Bayesian and Frequentist Inference in Partially Identified Models

Hyungsik Roger Moon
University of Southern California

Frank Schorfheide*
University of Pennsylvania
CEPR and NBER

May 20, 2010

*Correspondence: H.R. Moon: Department of Economics, KAP 300, University Park Campus University of Southern California, Los Angeles, CA 90089. Email: moonr@usc.edu. F. Schorfheide: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104. Email: schorf@ssc.upenn.edu. We thank seminar participants at the 2007 NASM, the 2008 SETA, Boston College, Johns Hopkins, Maryland, Ohio State, Rice, UC Davis, UC Irvine, UCLA, and Vanderbilt as well as three anonymous referees for helpful comments. We also thank Kyoungeun Kim and Martin Weidern for proofreading a draft of this paper. Moon gratefully acknowledges financial support from the National Science Foundation under Grant SES 0920903 and the USC Faculty Development Award. Schorfheide gratefully acknowledges financial support from the Alfred P. Sloan Foundation and the National Science Foundation under Grant SES 0617803. Gauss Programs that implement the numerical illustrations are available at *http://www.ssc.upenn.edu/~schorf*.

Abstract

A large sample approximation of the posterior distribution of partially identified structural parameters is derived for models that can be indexed by a finite-dimensional reduced form parameter vector. It is used to analyze the differences between Bayesian credible sets and frequentist confidence sets in partially identified models. We show that asymptotically Bayesian highest-posterior-density sets tend to be located inside the identified set, whereas it is well known that frequentist confidence sets extend beyond the boundaries of the estimated identified set. We strongly recommend that estimates of the identified set be reported along with Bayesian credible sets. A numerical illustration for a two-player entry game is provided.

JEL CLASSIFICATION: C11, C32, C33, C35

KEY WORDS: Bayesian Inference, Frequentist Confidence Sets, Partially Identified Models, Entry Games.

1 Introduction

In partially identified models one can only bound, but not point-identify the structural parameter vector of interest, θ. Such models arise in microeconometric as well as macroeconometric applications. Examples in the economics literature include models of censored sampling, models with missing observations or interval data (Manski and Tamer, 2002, and Manski, 2003), game-theoretic models with multiple equilibria such as entry games (Bresnahan and Reiss, 1991, Berry, 1994, Bajari, Benkard, and Levin, 2007, and Ciliberto and Tamer, 2009) and auctions (Haile and Tamer, 2003), vector autoregressions (VARs) that are identified by restricting only the sign of dynamic responses to structural shocks (Canova and De Nicolo, 2002, and Uhlig, 2005), and dynamic stochastic general equilibrium models (Lubik and Schorfheide, 2004). Given the lack of point identification researchers have focused on set estimators for the parameter of interest. While the macroeconometrics literature mostly applies Bayesian approaches, the microeconometric literature is dominated by frequentist procedures.

In regular identified models set estimates reported in the literature are often interpretable from both a Bayesian and a frequentist perspective. For instance, if Y_i, $i = 1, \ldots, n$ is a sequence of independent normal random variables with mean θ and unit variance, then the interval $[\hat{\theta}_n - 1.96/\sqrt{n},\ \hat{\theta}_n + 1.96/\sqrt{n}]$, where $\hat{\theta}_n = \sum Y_i/n$ is both a valid 95% frequentist confidence set as well as a 95% credible set under a flat prior for θ. While the finite-sample equivalence breaks down in more complicated models, it often holds approximately as the sample size tends to infinity. This (approximate) equivalence is very convenient for the dissemination of applied research. A Bayesian reader of a frequentist paper will find the numerical results useful despite disagreements about inference methods and so will a frequentist reader of a Bayesian paper. We will show that the equivalence breaks down in the context of partially identified models.[1] Thus, for instance, the credible sets reported in the macroeconometric literature for impulse responses in VARs identified with sign restrictions do not

[1] Non-stationary time series models provide another prominent case in which the asymptotic equivalence does not hold, e.g. Sims and Uhlig (1991).

provide valid frequentist confidence sets. Vice versa, the frequentist confidence sets reported in the microeconometric literature appear overly conservative from a Bayesian perspective.

