law \mathbb{P}^* and if $\mathbb{P}^*(|Z_n^*| > \epsilon) = o_{P_0}(1)$ for any $\epsilon > 0$.

To construct the confidence interval for $\theta_{j,0}$, which denotes the j th coordinate of the parameter θ_0 , we calculate the q -th quantile of the bootstrap distribution of $\tilde{\theta}_{n,j}^*$ as $Q_{n,j}^*(q) \equiv \inf\{\tau : \mathbb{P}^*(\tilde{\theta}_{n,j}^* \leq \tau) \geq q\}$. Then the percentile-type bootstrap $1 - \alpha$ confidence interval for

$\theta_{j,0}$ can be formed as

$$(2.10) \quad \mathcal{C}_n^{*(j)}(\alpha) = [Q_{n,j}^*(\alpha/2), Q_{n,j}^*(1 - \alpha/2)], \text{ for } j = 1, \dots, p.$$

Theorem 2.2 (Bootstrap Consistency). *In addition to Assumptions 2.1 to 2.8, we assume that the envelope functions for the functional classes in Parts (i) and (ii) in Assumption (2.7) have finite $2 + \iota$ moments, for $\iota > 0$. We have the following equivalence between the bootstrap posterior mean and the frequentist-type two-stage bootstrap estimator: $\tilde{\theta}_n^* = \hat{\theta}_n^* + o_{P^*}(n^{-1/2})$. Consequently, we have $P_0 \left\{ \theta_{0j} \in \mathcal{C}_n^{*(j)}(\alpha) \right\} \rightarrow 1 - \alpha$, for $j = 1, \dots, p$.*

3. The Petrin–Train Model

We demonstrate our theory by studying a class of endogenous discrete choice models originally proposed by Petrin and Train (2010). In particular, we revisit the Petrin–Train model with normal latent errors, which does not enforce the independence of irrelevant alternatives (IIA) assumption. The joint normality of the latent errors in both stages naturally leads to the identification strategy using the control function approach. We further relax the first-stage specification in Petrin and Train (2010) for flexible nonparametric estimation. Our study provides the theoretical foundation for the quasi-Bayes estimation and inferential tools. To the best of our knowledge, this has not been formally developed.

Let $(U_{ij})_{j=0,1,\dots,J}$ denote the latent utility of choice j for individual i , with the following linear specification:

$$(3.1) \quad U_{ij} = X_{ij}^\top \tilde{\beta}_0 + \varepsilon_{ij}, \quad i = 1, \dots, n, j = 0, \dots, J.$$

The individual i 's choice $C_i \in \{1, \dots, J\}$ is determined by

$$(3.2) \quad C_i = \sum_{j=1}^J j \mathbb{I} \{ U_{ij} > U_{ik} \text{ for } k \neq j \text{ and } U_{ij} > 0 \}.$$

And $C_i = 0$ if $U_{ij} \leq 0$ for all $j = 1, \dots, J$. Among the choice-specific covariates X_{ij} , some components X_{ij}^e can be correlated with the error term ε_{ij} . The endogenous regressors X_{ij}^e admit the following representation:

$$(3.3) \quad X_{ij}^e = \zeta_j(Z_{ij}) + v_{ij}, \quad j = 0, \dots, J,$$

where Z_{ij} collects the instrumental variables and other exogenous regressors, and v_{ij} stands for the unobservable error. The functions ζ_j 's are completely unspecified, except for some standard smoothness restrictions as detailed later. Following Example 1 of Petrin and Train (2010), the joint normality of latent errors ε_{ij} and v_{ij} leads to the decomposition of

ϵ_{ij} :

$$(3.4) \quad \epsilon_{ij} = \mathbb{E}[\epsilon_{ij} | v_{ij}] + \epsilon_{ij} = \lambda_0 v_{ij} + \epsilon_{ij}, \quad j = 0, \dots, J,$$

where $\lambda_0 v_{ij}$ is the control function², and the the reminder errors $(\epsilon_{ij})_{j=0, \dots, J}$ are also jointly normal. Using (3.4), the latent utility of the second stage can be written as:

$$(3.5) \quad U_{ij} = X_{ij}^\top \tilde{\beta}_0 + \lambda_0 v_{ij} + \epsilon_{ij}, \quad j = 0, \dots, J.$$

Taking difference relative to U_{i0} leads to

$$(3.6) \quad U_{ij} - U_{i0} = (X_{ij} - X_{i0})^\top \beta_0 + \lambda_0 (v_{ij} - v_{i0}) + \epsilon_{ij} - \epsilon_{i0}, \quad i = 1, \dots, n, j = 1, \dots, J.$$

The conditional choice probability for the j -th alternative in the second stage involves a multivariate integral:

$$(3.7) \quad \begin{aligned} & P_0(C_i = j | X_i, Z_i) \\ &= \int \mathbb{I} \{ \tilde{\epsilon}_{ikj} < (X_{ij} - X_{ik})^\top \beta_0 + \lambda_0 (v_{ij} - v_{ik}), \forall k \neq j \} dG(\tilde{\epsilon}_{ij}), \end{aligned}$$

where $\epsilon_{ikj} \equiv \epsilon_{ik} - \epsilon_{ij}$ for all $k \neq j$, and $\tilde{\epsilon}_{ij}$ denotes the vector $(\epsilon_{ikj})_{k \neq j}$. The function G is a J -dimensional normal distribution with mean zero and covariance matrix $M_j \Gamma_0(\eta_0) M_j^\top$, where M_j is a linear transformation that adds a column of -1 to the J -dimensional identity matrix, and $\Gamma_0(\eta_0)$ is the covariance matrix of $\epsilon_i \equiv (\epsilon_{ij})_{j=0, 1, \dots, J}$ that depends on the nuisance parameter η_0 . We collect the regression coefficient as $\beta_0^\top = (\tilde{\beta}_0^\top, \lambda_0)$.

The quasi-Bayesian estimation applied to the Petrin–Train model consists of two stages. The first stage estimates the functions ζ_j nonparametrically (e.g., kernel regression) and then obtains the residuals \hat{v}_{ij} for $j = 0, 1, \dots, J$. The second stage corresponds to the MNP model (3.6) with $(v_{ij} - v_{i0})$ replaced by the first stage residual estimates. Due to the analytically intractable conditional choice probability (3.7), the Bayesian approach becomes more appealing as it explores the conditional conjugate structure induced by MNP. The posterior of β can be drawn using Markov Chain Monte Carlo (MCMC) algorithms coupled with data augmentation techniques (Albert and Chib, 1993; McCulloch and Rossi, 1994; Nobile, 1998; Imai and Van Dyk, 2005).³ The point estimator for β can be constructed as the posterior mean. We use bootstrap to construct the confidence interval. Resampling the original data with replacement forms a bootstrap sample. For each bootstrap sample, we apply the first-stage frequentist and the second-stage Bayesian procedures to obtain the bootstrap quasi-posterior mean $\tilde{\beta}_n^*$ as in (2.9). Repeatedly drawing the bootstrap

²In principle, we can allow for different $\lambda_{0,j}$ over $j = 1, \dots, J$. In this case, the model can be written in the form of $U_{ij} = X_{ij}^\top \beta_0 + \tilde{v}_{ij}^\top \lambda_0 + \epsilon_{ij}$, by generating proper longer vectors of \tilde{v}_{ij}^\top and $\lambda_0^\top = (\lambda_{0,j})_{j=1, \dots, J}$. For ease of exposition, we do not consider such complications.

³Chapter 5 of Train (2009) presents a comprehensive discussion about the computational advantages of Bayesian methods in analyzing MNP-type models.

sample and calculating the bootstrap point estimator for many times leads to a bootstrap distribution of $\tilde{\beta}_n^*$, whose quantiles form the bootstrap confidence intervals described in (2.10). Note that the second-stage Bayesian procedure does not require maximization of any function. This reduces the computational burden of bootstrapping the quasi-Bayesian estimator compared to bootstrapping a frequentist two-stage estimator whose second stage involves maximization of the simulated likelihood function associated with the MNP model.

To spell out the regularity conditions that are sufficient to establish the asymptotic results, we introduce the following notations. Denote the collection of vectors across non-baseline choices by $\mathbf{U}_i = [U_{i1} - U_{i0}, \dots, U_{iJ} - U_{i0}]^\top$. We write

$$\mathbf{W}_i = \begin{bmatrix} W_{i1} & 0 & \cdots & 0 \\ 0 & W_{i2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & W_{iJ} \end{bmatrix} \quad \text{and} \quad \mathbf{W}_i^\top \beta = \begin{bmatrix} W_{i1}^\top \beta \\ W_{i2}^\top \beta \\ \vdots \\ W_{iJ}^\top \beta \end{bmatrix},$$

where $W_{ij}^\top = ((X_{ij} - X_{i0})^\top, v_{ij} - v_{i0})$. Let $R(C_i)$ specify the truncation region associated with the choice variable. If $C_i = 0$, $R(C_i)$ consists of the region such that each component of the latent utility is negative. For $C_i = j \neq 0$, $R(C_i)$ restricts U_{ij} to be positive and greater than all other $U_{ik}, k \neq j$. Referring to the expression (2.1), the score functions w.r.t. β and η take the following forms:

$$\frac{\partial \log p(Y_2, \theta; \zeta)}{\partial \beta^\top} = \mathbf{W}_i^\top \Sigma^{-1} \mathbb{E}[\mathbf{U}_i - \mathbf{W}_i^\top \beta | \mathbf{U}_i \in R(C_i)],$$

$$\frac{\partial \log p(Y_2, \theta; \zeta)}{\partial \eta^\top} = \text{vec}(\Sigma^{-1} (\mathbb{I}_{J \times J} - \mathbb{E}[(\mathbf{U}_i - \mathbf{W}_i^\top \beta)^{\otimes 2} | \mathbf{U}_i \in R(C_i)] \Sigma^{-1}) \Gamma(\eta)) \partial(\text{vec}(\Gamma(\eta))) / \partial \eta$$

cf. equations (4) and (5) of Hajivassiliou, McFadden, and Ruud (1996). Clearly, the score functions also involve multivariate integrals, which makes the analytical correction for the first-stage estimation error difficult.