The contribution of this paper is to make the previous statement rigorous and compare Bayesian credible sets and frequentist confidence sets for partially identified parameters. We focus on highest-posterior-density (HPD) sets, which have the smallest volume among all Bayesian credible sets and are widely used in the presentation of empirical results. Starting point of our analysis is a likelihood function indexed by a finite-dimensional, identifiable reduced-form parameter vector ϕ. Conditional on ϕ the structural parameter of interest θ is known to lie in the identified set $\Theta(\phi)$. Using a large sample approximation of the posterior of θ we show that under some suitable regularity conditions the HPD set of θ converges to the highest-density set associated with the prior distribution of θ conditional on the maximum-likelihood estimate $\hat{\phi}_n$. Since this conditional prior distribution has necessarily support on the estimated identified set $\Theta(\hat{\phi}_n)$, we deduce that HPD sets tend to exclude portions of the estimated set $\Theta(\hat{\phi}_n)$ in large samples. Frequentist confidence sets, on the other hand, have to extend beyond the boundaries of $\Theta(\hat{\phi}_n)$ to ensure the desired coverage probabilities for all elements $\theta \in \Theta(\phi_0)$, where ϕ_0 denotes the "true" value of the reduced form parameter.

The preceding analysis delivers our main conclusion as well as a recommendation for Bayesian inference. First, in the context of partially identified models $1-\tau$ credible sets tend to be smaller than $1-\tau$ confidence sets. Second, since the prior distribution assigned to $\Theta(\phi)$ conditional on ϕ does not get updated in view of the data – only the distribution of ϕ is updated by the likelihood function – it is important to report an estimate of the identified set, $\Theta(\hat{\phi}_n)$, and the prior of θ given $\phi = \hat{\phi}_n$, where $\hat{\phi}_n$ is a posterior estimate of ϕ.

Earlier work on Bayesian analysis of partially identified models, most notably Poirier (1998), has focused on the characterization of finite-sample posteriors in specific econometric models.[2] We contribute to the Bayesian literature by providing a large sample approximation

[2]After the first draft of this paper had been circulated Bollinger and van Hasselt (2008) derived finite-sample posteriors for a partially identified model with binary misclassification. Moreover, Liao and Jiang (2010) analyzed posterior distributions for moment inequality models based on limited-information likelihood

of the posterior distribution of a partially identified parameter θ, which is based on an insight that dates back at least to Kadane (1974): beliefs about the reduced form parameter ϕ are updated through the likelihood function, but the conditional distribution of θ given ϕ, which we denote by P_ϕ^θ, remains unchanged in view of new data. Since there exists a large literature dating back to Bernstein (1934), Le Cam (1953), and von Mieses (1965) that establishes the asymptotic normality of posterior distributions for finite-dimensional identifiable (reduced-form) parameters, we use this asymptotic normality of ϕ as the starting point for our analysis and approximate the posterior of θ as a mixture of conditional prior distributions P_ϕ^θ.

There is a rapidly growing literature on the construction of asymptotically valid frequentist confidence sets. Rather than providing a detailed review, we are highlighting some important developments. The literature distinguishes between confidence sets for the structural parameter θ and confidence sets for the identified set $\Theta(\phi)$. Our paper focuses on inference for θ. First, we view θ as the more interesting object in most applications. Second, in models which are expressed in terms of an identifiable reduced-form parameter credible or confidence sets for $\Theta(\phi)$ can easily be constructed by taking the union of identified sets for values of ϕ in a $1-\tau$ credible or confidence set. Imbens and Manski (2004) study frequentist inference for an interval-identified parameter in a treatment model, in which the length of the identified set is estimable at a faster rate than the location. Stoye (2009) extends this analysis to a setting in which both length and location of the identified interval are estimated at the same rate.