For any $\alpha \in \mathbb{R}$, let $\lfloor \tau \rfloor$ be its integer part. Considering a multi-index $k = (k_1, \dots, k_d)$, define $k \cdot = k_1 + k_2 + \dots + k_d$. For any $L > 0, \tau \geq 0$ and nonnegative function L on \mathbb{R}^d , define the τ -Hölder class with envelope L , denoted by $\mathcal{C}^{\tau, L}(\mathbb{R}^d)$ (Section 3 of Ichimura and Lee (2010)) to be the set of all functions $f : \mathbb{R}^d \mapsto \mathbb{R}$ that have finite mixed partial derivatives $D^k f$ of all orders up to $k \leq \lfloor \tau \rfloor$, such that for very $k \in \mathbb{N}_0^d$ with $k \cdot = \lfloor \tau \rfloor$,

$$(3.8) \quad |(D^k f)(x + y) - (D^k f)(x)| \leq L \|y\|^{\tau - \lfloor \tau \rfloor}, \quad \forall x, y \in \mathbb{R}^d.$$

In Assumption 3.4 below, \mathbb{W}_J denotes the Wishart distribution for $J \times J$ positive-definite random matrices.

Assumption 3.1. The parameter space Θ is a compact set, and the true θ_0 is in the interior of Θ . The latent error's covariance matrix $\Gamma_0(\eta_0)$ has its first element normalized to be 1, and it is non-singular. The matrix $V_0 = -P_0 \ddot{l}_\theta(Y_2, \theta_0; \zeta_0)$ is positive definite.

Assumption 3.2. The supports of covariates X and the control variable v are bounded. We assume the function $\zeta_{0j} \in \mathcal{C}^{\tau, L}(\mathbb{R}^d)$, for $j = 0, \dots, J$ with $\tau > d/2$. The functions $\zeta_{0j}, j = 0, \dots, J$ are uniquely identified from the first stage.

Assumption 3.3. For the first-stage estimator, we assume its convergence rate with respect to the supnorm is $r_n = o(n^{-1/4})$. Furthermore, it has the following linear representation:

$$(3.9) \quad \widehat{\zeta}_{n,j}(Y_{1i}) - \zeta_{0j}(Y_{1i}) = \frac{1}{n} \sum_{l=1}^n \phi_{n,j}(Y_{1l}) + b_{n,j}(Y_{1i}) + R_{n,j}(Y_{1i}), \quad 0 \leq j \leq J,$$

where $\phi_{n,j}(\cdot)$ is a stochastic term that has expectation zero, $b_{n,j}$ is a bias term satisfying $\max_j \sup_{y_1} |b_{n,j}(y_1)| = o(n^{-1/2})$ and $\max_j \sup_{y_1} |R_{n,j}(y_1)| = o_{P_0}(n^{-1/2})$.

Assumption 3.4. Priors for the finite-dimensional parameters follow the Gaussian–Wishart type:

$$(3.10) \quad \tilde{\beta} \sim \mathbb{N}(\mu_\beta, V_\beta), \text{ and } \Sigma^{-1} \sim \mathbb{W}_J([\rho V_0]^{-1}, \rho) \mathbb{I}\{\sigma_1^2 = 1\},$$

in which the scalar ρ represents the degree of freedom of the Wishart distribution⁴ and (μ_β, V_β) specify the mean and variance of the Gaussian distribution.

Assumption 3.5. The partial derivatives of $\dot{l}_\theta(\cdot, \theta; \zeta)$ with respect to control variables $(v_{ij} - v_{i0})_{j=1 \dots J}$ are uniformly bounded over the support of the covariates and control variables $(W_{ij})_{j=1 \dots J}$. In addition, $\dot{l}_\theta(\cdot, \theta; \zeta)$ is locally uniformly $L_2(P_0)$ -continuous with respect to θ and ζ in the sense that for all small positive $\delta = o(1)$,

$$(3.11) \quad \mathbb{E} \left[\sup_{d_\Theta(\theta, \theta_0) < \delta, d_G(\zeta, \zeta_0) < \delta} \left| \dot{l}_\theta(Y_2, \theta; \zeta) - \dot{l}_\theta(Y_2, \theta; \zeta) \right|^2 \right] \leq C\delta^2.$$

Furthermore, the second-order derivative $\{\ddot{l}_\theta(y_2, \theta; \zeta) : \theta \in \Theta, \zeta \in \mathcal{G}_n\}$ is upper semi-continuous for almost all y_2 and has an integrable envelope function.

Assumption 3.6. We express the pathwise derivative with respect to the first-stage estimation as

$$(3.12) \quad \dot{\Psi}_{\zeta_j}(\theta_0, \zeta_0)[\mathbf{h}] = \int \mathbf{h}(\cdot) g_j(\cdot, \theta) dP_0,$$

⁴The indicator function on the Wishart prior serves to enforce the identification restriction that the (1, 1) element of Σ is unity; see McCulloch, Polson, and Rossi (2000). This restriction is made for the identification purpose.

for some function $g_j(\cdot, \theta)$, with $j = 0, \dots, J$. We further assume that $\int \phi_{n,j}(\cdot, \theta_0) \frac{\partial g_j(\cdot, \theta_0)}{\partial \theta} dP_0$ has a finite second moment for $0 \leq j \leq J$.

The following proposition specializes Theorems 2.1 and 2.2 to the Petrin–Train Model. It establishes the asymptotic normality of the quasi-Bayesian point estimator and the asymptotic validity of the bootstrap confidence set.

Proposition 3.1. *Under Assumptions 3.1 to 3.6, the quasi-posterior mean is asymptotically normal:*

$$\sqrt{n}(\tilde{\theta}_n - \theta_0) \Rightarrow \mathbb{N}(0, V_0^{-1} \Omega_0 V_0^{-1}).$$

Moreover, the bootstrap confidence intervals are asymptotically valid.

Throughout the paper, we maintain the point identification assumption of all finite-dimensional parameters in the MNP model. This is also assumed in the classical literature on frequentist estimators (McFadden, 1989; Pakes and Pollard, 1989). When the identification fails, one might consider extending our approach to the direction of Chen, Christensen, and Tamer (2018) and develop a valid inferential procedure under partial identification. This is beyond the scope of the current work and will be pursued elsewhere. When it comes to the first-stage nonparametric estimation, one may use the kernel smoothing type estimator, including the local constant or local polynomial estimators in Chen, Linton, and Van Keilegom (2003); Ichimura and Lee (2010), or the sieve type estimators as displayed in Chen (2007). Regarding the prior choice, one can also impose the restriction on the eigenvalue rather than the first diagonal element as in Imai and Van Dyk (2005). Examples of other proper normalizations for the unrestricted or restricted covariance can be found in Train (2009, Chapter 5.2). We also refer interested readers to Anceschi, Fasano, Durante, and Zanella (2023) for a recent review on the contemporary development of various fast and scalable MCMC or deterministic algorithms.

4. Numerical Results

4.1. Monte Carlo Simulation

We conduct Monte Carlo simulations to show that the quasi-Bayesian credible set does not yield desirable coverage probabilities. However, this problem can be solved by bootstrapping. Our simulation design is a multinomial choice model with three or four choices and one endogenous regressor for each choice. The latent utility takes the form:

$$(4.1) \quad U_{ij} = \tilde{\beta} X_{ij}^e + \varepsilon_{ij}, \quad j = 0, 1, \dots, J.$$

The true value of $\tilde{\beta}$ is 1. The endogenous regressor X_{ij}^e depends on the instrument Z_{ij} as follows:

$$(4.2) \quad X_{ij}^e = \zeta(Z_{ij}) + v_{ij}, \quad j = 0, 1, \dots, J.$$

The instruments Z_{i0}, \dots, Z_{iJ} and error terms v_{i0}, \dots, v_{iJ} in the first stage (4.2) are independent standard normal random variables; the error term ε_{ij} in the latent utility (4.1) is generated by $\varepsilon_{ij} = \lambda v_{ij} + \epsilon_{ij}$ for $j = 0, 1, \dots, J$, where the centered errors $(\epsilon_{ij}, j = 0, 1, \dots, J)$ are jointly normal with means equal to zero, variances $\sigma_j^2 = 1$ if $j = 0$ and 0.5 if $j > 0$. In the case of three choices ($J = 2$), the correlation coefficient for any pair of errors is $\text{corr}(\epsilon_{ij}, \epsilon_{is}) = \rho$ for $j \neq s$. In the case of four choices ($J = 3$), $\text{corr}(\epsilon_{i0}, \epsilon_{i1}) = \text{corr}(\epsilon_{i2}, \epsilon_{i3}) = \rho$, and $\text{corr}(\epsilon_{ij}, \epsilon_{is}) = 0$ for other (j, s) pairs. We set $\rho = 0.4$ and $\lambda = 0.6$. We consider two functional forms for $\zeta(\cdot)$ in (4.2): (I). $\zeta(z) = 0.9z + 0.9z^2 + \ln(z + 1)^2$; (II). $\zeta(z) = 0.9z + 0.9z^2 + \exp(0.9z)$. Our first stage estimate for ζ uses a kernel regression with the bandwidth chosen by leave-one-out cross-validation. This step is implemented by the R package `np`. Controlling for the estimated $\hat{v}_{ij} = X_{ij}^e - \hat{\zeta}(Z_{ij})$, we draw the posterior of $\tilde{\beta}$ from the MNP model in the second stage, using Gibbs sampler with data augmentation (Imai and Van Dyk, 2005). This step uses the R package `MNP`.