Chernozhukov, Hong, and Tamer (2007) consider inference based on a generic criterion function $Q_n(\theta; Y^n) \geq 0$, where Y^n denotes the sample observations and the objective function has the property that $Q_n(\theta; Y^n) = 0$ whenever $\theta \in \hat{\Theta}$. Such an objective function arises, for instance, in the context of models defined by moment inequalities. A key challenge in this literature is to find critical values $c_{\tau,n}(\theta)$ such that sets of the form $C_{F,\tau}^\theta(Y^n) = \{\theta \mid Q_n(\theta; Y^n) \leq c_{\tau,n}^2(\theta)\}$ lead to asymptotically valid confidence sets. The literature has considered fixed critical values as well as θ-specific critical values based on

functions.

plug-in asymptotics, selection of binding moment conditions, sub-sampling, and bootstrap approaches, e.g., Rosen (2008), Andrews and Guggenberger (2009), Andrews and Soares (2010), Bugni (2010), Romano and Shaikh (2010). When we compare Bayesian credible sets to frequentist confidence sets, we consider confidence sets constructed from an objective function $Q_n(\theta; \hat{\phi}_n) \geq 0$ with the property that $Q_n(\theta; \hat{\phi}_n) = 0$ if and only if $\Theta(\hat{\phi}_n)$. Since the data Y^n are restricted to enter the objective function only through the estimator $\hat{\phi}_n$ of the reduced form parameter, the frequentist procedures considered in this paper can be viewed as minimum-distance approaches, which are quite natural in problems with well-defined reduced-form parameters. In Moon, Schorfheide, Granziera, and Lee (2009) we use the minimum-distance approach to construct frequentist error bands for a VAR identified with sign restrictions.

The remainder of the paper is organized as follows. In Section 2 we derive finite-sample Bayesian credible intervals and frequentist confidence intervals for a simple moment inequality model with a scalar structural parameter. The structure of this example resembles the above-mentioned interval-identified treatment model of Imbens and Manski (2004). The large sample analysis is presented in Section 3 and Section 4 considers a two-player entry game to illustrate the theoretical results. Finally, Section 5 concludes. The main proofs and key derivations are collected in an Appendix to this paper. Details of the analysis of the entry game are provided in a Technical Appendix that is available electronically.

Finally, a word on notation. We often use M to denote a generic finite constant and I to denote the identity matrix. The notation $\subseteq$ is used to denote weak inclusion and $\subset$ is used for strict inclusion. When X is a matrix and W is a positive-definite weight matrix, let $\|X\|_W = (tr\,(WX'X))^{1/2}$. We use $\|X\|$ (without the W subscript) to denote the Euclidean norm. When P and Q are probability measures, then $\|P - Q\|$ denotes their total variation or L_1 distance. We use $N(\mu, \Sigma)$ to denote the multivariate normal distribution with mean μ and covariance matrix Σ. We let $\varphi_N(\cdot)$ and $\Phi_N(\cdot)$ denote the probability density (pdf) and cumulative density (cdf) functions of a vector of standard normal random variables. Moreover, we denote the one-sided critical value for a standard normal random variable by

$z_\tau = |\Phi_N^{-1}(\tau)|$. $\mathcal{U}[a,b]$ signifies the uniform distribution on the interval $[a,b]$. We use P_b^a to denote the probability distribution of a random variable a conditional on the realization of a random variable b. $I\{X \leq \xi\}$ is the indicator function that is equal to one if $X \leq \xi$ and zero otherwise. "$\stackrel{\mathbb{P}}{\longrightarrow}$" indicates convergence in probability, "$\Longrightarrow$" is convergence in distribution, and w.p.a. 1 abbreviates "with probability approaching one."

2 A Simple Example

We provide a simple example of a partially identified model to illustrate that Bayesian credible sets and frequentist confidence sets are numerically different even as the sample size tends to infinity. Consider the Gaussian location model $Y_i = \phi + U_i$, where $U_i \sim iidN(0,1)$. We shall refer to $\phi \in \mathbb{R}$ as reduced-form parameter. The log likelihood function for this model is quadratic and maximized at the sample mean $\hat{\phi}_n = \frac{1}{n}\sum_{i=1}^n Y_i$. Under a flat prior $p(\phi) \propto c$ the Bayesian posterior distribution of $\sqrt{n}(\phi - \hat{\phi}_n)$ is $N(0,1)$ and thereby identical to the sampling distribution of $\sqrt{n}(\hat{\phi}_n - \phi)$. As a consequence the Bayesian $1-\tau$ credible interval $CS_B^\phi = \left[\hat{\phi}_n - z_{\tau/2}/\sqrt{n},\ \hat{\phi}_n + z_{\tau/2}/\sqrt{n}\right]$ is also a valid frequentist confidence interval.