Table 1 presents the empirical coverage probabilities and the average lengths for the quasi-Bayesian credible interval (QB CI) and the bootstrap quasi-Bayesian confidence interval (BQB CI). The sample size is fixed at 1,000. When simulating the posterior, the initial 2,000 Gibbs draws are discarded, and the following 5,000 draws are stored. BQB CI is constructed following (2.10). The number of bootstrap repetitions is 500. We make the following observations. First, QB CI systematically under-covers the true parameter $\tilde{\beta}$ in all occasions. Consider the nominal coverage $1 - \alpha = 0.90$ as an example; its coverage probabilities are around 0.8 in three of the four scenarios and about 0.85 in the last scenario. Second, the bootstrap version BQB CI significantly improves the coverage performance. Again, considering $\alpha = 0.10$, the empirical coverage probability of BQB CI ranges between 88% and 92%. Third, the average lengths of BQB CI are longer than those of QB CI, as the former incorporates the estimation uncertainty from the first stage.


Results in Table 1 are computed using an improper non-informative prior distribution on the coefficient $\tilde{\beta}$. Figure 1 illustrates how the coverage probabilities are affected by the variance of the prior distribution. In particular, we consider the prior variance for $\tilde{\beta}$ equal to 100, 1,000, and $+\infty$ (Table 1 uses $+\infty$). The prior mean for $\tilde{\beta}$ is set as zero in all cases. Figure 1 shows that the under-coverage of QB and the improvement achieved by BQB also occur under more informative priors.

Overall, our simulation findings suggest that the under-coverage of the quasi-Bayesian credible interval is ubiquitous for different first-stage relationships, number of choices, and

TABLE 1. Coverage probabilities and average lengths of quasi-Bayesian (QB) credible intervals and bootstrapped quasi-Bayesian (BQB) confidence intervals for the coefficient β , $n = 1,000$, nominal coverage $1 - \alpha = 0.90, 0.95$, and 0.99 .

# of choices	$\zeta(\cdot)$ form	Methods	Coverage probability			Average length		
			0.90	0.95	0.99	0.90	0.95	0.99
3	I	QB	0.813	0.882	0.954	0.385	0.465	0.628
		BQB	0.888	0.932	0.978	0.454	0.555	0.784
	II	QB	0.819	0.882	0.948	0.367	0.447	0.609
		BQB	0.887	0.932	0.984	0.459	0.564	0.791
4	I	QB	0.794	0.874	0.952	0.470	0.578	0.771
		BQB	0.878	0.941	0.990	0.634	0.696	1.043
	II	QB	0.847	0.914	0.972	0.479	0.597	0.774
		BQB	0.921	0.962	0.996	0.692	0.752	1.099

FIGURE 1. Coverage probability of quasi-Bayesian (QB) credible intervals (gray, dashed) and bootstrapped quasi-Bayesian (BQB) confidence intervals (blue, solid), with the prior variance for β being 100, 1,000 and ∞ , and the nominal coverage probability ($1 - \alpha$) being 0.90 (\bullet), 0.95 (\blacksquare), and 0.99 (\blacktriangle); # of choices = 3, $n = 1,000$.



prior variance values. On the other hand, our recommended bootstrap procedure successfully restores the coverage probabilities to the nominal levels.

4.2. Empirical Application

We apply the quasi-Bayesian approach to a real dataset about the firms’ incorporation decisions, initially constructed by Eldar and Magnolfi (2020). We focus on the effect of states’ anti-takeover laws on firms’ choice of where to incorporate their business. The data includes an *ATS* index scoring from 0 to 5 that counts the number of anti-takeover statutes in each of 51 states (including the District of Columbia). Two states dominate the incorporation market in the US: in 2013, 63.9% firms were incorporated in Delaware, and 8.5% in Nevada. Letting i denote firm and j denote the incorporation choice, we consider an MNP model with three choices: Delaware ($j = 1$), Nevada ($j = 2$), and other states ($j = 0$). The regressors contain *ATS* and its interaction with firm-level characteristics. We use the data in the year 2013, which contains 2,922 firms. We specify the firm’s latent utility as follows:

$$(4.3) \quad U_{ij} = \tilde{\beta}_1 ATS_j + (ATS_j \times X_i)^\top \tilde{\beta}_2 + \varepsilon_{ij}, \quad j = 0, 1, 2.$$

where X_i is a vector of firm characteristics: share of institutional ownership (*Institutional ownership*), dummies for small and median-sized firms (*small* and *medium*), and takeover premium (*premium*) in the industry. The main concern for the endogeneity arises from institutional ownership. Following Eldar and Magnolfi (2020), we use the dummy instrumental variable S&P 500 that indicates whether the firm is included in the S&P 500 index. Inclusion in the S&P 500 index is positively correlated with a firm’s institutional ownership. On the other hand, S&P 500 is mainly decided by the market’s view about the firm’s representativeness and thus does not depend on the firms’ decisions. The first-stage is

$$(4.4) \quad Institutional\ ownership_i = \zeta(S\&P500, X_{-1i}) + v_i,$$

where ζ is an unknown function and $X_{-1i} = \{Small_i, Medium_i, Premium_i\}$. The control function v_i decomposes the error term in (4.3) by $\varepsilon_{ij} = \lambda_j v_i + \epsilon_{ij}$. We estimate the first stage (4.4) by a kernel regression and then draw the posteriors of $\tilde{\beta}_1$ and $\tilde{\beta}_2$ in the same way as our simulation exercise. The point estimates in Table 2 are posterior means, and the 95% confidence intervals reported in brackets are constructed using bootstrap. The significantly negative estimate for $\tilde{\beta}_1$ suggests that firms on average prefer less legal restrictions on the takeover. Moreover, such preference is heterogeneous in firm characteristics. In particular, small firms are more likely to welcome the anti-takeover laws, which explains why Nevada, with $ATS = 5$, occupies a notable market share for incorporation. Our results align with the findings of Eldar and Magnolfi (2020), who estimated a multinomial logit model.

TABLE 2. Quasi-Bayesian estimates for incorporation decision, bootstrapped 95% confidence intervals in brackets, $n = 2,922$.

<i>ATS</i>	-0.065 [-0.108, -0.018]
<i>ATS</i> \times <i>Institutional ownership</i>	-0.253 [-0.311, -0.200]
<i>ATS</i> \times <i>Small</i>	0.080 [0.024, 0.127]
<i>ATS</i> \times <i>Medium</i>	-0.005 [-0.040, 0.030]
<i>ATS</i> \times <i>Premium</i>	-0.043 [-0.072, -0.012]

5. Conclusion

Our research highlights two key aspects of modern econometric models. Firstly, the complexity of these models leads to analytically intractable likelihoods, which can be more conveniently dealt with by simulation-based procedures like the Bayesian method. Secondly, the control function approach effectively handles the endogeneity problem by including the first-stage residuals as additional covariates, rectifying the second stage’s inconsistency. Our study extensively investigates this novel quasi-Bayesian method. Given the widespread use of the Bayesian approach in contemporary econometrics and machine learning, our method provides practitioners with enhanced flexibility to integrate cutting-edge algorithms from various estimation stages. This methodology can be beneficial for other structural models when the general theory is well understood.

6. Appendix A: Proofs of Main Results

Proof of Theorem 2.1. Our proof of (2.3) is patterned in line with generic arguments in the proof of Theorem 1.4.2 from Ghosh and Ramamoorthi (2002) or Theorem 1 from Chernozhukov and Hong (2003). The crux is to address the first-stage nonparametric estimator $\widehat{\zeta}_n$. This is reflected in a different choice of splitting ranges of the integration and a more delicate expansion of the quasi-log-likelihood ratio in our proof.