Now suppose that the object of interest is not the reduced-form parameter ϕ but rather a structural parameter θ that can be bounded based on ϕ as follows:

$$\phi \leq \theta \quad \text{and} \quad \theta \leq \phi + \lambda.$$

The interval $\Theta(\phi) = [\phi,\ \phi + \lambda]$ is called the identified set and in our simple example its length is known to be λ.[3] If we maintain the flat prior for ϕ and use $p(\theta|\phi)$ to denote the properly normalized prior density of θ conditional of ϕ with support $\Theta(\phi)$, then the joint

[3]Our example resembles the treatment model in Imbens and Manski (2004). In their model $Y_i \in [0,1]$ is a random outcome. Outcomes are only observed for individuals that received treatment as indicated by a binary variable $D_i \in \{0,1\}$. The identifiable reduced form parameter can be defined as $\phi = \mathbb{E}[Y_i D_i]$ and the structural parameter of interest is $\theta = \mathbb{E}[Y_i]$. Under the assumption that the probability of treatment $p = \mathbb{E}[D_i]$ lies between zero and one and is known, the identified set is given by $\Theta(\phi) = [p\phi,\ p\phi + (1-p)]$.

distribution of data and parameters is given by

$$p(Y^n, \phi, \theta) \propto p(Y^n|\phi)p(\theta|\phi), \tag{1}$$

where $Y^n = [y_1, \ldots, y_n]$. From integrating (1) with respect to θ we deduce that the marginal posterior of ϕ remains normal. It is also immediately apparent from

$$p(Y^n, \theta|\phi) \propto p(Y^n|\phi)p(\theta|\phi) \tag{2}$$

that the data Y^n and the parameter vector θ are independent conditional on ϕ because θ does not enter the likelihood function. Consequently, the conditional distribution of θ given ϕ is not updated in view of the data:

$$p(\theta|Y^n, \phi) \propto p(\theta|\phi). \tag{3}$$

Let $s = \sqrt{n}(\phi - \hat{\phi}_n)$. Then the marginal posterior density of θ can be expressed as the mixture

$$p(\theta|Y^n) = \int p(\theta|\hat{\phi}_n + n^{-1/2}s)\varphi_N(s)ds. \tag{4}$$

This equation suggests that as the sample size $n \longrightarrow \infty$ the posterior density of θ converges to $p(\theta|\hat{\phi}_n)$.

Before we present a formal derivation of the asymptotic approximation, we proceed with the finite-sample analysis of our simple model. For concreteness, we assume that the prior for θ conditional on the reduced form parameter ϕ is uniform on the interval $\Theta(\phi)$:

$$p(\theta|\phi) = \frac{1}{\lambda} I\{\phi \leq \theta \leq \phi + \lambda\}.$$

It can be shown by direct calculation (see Appendix) that

$$p(\theta|Y^n) = \frac{1}{\lambda}\left[\Phi_N\left(\sqrt{n}(\theta - \hat{\phi}_n)\right) - \Phi_N\left(\sqrt{n}(\theta - \hat{\phi}_n - \lambda)\right)\right]. \tag{5}$$

This density is depicted in Figure 1 for $\hat{\phi}_n = 0$, $n = 100$, and various choices of λ. If λ is small (relative to the sample size) the posterior density is approximately normal whereas for larger values of λ it resembles a step function. Theorem 1(ii) presented in Section 3 implies that for any fixed value of $\lambda > 0$

$$\int_\theta \left| p(\theta|Y^n) - \frac{1}{\lambda} I\{\hat{\phi}_n \leq \theta \leq \hat{\phi}_n + \lambda\} \right| d\theta = o_p(1). \tag{6}$$

Thus, the L_1 distance between the posterior density and the prior density $p(\theta|\hat{\phi}_n)$ on the estimated set $\Theta(\hat{\phi}_n)$ converges to zero.

The posterior density characterized by (5) is symmetric around $\hat{\phi}_n + \lambda/2$ and so are HPD sets constructed from this posterior. Since the posterior density is continuous, we can express the HPD set as

$$CS_B^\theta(Y^n) = \left[\hat{\phi}_n + \tau/2 - \eta_n,\ \hat{\phi}_n + \lambda - \tau/2 + \eta_n \right]. \tag{7}$$

Here η_n is chosen to guarantee that $P^\theta_{Y^n}\{\theta \in CS_B^\theta(Y^n)\} = 1 - \tau$. The convergence in (6) implies that

$$\left| P^\theta_{Y^n}\{\theta \in CS_B^\theta(Y^n)\} - P^\theta_{\hat{\phi}_n}\{\theta \in CS_B^\theta(Y^n)\} \right| = 2|\eta_n| \xrightarrow{p} 0, \tag{8}$$

where $P^\theta_{\hat{\phi}_n}$ is the conditional prior distribution of θ given that $\phi = \hat{\phi}_n$. Thus, for n sufficiently large, the HPD interval has to lie in the interior of the estimated set $\Theta(\hat{\phi}_n)$.

The frequentist analysis is markedly different. We can parameterize the correspondence as $\theta = \phi + \alpha$, where $\alpha \in [0, \lambda]$. Following Chernozhukov, Hong, and Tamer (2007) we construct a frequentist confidence set as a level set of the concentrated log inverse of the likelihood function. Define[4]

$$\begin{aligned} Q_n(\theta; \hat{\phi}_n) &= \inf_{\alpha \in [0,\lambda]} -2[\ln p(Y^n|\theta + \alpha) - \ln p(Y^n|\hat{\phi}_n)] \\ &= \begin{cases} n(\hat{\phi}_n - \theta)^2 & \text{if } \theta \leq \hat{\phi}_n \\ 0 & \text{if } \hat{\phi}_n < \theta < \hat{\phi}_n + \lambda \\ n(\hat{\phi}_n - \theta + \lambda)^2 & \text{if } \hat{\phi}_n + \lambda \leq \theta \end{cases} \end{aligned} \tag{9}$$

[4] In this simple example $\hat{\phi}_n$ is a sufficient statistic. Thus, conditioning on $\hat{\phi}_n$ instead of Y^n imposes no restrictions.

and let

$$CS_F^\theta(Y^n) = \left\{ \theta \;\middle|\; Q_n(\theta; \hat{\phi}_n) \leq c_\tau^2 \right\}. \tag{10}$$

The finite-sample distribution of the maximum likelihood estimator is $\sqrt{n}(\hat{\phi}_n - \phi) \sim \mathcal{Z}$, where $\mathcal{Z} \sim N(0,1)$. It is convenient to re-scale θ according to $\vartheta = \sqrt{n}(\theta - \phi)$. In terms of the ϑ transform, the identified set $\Theta(\phi)$ is given by $0 \leq \vartheta \leq \sqrt{n}\lambda$. We can now characterize the distribution of the profile objective function as

$$Q_n(\phi + n^{-1/2}\vartheta; \hat{\phi}_n) \sim \begin{cases} (\mathcal{Z} - \vartheta)^2 & \text{if } \vartheta \leq \mathcal{Z} \\ 0 & \text{if } \mathcal{Z} < \vartheta < \mathcal{Z} + \sqrt{n}\lambda \\ (\mathcal{Z} - \vartheta + \sqrt{n}\lambda)^2 & \text{if } \mathcal{Z} + \sqrt{n}\lambda \leq \vartheta \end{cases}. \tag{11}$$

Define c_τ as the solution of the following equation:

$$\Phi_N(\sqrt{n}\lambda + c_\tau) - \Phi_N(-c_\tau) = 1 - \tau. \tag{12}$$

In view of (11) we deduce that c_τ is the correct critical value that generates a uniformly valid confidence set:

$$\begin{aligned} & \inf_{\phi \in \Phi} \inf_{\theta \in \Theta(\phi)} P_\phi^{Y^n}\{\theta \in CS_F^\theta(Y^n)\} \\ = & \inf_{0 \leq \vartheta \leq \sqrt{n}\lambda} P\{\vartheta - \sqrt{n}\lambda - c_\tau \leq \mathcal{Z} \leq \vartheta + c_\tau\} \\ = & \inf_{0 \leq \vartheta \leq \sqrt{n}\lambda} \Phi_{\mathcal{N}}(\vartheta + c_\tau) - \Phi_{\mathcal{N}}(\vartheta - \sqrt{n}\lambda - c_\tau) \\ = & \; 1 - \tau. \end{aligned}$$