Consider the following normalized quasi-posterior distribution of $h \equiv \sqrt{n}(\theta - \hat{\theta}_n)$:

$$\begin{aligned}\tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n) &= \frac{\prod_{i=1}^n p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)\pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})}{\int \prod_{i=1}^n p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)\pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})dh} \\ &= \frac{\prod_{i=1}^n \left[\frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n)} \right] \pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})}{\int \prod_{i=1}^n \left[\frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n)} \right] \pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})dh}.\end{aligned}$$

Denote its denominator as

$$(6.1) \quad C_n \equiv \int \prod_{i=1}^n \left[\frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n)} \right] \pi(\hat{\theta}_n + \frac{h}{\sqrt{n}})dh.$$

Recall that the limiting normal probability density function is as follows:

$$\pi_\infty(h) \equiv (2\pi)^{-p/2} |\det V_0|^{1/2} \exp \left\{ -\frac{h^\top V_0 h}{2} \right\}.$$

The total variation norm we employ can be bounded from above by

$$\begin{aligned}(6.2) \quad & \int |h|^a \left| \tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n) - \pi_\infty(h) \right| dh \\ & \leq C_n^{-1} \underbrace{\int |h|^a \left| \prod_{i=1}^n \left[\frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n)} \right] \pi(\hat{\theta}_n + \frac{h}{\sqrt{n}}) - e^{-\frac{h^\top V_0 h}{2}} \pi(\theta_0) \right| dh}_{N_{n1}} \\ & + C_n^{-1} \underbrace{\int |h|^a e^{-\frac{h^\top V_0 h}{2}} \left(\pi(\theta_0) - C_n (2\pi)^{-p/2} |\det V_0|^{1/2} \right) dh}_{N_{n2}}.\end{aligned}$$

It is sufficient to show that

$$(6.3) \quad N_{n1} \equiv \int |h|^a \left| \prod_{i=1}^n \left[\frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n)} \right] \pi(\hat{\theta}_n + \frac{h}{\sqrt{n}}) - e^{-\frac{h^\top V_0 h}{2}} \pi(\theta_0) \right| dh = o_{P_0}(1).$$

This is because if (6.3) holds, taking $a = 0$ leads to

$$(6.4) \quad C_n = (2\pi)^{p/2} |\det V_0|^{-1/2} \pi(\theta_0) + o_{P_0}(1),$$

which yields $C_n = O_{P_0}(1)$. The convergence in (6.4) also implies

$$N_{n2} = |C_n (2\pi)^{-p/2} |\det V_0|^{1/2} - \pi(\theta_0)| \int_{\mathbb{R}^p} |h|^a \exp \left\{ -\frac{h^\top V_0 h}{2} \right\} dh = o_{P_0}(1),$$

which makes the upper bound in (6.2) for the total variation norm satisfy $(N_{n1} + N_{n2})/C_n = o_{P_0}(1)$. This completes the proof of 2.3.

To show (6.3), we split its integration range into three mutually exclusive areas:

- $H_{1,n} \equiv \{h : |h| \leq C\sqrt{nr_n}\}$;
- $H_{2,n} \equiv \{h : Cr_n\sqrt{n} < |h| \leq \delta\sqrt{n}\}$;
- $H_{3,n} \equiv \{h : |h| > \delta\sqrt{n}\}$;

for a large positive constant C and a small $\delta > 0$ to be specified in the sequel. Note that

$$\int_{H_{2,n} \cup H_{3,n}} \exp\left(-\frac{h^\top V_0 h}{2}\right) \pi(\theta_0) dh = o(1),$$

given the fact that the smallest eigenvalue of V_0 is positive and $r_n\sqrt{n} \rightarrow \infty$. We further define

$$(6.5) \quad \Delta_n^r \equiv \sup_{d_\Theta(\theta, \theta_0) \geq r} \left\{ \frac{1}{n} \sum_{i=1}^n [\log p_\theta(Y_{2i}; \hat{\zeta}_n) - \log p_{\hat{\theta}_n}(Y_{2i}; \hat{\zeta}_n)] \right\}.$$

First consider the integration over the outer range $H_{3,n}$. By Lemma 7.1, for any small $\epsilon > 0$, we have

$$(6.6) \quad \int_{H_{3,n}} |h|^a \prod_{i=1}^n \frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i} | \hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i} | \hat{\zeta}_n)} \pi(h) dh \leq \mathbb{I}\{\Delta_n^\delta < -\epsilon\} \sqrt{n}^{a+1} e^{-n\epsilon} \int_{\theta \in \Theta} |\theta|^a \pi(\theta) d\theta, \quad w.p.a.1.$$

The right hand side of the above inequality is of $o_{P_0}(1)$.

Then consider the integration over the middle range $H_{2,n}$. Applying Lemma 7.2, we have

$$\int_{H_{2,n}} |h|^a \prod_{i=1}^n \frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i} | \hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i} | \hat{\zeta}_n)} \pi(h) dh \leq \mathbb{I}\{\Delta_n^{Cr_n} < -cr_n^2\} \sqrt{n}^{a+1} e^{-cnr_n^2} \int_{\theta \in \Theta} |\theta|^a \pi(\theta) d\theta,$$

for any given positive c and a large enough constant C , *w.p.a.1*. The right hand side of the above inequality is again of $o_{P_0}(1)$.

Lastly, we consider the inner range $H_{1,n}$. We first switch from $\pi(\theta_0)$ to $\pi(\hat{\theta}_n + h/\sqrt{n})$ by noting that $\sup_{h \in \mathcal{H}_n} \limsup \left| \hat{\theta}_n + \frac{h}{\sqrt{n}} - \theta_0 \right| = o_{P_0}(1)$, which implies

$$\int_{H_{1,n}} |h|^a \exp\left(-\frac{h^\top V_0 h}{2}\right) \left| \pi\left(\hat{\theta}_n + \frac{h}{\sqrt{n}}\right) - \pi(\theta_0) \right| dh = o_{P_0}(1),$$

by the continuity of the prior density $\pi(\cdot)$ and the dominated convergence. The remaining analysis is about

$$\begin{aligned} & \int_{H_{1,n}} |h|^a \left| \prod_{i=1}^n \left[\frac{p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n)}{p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n)} \right] - \exp\left(-\frac{h^\top V_0 h}{2}\right) \right| \pi\left(\hat{\theta}_n + \frac{h}{\sqrt{n}}\right) dh \\ &= \int_{H_{1,n}} |h|^a e^{-\frac{1}{2}h^\top V_0 h} \left| \exp\left\{ \sum_{i=1}^n \log p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n) - \log p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n) + \frac{1}{2}h^\top V_0 h \right\} - 1 \right| \pi\left(\hat{\theta}_n + \frac{h}{\sqrt{n}}\right) dh \\ &\leq A_n \times B_n, \end{aligned}$$

where

$$\begin{aligned} A_n &\equiv \int_{\mathcal{H}_{1,n}} |h|^a \exp\left\{-\frac{h^\top V_0 h}{4}\right\} \pi\left(\hat{\theta}_n + \frac{h}{\sqrt{n}}\right) dh, \\ B_n &\equiv \sup_{h \in H_{1,n}} \left[e^{-\frac{1}{4}h^\top V_0 h} \left| \exp\left\{ \sum_{i=1}^n \log p_{\hat{\theta}_n + \frac{h}{\sqrt{n}}}(Y_{2,i}|\hat{\zeta}_n) - \log p_{\hat{\theta}_n}(Y_{2,i}|\hat{\zeta}_n) + \frac{1}{2}h^\top V_0 h \right\} - 1 \right| \right]. \end{aligned}$$

It is clear that A_n is stochastically bounded due to the boundedness of the prior density. Lemma 7.4 shows $B_n = o_{P_0}(1)$. Then we obtain $A_n \times B_n = o_{P_0}(1)$, which completes the proof of (6.3).

The convergence in the norm $TVM(a)$ stated in (2.3) immediately implies the convergence of the distribution in the total variation as stated in (2.4); see Ghosh and Ramamoorthi (2002). \square

Proof of Corollary 2.1. We denote the p -dimensional standard normal measure of a generic set A by $\mathbb{N}(A) \equiv (2\pi)^{-p/2} \int_A e^{-h^\top h/2} dh$. Recall the definition of $\Delta_{n,0}$ in Assumption 2.8. Consider the quasi-Bayesian credible set $\mathcal{C}_n(\alpha)$, which satisfies $\Pi(\theta \in \mathcal{C}_n(\alpha) | \mathbf{Y}_2^n, \hat{\zeta}_n) = 1 - \alpha$. Applying (2.3) in Theorem 2.1 with $a = 0$ and using (7.2) in Appendix B, we have

$$\mathbb{N}(V_0^{1/2} n^{1/2} (\mathcal{C}_n(\alpha) - \theta_0 - n^{-1/2} V_0^{-1} \Delta_{n,0})) \rightarrow 1 - \alpha \text{ in probability.}$$

Thus, $\mathcal{C}_n(\alpha) = \theta_0 + n^{-1/2} V_0^{-1} \Delta_{n,0} + n^{-1/2} V_0^{-1/2} B_n$, where B_n satisfies $\mathbb{N}(B_n) \rightarrow 1 - \alpha$, as $n \rightarrow \infty$. Therefore, the frequentist coverage of the Bayesian credible set is

$$P_0\{\theta_0 \in \mathcal{C}_n(\alpha)\} = P_0\{\theta_0 \in \theta_0 + n^{-1/2} V_0^{-1} \Delta_{n,0} + n^{-1/2} V_0^{-1/2} B_n\} = P_0\{-V_0^{-1/2} \Delta_{n,0} \in B_n\}.$$

The above coverage probability of the set $\mathcal{C}_n(\alpha)$ does not converge to $1 - \alpha$ in general, unless the matrix V_0 is equal to Ω_0 , which is the asymptotic covariance matrix of $\Gamma_1(Y_1) + \Gamma_2(Y_2)$ in our Assumption 2.8. \square