The last equality follows from the definition c_τ in (12) and because the infimum is achieved at $\vartheta = 0$ and $\vartheta = \sqrt{n}\lambda$. Therefore, the resulting confidence interval is of the form

$$CS_F^\theta(Y^n) = \left[\hat{\phi}_n - c_\tau/\sqrt{n}, \; \hat{\phi}_n + \lambda + c_\tau/\sqrt{n} \right]. \tag{13}$$

As pointed out by Imbens and Manski (2004), if the re-scaled length of the identified set, $\sqrt{n}\lambda$ is large, then a $1 - \tau$ confidence set for the parameter θ is obtained by expanding the boundaries of the interval $\Theta(\hat{\phi}_n)$ using a one sided critical value of a standard normal distribution.

A comparison of (7) and (13) indicates that the (numerical) difference between credible and confidence sets does not vanish asymptotically. The Bayesian HPD set $CS_B^\theta(Y^n)$ is eventually located strictly inside the estimated identified set $\Theta(\hat{\phi}_n) = [\hat{\phi}_n,\ \hat{\phi}_n + \lambda]$, while the frequentist confidence set $CS_F^\theta(Y^n)$ extends beyond the boundaries of $\Theta(\hat{\phi}_n)$:

$$CS_B^\theta(Y^n) \subset \Theta(\hat{\phi}_n) \subset CS_F^\theta(Y^n)$$

eventually. This implies that the posterior probability of the confidence set tends to one

$$\text{plim}_{n\longrightarrow\infty}\, P^\theta_{Y^n}\{\theta \in CS_F^\theta(Y^n)\} = 1,$$

while the asymptotic coverage probability of the credible set is zero,

$$\lim_{n\to\infty} \inf_{\phi\in\mathbb{R}} \inf_{\theta\in\Theta(\phi)} P_\phi^{Y^n}\{\theta \in CS_B^\theta(Y^n)\} = 0.$$

Thus, from the Bayesian perspective, the frequentist confidence set is too wide, while from the frequentist perspective, the Bayesian credible set is too narrow.

In Table 1 we compare 80% Bayesian credible sets and frequentist confidence intervals, computed based on (7) and (13), for various choices of n and the length of the identified set λ. Since in our example both the confidence interval as well as the credible interval are centered at $\hat{\phi}_n + \lambda/2$ we restrict the comparison to the length of the intervals and the posterior probability that θ falls into the frequentist confidence set. If the length of the identified set is small relative to the sample size, that is, θ is approximately point identified, then confidence and credible intervals are essentially numerically identical, that is they have the same length and and the posterior probability assigned to the interval $CS_F^\theta(Y^n)$ is approximately equal to the frequentist coverage probability of 80%. As the length of the identified interval increases or the uncertainty about ϕ decreases, the table illustrates that an 80% Bayesian credible set is shorter than an 80% frequentist confidence interval.

3 Large Sample Approximations

We will now generalize the analysis presented in the previous section. We shall assume that the joint prior distribution for the $m \times 1$ reduced form parameter vector ϕ and the $k \times 1$

structural parameter vector θ can be decomposed into a marginal distribution for $\phi \in \Phi$, denoted by P^{ϕ}, and a conditional distribution P^{θ}_{ϕ} of θ. The conditional distribution has support on the identified set $\Theta(\phi)$. Let $l_n(\phi)$ denote the log likelihood function $\ln p(Y^n|\phi)$. Kadane (1974) emphasized that the derivation of the posterior distribution can be done on the space of the reduced form parameter ϕ. The calculations for the example in Section 2 can be generalized as follows. For any measurable set $A \subseteq \Theta$:

$$\begin{aligned} P^{\theta}_{Y^n}(A) &= \frac{\int_{\Phi}\int_{\Theta(\phi)} I\{\theta \in A\}\exp[l_n(\phi)]dP^{\theta}_{\phi}dP^{\phi}}{\int_{\Phi}\int_{\Theta(\phi)}\exp[l_n(\phi)]dP^{\theta}_{\phi}dP^{\phi}} \\ &= \int_{\Phi}\left[\int_{\Theta(\phi)} I\{\theta \in A\}dP^{\theta}_{\phi}\right]\frac{\exp[l_n(\phi)]}{\int_{\Phi}\exp[l_n(\phi)]dP^{\phi}}dP^{\phi} \\ &= \int_{\Phi} P^{\theta}_{\phi}(A)dP^{\phi}_{Y^n}. \end{aligned} \tag{14}$$

Since conditional on ϕ the structural parameter θ does not enter the likelihood function the prior distribution of θ given ϕ, P^{θ}_{ϕ}, is not updated in view of the data Y^n. This point also had been stressed by Poirier (1998).