Proof of Corollary 2.2. The convergence in the total variation of moments norm with $a = 1$ implies that

$$\int h \left(\tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n) - \pi_\infty(h) \right) dh \rightarrow_{P_0} 0.$$

As the limiting normal distribution is centered around zero, i.e., $\int h\pi_\infty(h)dh = 0$, we have

$$\int h\tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n)dh \rightarrow_{P_0} 0.$$

Translating the above result to the quasi-posterior mean, we note that

$$\tilde{\theta}_n \equiv \int \theta d\Pi_n(\theta|\mathbf{Y}_2; \hat{\zeta}_n) = \int \left(\hat{\theta}_n + \frac{h}{\sqrt{n}} \right) \tilde{\pi}(h|\mathbf{Y}_2^n; \hat{\zeta}_n) dh = \hat{\theta}_n + o_{P_0}(n^{-1/2}),$$

which leads to the desired result. The proof of the posterior variance follows along similar lines involving the second-order moment. It is straightforward and hence omitted. \square

Proof of Theorem 2.2. In Lemma 7.5, we have shown the convergence in total variation norm for the bootstrapped posterior density function. This implies the asymptotic equivalence of the mean of the bootstrapped quasi-posterior $\tilde{\theta}_n^*$ and the bootstrap frequentist-type two-stage estimator $\hat{\theta}_n^*$, i.e., $\sqrt{n}(\tilde{\theta}_n^* - \hat{\theta}_n^*)$ is $o_{P^*}(1)$. For the two-stage frequentist estimators, we have the following:

$$(6.7) \quad \sup_{\xi \in \mathbb{R}^p} |P_0(\sqrt{n}(V_0\Omega_0V_0)^{-1/2}(\hat{\theta}_n - \theta_0) \leq \xi) - \Phi_p(\xi)| \rightarrow_{P_0} 0,$$

and its bootstrap analog:

$$(6.8) \quad \sup_{\xi \in \mathbb{R}^p} |P^*(\sqrt{n}(V_0\Omega_0V_0)^{-1/2}(\hat{\theta}_n^* - \hat{\theta}_n) \leq \xi) - \Phi_p(\xi)| \rightarrow_{P^*} 0.$$

Using (6.7), (6.8), and the asymptotic equivalence between $\tilde{\theta}_n^*$ and $\hat{\theta}_n^*$, we conclude the proof by applying Corollary 1 in Cheng and Huang (2010). \square

Proof of Proposition 3.1. We verify the high-level assumptions that lead to our Theorems 2.1 and 2.2. Regarding Assumption 2.1, we have a well-separated maximum point, due to the identification and the compactness of the parameter space in Assumption 3.1, as well as the continuity of the log-likelihood function (Newey and McFadden, 1994). Assumption 3.3 also satisfies the convergence rate condition in Assumption 2.2. Referring to the frequentist's estimator $\hat{\theta}_n$, one can take the simulated score estimator by Hajivassiliou and McFadden (1998) in the second stage. The prior specification given in Assumption 3.4 satisfies the requirement in Assumption 2.4. The normality of the error term generates sufficient smoothness of the conditional choice probabilities, which satisfy Assumption 2.6. Therein, the function $e_n(\cdot, \cdot)$ in Assumption 2.5 can be taken as the cross product term $(\theta - \theta_0)^\top \mathbb{E}[\ddot{l}_{\theta, \zeta}][\zeta - \zeta_0]$ as in Lemma C.2 of Chen, Lee, and Sung (2014). The smoothness

of MNP implies a stronger notion of differentiability than what is required in Assumption 2.6, i.e., the likelihood is Frechét differentiable w.r.t. ζ (Ichimura and Lee, 2010).

To verify Assumption 2.7, recall that the bracketing number $N_{[]}(\epsilon, \mathcal{F}, \|\cdot\|_2)$ for a functional class \mathcal{F} is defined to be the minimum of m such that $\exists f_1^L, f_1^U, \dots, f_m^L, f_m^U$ for $\forall f \in \mathcal{F}$, $f_j^L \leq f \leq f_j^U$ for some j , and $\|f_j^U - f_j^L\|_2 \leq \epsilon$ (van der Vaart and Wellner, 1996). Let $\mathcal{P} \equiv \{\log p_\theta(Y_{2i}; \zeta) : \theta \in \Theta, \zeta \in \mathcal{G}_n\}$. Our Assumption 3.2 on the Hölder class implies the P_0 -Glivenko-Cantelli property of \mathcal{P} . Given the smoothness requirement in Assumption 3.2, we can utilize the Lipschitz continuity and apply Theorem 2.7.11 and Theorem 2.7.1 in van der Vaart and Wellner (1996) to bound its overall entropy by

(6.9)

$$\log N_{[]}(\epsilon, \mathcal{P}, \|\cdot\|_2) \lesssim \log N_{[]}(\epsilon/2C, \Theta, \|\cdot\|_2) + \log N_{[]}(\rho/2C, \mathcal{G}_n, \|\cdot\|_2) \lesssim p \log(\epsilon^{-1}) + \epsilon^{-d/\tau}.$$

To check the P_0 -Donsker property as in Assumption 2.7(ii), we follow the route in Example 19.7 from van der Vaart (1998) given our Assumption 3.5, along with the restriction that $\tau > d/2$ in Assumption 3.2. Other properties, such as the functional class of the second derivative \ddot{l}_θ being P_0 -Glivenko-Cantelli, can be checked along similar lines, utilizing the smoothness of the MNP in the second stage, as well as the entropy bound for the Hölder class from the first stage in (6.9); see McFadden (1989) and Pakes and Pollard (1989). The score functions of β and η (equations (14) and (15) in Hajivassiliou and McFadden (1998)) have finite second-order moments given the bounded support of covariates and control variables. Given the smoothness and boundedness of covariates support, we also have the envelope functions being bounded. Thus, the slightly strong assumption used in the bootstrap part is also satisfied. When it comes to Assumption 2.7(iii), one can apply Lemma 2.14.3 in van der Vaart and Wellner (1996) to obtain $\phi_n(\rho) = \rho^{1-\frac{d}{2\tau}}$, which satisfy the restrictions under the maintained assumption $\tau > d/2$. By equation (3.15) of Ichimura and Lee (2010), the influence function of the first-stage estimation defined by $\Gamma_1(Y_{1i}) = \sum_{j=0}^J \int \phi_{n,j}(\cdot, \theta_0) \frac{\partial g_j(\cdot, \theta_0)}{\partial \theta} dP_0$. The asymptotic normality in Assumption 2.8 follows from the Lindeberg-Lévy CLT under Assumption 3.6. \square

7. Appendix B: Proofs of Technical Lemmas

We prove several technical lemmas needed in the proof of Theorems 2.1 and 2.2 in this part. For simplicity, we state the lemmas under all of our maintained Assumptions 2.1 to 2.8. From the proofs, it is clear which specific conditions are actually invoked.

Lemma 7.1. *Under Assumptions 2.1 to 2.8, for any $\delta > 0$, we have the following inequality holds w.p.a.1,*

$$(7.1) \quad \sup_{d_{\Theta}(\theta, \theta_0) > \delta} \mathbb{P}_n[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n) - \log p_{\widehat{\theta}_n}(Y_{2i}; \widehat{\zeta}_n)] \leq -c,$$

for some positive constant term c .

Proof. We start with the following decomposition:

$$\begin{aligned} & \mathbb{P}_n[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n) - \log p_{\widehat{\theta}_n}(Y_{2i}; \zeta_0)] \\ &= \mathbb{P}_n[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_{2i}; \zeta_0)] + \mathbb{P}_n[\log p_{\theta_0}(Y_{2i}; \zeta_0) - \log p_{\widehat{\theta}_n}(Y_{2i}; \widehat{\zeta}_n)]. \end{aligned}$$

The second term above can be shown to be $o_{P_0}(1)$ following the standard consistency argument for the two-stage estimation given the identification in Assumption 2.1 and P_0 -Glivenko-Cantelli property in Assumption 2.7; see Chen (2007). Therefore, we focus on the first term and decompose it as follows:

$$\begin{aligned} & \mathbb{P}_n[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_{2i}; \zeta_0)] \\ &= (\mathbb{P}_n - P_0)[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n)] - (\mathbb{P}_n - P_0)[\log p_{\theta_0}(Y_{2i}; \zeta_0)] \\ & \quad + P_0[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n) - \log p_{\theta}(Y_{2i}; \zeta_0)] + P_0[\log p_{\theta}(Y_{2i}; \zeta_0) - \log p_{\theta_0}(Y_{2i}; \zeta_0)]. \end{aligned}$$

By the P_0 -Glivenko-Cantelli property, i.e., $\sup_{\theta \in \Theta, \zeta \in \mathcal{G}_n} |(\mathbb{P}_n - P_0)[\log p_{\theta}(Y_{2i}; \zeta)]| = o_{P_0}(1)$, we have the first two terms on the right hand side of the above equality converging to zero in probability uniformly for any $\theta \in \Theta$. When it comes to the third term, we utilize the uniform (regarding θ) continuity of the criterion function with respect to the first stage nuisance function in our Assumption 2.6, so that $P_0[\log p_{\theta}(Y_{2i}; \widehat{\zeta}_n) - \log p_{\theta}(Y_{2i}; \zeta_0)] = o_{P_0}(1)$. For any given $\delta > 0$, the last term satisfies

$$\sup_{d_{\Theta}(\theta, \theta_0) > \delta} |P_0[\log p_{\theta}(\cdot; \zeta_0) - \log p_{\theta_0}(\cdot; \zeta_0)]| \leq -c,$$

with a proper choice of constant $c > 0$, which concludes the proof. \square

Our assumptions imply Conditions (i)-(iv) in Corollary 1 of Nan and Wellner (2013). Thus, the frequentist's two-stage estimator satisfies the following expansion:

$$(7.2) \quad \begin{aligned} (\widehat{\theta}_n - \theta_0) &= -V_0^{-1} \left(S_{\theta, n}(\theta_0, \zeta_0) + \dot{\Psi}_{\zeta}(\theta_0, \zeta_0)[\widehat{\zeta}_n - \zeta_0] \right) + o_{P_0}(n^{-1/2}) \\ &\equiv -V_0^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n [\Gamma_2(Y_{2i}) + \Gamma_1(Y_{1i})] + o_{P_0}(n^{-1/2}), \end{aligned}$$

which implies that $\widehat{\theta}_n$ is root- n consistent and asymptotically normal.

Lemma 7.2. Recall the definition of Δ_n^r in (6.5). Under Assumptions 2.1 to 2.8, for the decreasing sequence $r_n = o(n^{-1/4})$ and $nr_n^2 \rightarrow \infty$, we have

$$(7.3) \quad \limsup_{n \rightarrow \infty} P_0\{\Delta_n^{Cr_n} \geq -cr_n^2\} = 0,$$

for a given positive c and a large enough positive constant C .

Proof of Lemma 7.2. Because $\hat{\theta}_n$ is root- n consistent and $nr_n^2 \rightarrow \infty$, we have

$$\{\theta : d_{\Theta}(\theta, \hat{\theta}_n) \geq Cr_n\} \subset \{\theta : d_{\Theta}(\theta, \theta_0) \geq C'r_n\},$$

for a proper choice of the constant term C' . Therefore, we obtain

$$\begin{aligned} \{\Delta_n^{Cr_n} \geq -cr_n^2\} &\subset \left\{ \sup_{\theta: d_{\Theta}(\theta, \theta_0) \geq C'r_n} \mathbb{P}_n[\log p_{\theta}(Y_2; \hat{\zeta}_n) - \log p_{\hat{\theta}_n}(Y_2; \hat{\zeta}_n)] \geq -cr_n^2 \right\} \\ &\subset \left\{ \sup_{\theta: d_{\Theta}(\theta, \theta_0) \geq C'r_n} \mathbb{P}_n[\log p_{\theta}(Y_2; \hat{\zeta}_n) - \log p_{\theta_0}(Y_2; \hat{\zeta}_n)] \geq -cr_n^2 \right\}, \end{aligned}$$

where the second inclusion follows from $\mathbb{P}_n[\log p_{\hat{\theta}_n}(Y_2; \hat{\zeta}_n)] \geq \mathbb{P}_n[\log p_{\theta_0}(Y_2; \hat{\zeta}_n)]$ by definition of $\hat{\theta}_n$ from Assumption 2.3. We can further restrict to our attention to the set $B_n \equiv \{d_G(\zeta, \zeta_0) \leq Cr_n\}$, as its complement is asymptotically negligible by Assumption 2.2. It suffices to examine

$$(7.4) \quad \left\{ \sup_{d_{\Theta}(\theta, \theta_0) \geq C'r_n, d_G(\zeta, \zeta_0) \leq Cr_n} \mathbb{P}_n[\log p_{\theta}(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)] \geq -cr_n^2 \right\}$$

Thanks to our Lemma 7.1, we can localize on the set $\{\theta : d_{\Theta}(\theta, \theta_0) \leq \delta\}$ for a given small δ . This localized set will be partitioned into the shells for $\rho = C'r_n$:

$$S_{n,j,M} = \{(\theta, \zeta) \in \Theta_n \times \mathcal{G}_n : 2^{j-1}\rho < d_{\Theta}(\theta, \theta_0) \leq 2^j\rho, d_G(\zeta, \zeta_0) \leq 2^{-M}d_{\Theta}(\theta, \theta_0)\}.$$

In the slicing argument, we sum over the shells that $2^j\rho < \delta$. For the j -th set involved, we have

$$\begin{aligned} &P_0(\log p_{\theta}(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)) \\ &= P_0(\log p_{\theta}(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta_0)) + P_0(\log p_{\theta_0}(Y_2; \zeta_0) - \log p_{\theta_0}(Y_2; \zeta)) \\ &\leq P_0(\log p_{\theta}(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta_0)) + e_n(\theta, \zeta) + C_1r_n^2 \\ (7.5) \quad &\lesssim -d_{\Theta}^2(\theta, \theta_0) + d_G^2(\zeta, \zeta_0) + C_22^j\rho r_n + C_1r_n^2 \lesssim -(1 - 2^{-2M})d_{\Theta}^2(\theta, \theta_0) \lesssim -2^{2j-2}\rho^2, \end{aligned}$$

under Assumption 2.5. In other words, we have

$$\inf_{(\theta, \zeta) \in S_{n,j,M}} -P_0(\log p_{\theta}(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)) \geq 2^{2j-2}\rho^2.$$

Thus, the probability of (7.4) can be bounded by

$$\begin{aligned}
& \sum_{j \geq M, 2^j \rho_n \leq \delta} P_0 \left(\sup_{(\theta, \zeta) \in S_{n,j,M}} \mathbb{P}_n[\log p_\theta(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)] \geq -c_1 r_n^2 \right) \\
& \leq \sum_{j \geq M} P_0 \left(\sup_{(\theta, \zeta) \in S_{n,j,M}} \mathbb{G}_n[\log p_\theta(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)] \geq C 2^{2j-2} \sqrt{n} \rho^2 \right) \\
& \lesssim \sum_{j \geq M} \frac{1}{2^{2j} \sqrt{n} \rho^2} \mathbb{E}_0 \sup_{(\theta, \zeta) \in S_{n,j,M}} |\mathbb{G}_n[\log p_\theta(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)]| \\
& \lesssim \sum_{j \geq M} \frac{1}{2^{2j} \sqrt{n} \rho^2} \phi_n(2^j \rho) \leq \sum_{j \geq M} \frac{1}{2^{2j} \sqrt{n} \rho^2} 2^{j\gamma} \phi_n(\rho) \lesssim \sum_{j \geq M} 2^{j(\gamma-2)}.
\end{aligned}$$

In the first inequality, we add/subtract $P_0[\log p_\theta(\cdot; \zeta) - \log p_{\theta_0}(\cdot; \zeta)]$ and apply (7.5). The fourth inequality makes use of the monotonicity of the mapping $\delta \mapsto \phi_n(\delta)/\delta^\gamma$ for some $\gamma < 2$, so that we have $\phi_n(2^j \rho) \leq 2^{j\gamma} \phi_n(\rho)$. The final inequality follows from $\phi_n(\rho) \leq \sqrt{n} \rho^2$ for every n . The desired result follows by choosing a large enough M , which is induced by a large C . \square

Lemma 7.3. *Under Assumptions 2.1 to 2.8, for any θ_n converging to θ_0 such that $d_\Theta(\theta_n, \theta_0) \leq Cr_n$ w.p.a.1, one has*

$$(7.6) \quad \ell_n(\theta_n; \widehat{\zeta}_n) - \ell_n(\widehat{\theta}_n; \widehat{\zeta}_n) = -\frac{1}{2}(\theta_n - \widehat{\theta}_n)^\top V_0(\theta_n - \widehat{\theta}_n) + o_{P_0}(1 + \sqrt{nd_\Theta}(\theta_n, \theta_0))^2.$$

Proof. Recall that $\Delta_{n,0} \equiv \sqrt{n} \left(S_{\theta,n}(\theta_0, \zeta_0) + \dot{\Psi}_\zeta(\theta_0, \zeta_0)[\widehat{\zeta}_n - \zeta_0] \right)$. We start with the following string of equations:

$$\begin{aligned}
& \ell_n(\theta_n; \widehat{\zeta}_n) - \ell_n(\theta_0; \widehat{\zeta}_n) \\
& = n \mathbb{P}_n[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n)] \\
& = \sqrt{n} \mathbb{G}_n[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n)] + n P_0[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n)] \\
& = \sqrt{n} \mathbb{G}_n \left[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n) - (\theta_n - \theta_0)^\top \dot{l}_\theta(Y_2, \theta_0; \widehat{\zeta}_n) \right] \\
& \quad + \sqrt{n} h_n^\top \mathbb{G}_n \left[\dot{l}_\theta(Y_2, \theta_0; \widehat{\zeta}_n) - \dot{l}_\theta(Y_2, \theta_0; \zeta_0) \right] \\
& \quad + \sqrt{n} h_n^\top \mathbb{G}_n \left[\dot{l}_\theta(Y_2, \theta_0; \zeta_0) \right] + n P_0 \left[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n) \right]
\end{aligned}$$

By the P_0 -Donsker property in our Assumption 2.7, we have

$$(7.7) \quad \sup_{d_\Theta(\theta_n, \theta_0) \leq Cr_n} \left| \mathbb{G}_n \left[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n) - (\theta_n - \theta_0)^\top \dot{l}_\theta(Y_2, \theta_0; \widehat{\zeta}_n) \right] \right| = o_{P_0}(d_\Theta(\theta_n, \theta_0)),$$

and

$$(7.8) \quad \mathbb{G}_n \left[\dot{l}_\theta(Y_2, \theta_0; \widehat{\zeta}_n) - \dot{l}_\theta(Y_2, \theta_0; \zeta_0) \right] = o_{P_0}(n^{-1/2}).$$

In addition, the smoothness in Assumption 2.6 leads to

$$\begin{aligned} & \sup_{d_\Theta(\theta_n, \theta_0) \leq Cr_n} \left| nP_0 \left[\log p_{\theta_n}(Y_2; \widehat{\zeta}_n) - \log p_{\theta_0}(Y_2; \widehat{\zeta}_n) \right] \right. \\ & \quad \left. - \frac{1}{2}(\theta_n - \theta_0)^\top V_0(\theta_n - \theta_0) - (\theta_n - \theta_0)^\top \dot{\Psi}_\zeta(\theta_0, \zeta_0)[\widehat{\zeta}_n - \zeta_0] \right| \\ & = o_{P_0}(1 + \sqrt{nd}_\Theta(\theta_n, \theta_0))^2, \end{aligned}$$

where $r_n = o(n^{-1/4})$. Combining the previous steps, we have

$$\begin{aligned} \ell_n(\theta_n; \widehat{\zeta}_n) - \ell_n(\theta_0; \widehat{\zeta}_n) &= \sqrt{n}(\theta_n - \theta_0)^\top \Delta_{n,0} - \frac{n}{2}(\theta_n - \theta_0)^\top V_0(\theta_n - \theta_0) \\ & \quad + o_{P_0}(1 + \sqrt{nd}_\Theta(\theta_n, \theta_0))^2 \\ \ell_n(\widehat{\theta}_n; \widehat{\zeta}_n) - \ell_n(\theta_0; \widehat{\zeta}_n) &= \frac{n}{2}(\widehat{\theta}_n - \theta_0)^\top V_0(\widehat{\theta}_n - \theta_0) + o_{P_0}(1 + \sqrt{nd}_\Theta(\widehat{\theta}_n, \theta_0))^2. \end{aligned}$$

Now we take the difference of the above two equations and utilize the linear representation of $\widehat{\theta}_n$ to get

$$\begin{aligned} \ell_n(\theta_n; \widehat{\zeta}_n) - \ell_n(\widehat{\theta}_n; \widehat{\zeta}_n) &= n(\theta_n - \theta_0)^\top V_0(\widehat{\theta}_n - \theta_0) - \frac{n}{2}(\theta_n - \theta_0)^\top V_0(\theta_n - \theta_0) \\ & \quad - \frac{n}{2}(\widehat{\theta}_n - \theta_0)^\top V_0(\widehat{\theta}_n - \theta_0) + o_{P_0}(1 + \sqrt{nd}_\Theta(\widehat{\theta}_n, \theta_0))^2 \\ & = -\frac{1}{2}(\theta_n - \widehat{\theta}_n)^\top V_0(\theta_n - \widehat{\theta}_n) + o_{P_0}(1 + \sqrt{nd}_\Theta(\theta_n, \theta_0))^2. \end{aligned}$$

The last step follows from completing the square of $(\theta_n - \widehat{\theta}_n)^\top V_0(\theta_n - \widehat{\theta}_n)$. \square

Lemma 7.4. *Let*

$$\tilde{B}_n(\theta) \equiv e^{-\frac{n}{4}(\theta - \widehat{\theta}_n)^\top V_0(\theta - \widehat{\theta}_n)} \left| e^{\sum_{i=1}^n \log p_{\theta}(Y_{2,i}|\widehat{\zeta}_n) - \log p_{\widehat{\theta}_n}(Y_{2,i}|\widehat{\zeta}_n) + \frac{n}{2}(\theta - \widehat{\theta}_n)^\top V_0(\theta - \widehat{\theta}_n)} - 1 \right|.$$

Under Assumptions 2.1 to 2.8, we have $B_n \equiv \sup_{d(\theta, \theta_0) \leq Cr_n} \tilde{B}_n(\theta) = o_{P_0}(1)$.

Proof. For any random sequence θ_n such that $d_\Theta(\theta_n, \widehat{\theta}_n) \leq Cr_n$ for all n and some fixed $C < \infty$, we have

$$\tilde{B}_n(\theta_n) \leq \exp \left\{ -\frac{n}{4}(\theta_n - \widehat{\theta}_n)^\top V_0(\theta_n - \widehat{\theta}_n) \right\} \times \left| \exp \{ o_{P_0}(1 + \sqrt{nd}_\Theta(\theta_n, \theta_0))^2 \} - 1 \right|,$$

where the remainder term satisfies

$$\begin{aligned} o_{P_0}(1 + \sqrt{nd}_\Theta(\theta_n, \theta_0))^2 &\leq o_{P_0}(1)[(1 + \sqrt{nd}_\Theta(\theta_n, \widehat{\theta}_n))^2 + nd(\theta_0, \widehat{\theta}_n)^2] \\ &= o_{P_0}(1)(1 + \sqrt{nd}_\Theta(\theta_n, \widehat{\theta}_n))^2, \end{aligned}$$

due to Lemma 7.3. Note that $|e^x - 1| \leq e^{|x|}$ for any x , we can proceed as follows for a given large positive constant M :

$$\begin{aligned} \tilde{B}_n(\theta_n) &\leq \exp \left\{ -\frac{n}{4}(\theta_n - \hat{\theta}_n)^\top V_0(\theta_n - \hat{\theta}_n) \right\} |\exp\{o_{P_0}(1)\} - 1| \times \mathbb{I}\{\sqrt{n}d_\Theta(\theta_n, \hat{\theta}_n) \leq M\} \\ &\quad + \exp \left\{ -\frac{n}{4}(\theta_n - \hat{\theta}_n)^\top V_0(\theta_n - \hat{\theta}_n) + o_{P_0}(1)(1 + \sqrt{n}d(\theta_n, \theta_0))^2 \right\} \times \mathbb{I}\{\sqrt{n}d_\Theta(\theta_n, \hat{\theta}_n) \geq M\} \\ &\leq o_{P_0}(1) + \exp \left\{ \frac{M^2}{4} \lambda_1(1 - o_{P_0}(1)) \right\} \rightarrow \exp\{-\lambda_1 M^2/4\}, \end{aligned}$$

where λ_1 is the smallest eigenvalue of V_0 which is strictly positive by our Assumption 2.1. Therefore, $B_n = o_{P_0}(1)$, as M can be made arbitrarily large. \square

Lemma 7.5. *Let the bootstrap posterior density function of $h \equiv \sqrt{n}(\theta - \hat{\theta}_n^*)$ be*

$$(7.9) \quad \tilde{\pi}^*(h | \mathbf{Y}_2^{n*}; \hat{\zeta}_n^*) \equiv \frac{\prod_{i=1}^n p_{\hat{\theta}_n^* + \frac{h}{\sqrt{n}}}^{M_{ni}}(\mathbf{Y}_2^n | \hat{\zeta}_n^*) \pi(\hat{\theta}_n^* + \frac{h}{\sqrt{n}})}{\int \prod_{i=1}^n p_{\hat{\theta}_n^* + \frac{h}{\sqrt{n}}}^{M_{ni}}(\mathbf{Y}_2^n | \hat{\zeta}_n^*) \pi(\hat{\theta}_n^* + \frac{h}{\sqrt{n}}) dh}.$$

In addition to Assumptions 2.1-2.8, we assume that the envelope functions for the functional classes in Parts (i) and (ii) in Assumption (2.7) have finite $2+\iota$ moments, for $\iota > 0$. Then, we have the following convergence in total variation norm:

$$(7.10) \quad \|\tilde{\pi}^*(h | \mathbf{Y}_2^{n*}; \hat{\zeta}_n^*) - \pi_\infty(h)\|_{TVM(a)} = o_{P^*}(1).$$

Proof. The bootstrap version follows along similar lines of the proof of our Theorem 2.1. We highlight two main modifications needed herein.

First, we show the posterior mass outside an shrinking neighborhood around the truth is negligible conditional on the observed data w.p.a.1. Specifically, we need to characterize the order of $\Delta_n^{*Cr_n}$, where

$$\Delta_n^{*Cr_n} = \sup_{\theta: d_\Theta(\theta, \hat{\theta}_n^*) \geq Cr_n} \mathbb{P}_n^*[\log p_\theta(Y_2; \hat{\zeta}_n^*) - \log p_{\hat{\theta}_n^*}(Y_2; \hat{\zeta}_n^*)].$$

A key step in the chaining argument is to use a proper maximal inequality related to the bootstrap empirical process. For this purpose, we resort to Corollary 1 from Han and Wellner (2019) to obtain

$$(7.11) \quad \mathbb{E}_0^* \sup_{(\theta, \zeta) \in S_{n,j,M}} |\mathbb{G}_n^*[\log p_\theta(Y_2; \zeta) - \log p_{\theta_0}(Y_2; \zeta)]| \leq C\phi_n(2^j\rho),$$

with the multinomial weights.

Second, when we carry out an expansion of the bootstrap log-likelihood, we need to verify the following stochastic equicontinuity as

$$(7.12) \quad \mathbb{G}_n^* \left[\dot{l}_\theta(Y_2, \theta_0; \hat{\zeta}_n^*) - \dot{l}_\theta(Y_2, \theta_0; \zeta_0) \right] = o_{P^*}(n^{-1/2}).$$

Denote $\mathcal{F} \equiv \{f = \dot{l}_\theta(\cdot, \theta_0; \zeta) - \dot{l}_\theta(\cdot, \theta_0; \zeta_0), \zeta \in \mathcal{G}_n\}$. We apply the multiplier inequality in Lemma 3.6.7 of van der Vaart and Wellner (1996), which states for any $1 \leq n_0 \leq n$

$$(7.13) \quad \mathbb{E}_{ZM} \|\mathbb{G}_n^* f\| \lesssim 2(n_0 - 1) \mathbb{E}_Z |F| \frac{1}{\sqrt{n}} \mathbb{E}_M \left| \max_i M_{ni} \right| + 2 \mathbb{E}_M M_{n1}^2 \left| \max_{n_0 \leq k \leq n} \mathbb{E}_Z \|\mathbb{G}_k f\| \right|,$$

where F is the corresponding envelope function for the functional class \mathcal{F} , and the subscript on the expectation signifies the source of randomness.

For the multinomial weights, we have $\mathbb{E}|M_{ni}|^2 < \infty$. We start with applying the Jensen's inequality to obtain

$$\left(\mathbb{E}_M \left| \max_{1 \leq i \leq n} \frac{M_{ni}}{\sqrt{n}} \right| \right)^2 \leq \mathbb{E}_M \left| \max_{1 \leq i \leq n} \frac{M_{ni}^2}{n} \right| \leq \frac{m^2}{n} + \frac{1}{n} \sum_{i=1}^n \mathbb{E}_M [M_{ni}^2 \mathbb{I}\{M_{ni} > m\}],$$

which goes to 0. Thus, the first term in (7.13) is of smaller order as $\frac{1}{\sqrt{n}} \mathbb{E}_M |\max_i M_{ni}| = o(1)$. Regarding the term $\max_{n_0 \leq k \leq n} \mathbb{E}_Z \|\mathbb{G}_k f\|$, it suffices to bound $\mathbb{E}_Z \|\mathbb{G}_n f\|$ by Lemma 4.3 of Praestgaard and Wellner (1993). Because of the stochastic equicontinuity of the empirical process and the Hoffmann-Jorgensen inequality, the first-order moment version also holds, i.e., $\mathbb{E}_Z \|\mathbb{G}_n f\| = o_{P_0}(1)$. In sum, this makes the second term of (7.13) negligible, given the original empirical process satisfies the P_0 -Donsker property. \square

References

- AGARWAL, N., AND P. SOMAINI (2018): "Demand analysis using strategic reports: An application to a school choice mechanism," *Econometrica*, 86, 391–444.
- ALBERT, J. H., AND S. CHIB (1993): "Bayesian analysis of binary and polychotomous response data," *Journal of the American statistical Association*, 88(422), 669–679.
- ANCESCHI, N., A. FASANO, D. DURANTE, AND G. ZANELLA (2023): "Bayesian conjugacy in Probit, Tobit, Multinomial Probit and extensions: A review and new results," *Journal of the American Statistical Association*, 118(542), 1451–1469.
- CHEN, L. Y., S. LEE, AND M. J. SUNG (2014): "Maximum score estimation with non-parametrically generated regressors," *Econometrics Journal*, 17, 271–300.
- CHEN, X. (2007): "Large sample sieve estimation of semi-nonparametric models," *Handbook of Econometrics*, 6, 5549–5632.
- CHEN, X., T. M. CHRISTENSEN, AND E. TAMER (2018): "Monte Carlo confidence sets for identified sets," *Econometrica*, 86, 1965–2018.
- CHEN, X., O. LINTON, AND I. VAN KEILEGOM (2003): "Estimation of semiparametric models when the criterion function is not smooth," *Econometrica*, 71, 1591–1608.
- CHENG, G., AND J. HUANG (2010): "Bootstrap consistency for general semiparametric M-estimate," *Annals of Statistics*, 38, 2884–2915.

- CHENG, G., AND M. R. KOSOROK (2008): “General frequentist properties of the posterior profile distribution,” *Annals of Statistics*, 36, 1819–1853.
- CHERNOZHUKOV, V., AND H. HONG (2003): “An MCMC approach to classical estimation,” *Journal of Econometrics*, 115, 293–346.
- CHIB, S., M. SHIN, AND A. SIMONI (2018): “Bayesian estimation and comparison of moment condition models,” *Journal of the American Statistical Association*, 113, 1656–1668.
- EFRON, B. (1979): “Bootstrap methods: Another look at the jackknife,” *The Annals of statistics*, 7, 1–26.
- ELDAR, O., AND L. MAGNOLFI (2020): “Regulatory competition and the market for corporate law,” *American Economic Journal: Microeconomics*, 12(2), 60–98.
- GHOSAL, S., AND A. VAN DER VAART (2017): *Fundamentals of nonparametric Bayesian inference*. Cambridge University Press.
- GHOSH, J., AND R. RAMAMOORTHI (2002): *Bayesian Nonparametrics*. Springer.
- HAJIVASSILIOU, V., D. MCFADDEN, AND P. RUUD (1996): “Simulation of multivariate normal rectangle probabilities and their derivatives theoretical and computational results,” *Journal of econometrics*, 72(1-2), 85–134.
- HAJIVASSILIOU, V. A., AND D. MCFADDEN (1998): “The method of simulated scores for the estimation of LDV models,” *Econometrica*, 66, 863–896.
- HAN, Q., AND J. WELLNER (2019): “Convergence rates of least squares regression estimators with heavy-tailed errors,” *The Annals of Statistics*, 47, 2286–2319.
- HECKMAN, J. J., AND R. ROBB (1985): “Alternative methods for evaluating the impact of interventions: An overview,” *Journal of Econometrics*, pp. 239–267.
- ICHIMURA, H., AND S. LEE (2010): “Characterization of the asymptotic distribution of semiparametric M-estimators,” *Journal of Econometrics*, 159, 252–266.
- IMAI, K., AND D. A. VAN DYK (2005): “A Bayesian analysis of the multinomial probit model using marginal data augmentation,” *Journal of econometrics*, 124(2), 311–334.
- KIM, J.-Y. (2002): “Limited information likelihood and Bayesian analysis,” *Journal of Econometrics*, 107(1-2), 175–193.
- (2014): “An alternative quasi likelihood approach, Bayesian analysis and data-based inference for model specification,” *Journal of Econometrics*, 178, 132–145.
- KLEIJN, B. J. K., AND A. W. VAN DER VAART (2012): “The Bernstein-von-Mises theorem under misspecification,” *Electronic Journal of Statistics*, 6, 354–381.
- LEE, B., M. KOSOROK, AND J. FINE (2005): “The profile sampler,” *Journal of the American Statistical Association*, 100, 960–969.
- MCCULLOCH, R., N. POLSON, AND P. E. ROSSI (2000): “A Bayesian analysis of the multinomial probit model with fully identified parameters,” *Journal of Econometrics*,

- 99(1), 173–193.
- MCCULLOCH, R., AND P. E. ROSSI (1994): “An exact likelihood analysis of the multinomial probit model,” *Journal of Econometrics*, 64(1-2), 207–240.
- McFADDEN, D. (1989): “A method of simulated moments for estimation of discrete response models without numerical integration,” *Econometrica*, 57, 995–1026.
- MÜLLER, U. K. (2013): “Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix,” *Econometrica*, 81(5), 1805–1849.
- MURPHY, K. M., AND R. H. TOPEL (2002): “Estimation and inference in two-step econometric models,” *Journal of Business & Economic Statistics*, 20, 88–97.
- NAN, B., AND J. A. WELLNER (2013): “A general semiparametric Z-estimation approach for case-cohort studies,” *Statistica Sinica*, 23, 1155.
- NEWKEY, W. K. (1994): “The asymptotic variance of semiparametric estimators,” *Econometrica*, 62, 1349–1382.
- NEWKEY, W. K., AND D. McFADDEN (1994): “Large sample estimation and hypothesis testing,” *Handbook of Econometrics*, 4, 2111–2245.
- NOBILE, A. (1998): “A hybrid Markov chain for the Bayesian analysis of the multinomial probit model,” *Statistics and Computing*, 8(3), 229–242.
- PAKES, A., AND D. POLLARD (1989): “Simulation and the asymptotics of optimization estimators,” *Econometrica*, 57, 1027–1057.
- PETRIN, A., AND K. TRAIN (2010): “A control function approach to endogeneity in consumer choice models,” *Journal of marketing research*, 47(1), 3–13.
- PRAESTGAARD, J., AND J. WELLNER (1993): “Exchangably weighted bootstraps of the general empirical process,” *Annals of Probability*, 21, 2053–2086.
- TRAIN, K. E. (2009): *Discrete choice methods with simulation*. Cambridge university press.
- VAN DER VAART, A. (1998): *Asymptotic statistics*. Cambridge University Press.
- VAN DER VAART, A., AND J. A. WELLNER (1996): *Weak convergence and empirical processes*. Springer.
- WOOLDRIDGE, J. M. (2015): “Control function methods in applied econometrics,” *Journal of Human Resources*, 50(2), 420–445.