The goal of this section is to provide a rigorous and insightful large sample approximation of the marginal posterior distribution $P^{\theta}_{Y^n}$ as well as Bayesian HPD sets. In a nutshell the approximation of the posterior takes the following form: As the sample size increases, the posterior distribution of ϕ concentrates in a small neighborhood of the maximum likelihood estimator $\hat{\phi}_n$ and in view of (14) we can approximate $P^{\theta}_{Y^n}$ by $P^{\theta}_{\hat{\phi}_n}$. This approximation is obtained in two steps. First, in Section 3.1 $P^{\phi}_{Y^n}$ in (14) is replaced by its normal limit distribution. Second, under the assumption that the conditional prior distribution P^{θ}_{ϕ} is Lipschitz in ϕ we show in Section 3.2 that the posterior of θ indeed converges to $P^{\theta}_{\hat{\phi}_n}$ or to $P^{\theta}_{\phi_0}$ if $\hat{\phi}_n$ has a probability limit ϕ_0. Section 3.3 examines the convergence of HPD sets and delivers the main result: under certain regularity conditions the difference between finite-sample HPD sets and highest density sets constructed from the distribution $P^{\theta}_{\hat{\phi}_n}$ vanishes asymptotically. Thus, the HPD sets are asymptotically located inside the estimated identified set and thereby smaller than frequentist confidence sets, which extend beyond the boundaries of $\Theta(\hat{\phi}_n)$. The comparison between HPD and confidence sets is presented in Section 3.4.

Our approximations are stated in terms of the total-variation or L_1 distance between probability measures. Let P and Q be two probability measures on the space $(\Omega, \mathcal{F})$ and f be mapping from Ω into the real line. The L_1 distance can be defined, e.g. Pollard (2002), as

$$\|P - Q\| = \sup_{|f| \leq 1} \left| \int f dP - \int f dQ \right|,$$

where $\sup_{|f|\leq 1}$ denotes the supremum over all functions f that are bounded by one in absolute value. If P and Q have densities $p(\omega)$ and $q(\omega)$ with respect to Lebesgue measure, then

$$\|P - Q\| = \int |p(\omega) - q(\omega)| d\omega.$$

We decided to adopt the concept of L_1 convergence instead of the weaker notion of convergence in distribution because we want to establish the convergence of HPD sets in Section 3.3. Convergence in distribution results can be found in an earlier version of this paper (Moon and Schorfheide, 2009).

3.1 First Approximation: $p(\phi|Y^n)$

There exists a large literature on the asymptotic normality of posterior distributions in identified models dating back to Bernstein (1934), Le Cam (1953), and von Mieses (1965). Since the goal of our paper is not to make an independent contribution to this literature, we will state the asymptotic normality of $P^{\phi}_{Y^n}$ as an assumption rather than derive it from more low-level conditions on the likelihood function. Such fundamental conditions can be found, for instance, in Johnson (1970) and in the textbook treatments of Hartigan (1983, Section 11.2), van der Vaart (1998, Theorem 10.1), or Ghosh and Ramamoorthi (2003, Theorem 1.4.2). Extensions to time series models are provided by Kim (1998) and Phillips and Ploberger (1998). To state our basic assumption, we need to introduce some additional notation. Below we shall assume that the log-likelihood function is twice continuously differentiable with respect to ϕ such that we can define

$$\hat{J}_n = D_n^{-1} \left[-\frac{\partial^2 l_n(\phi)}{\partial \phi \partial \phi'} \right]_{\phi = \hat{\phi}_n} D_n^{-1'}.$$

The matrix $\hat{J}_n^{-1}$ plays the role of an asymptotic posterior covariance matrix of ϕ. We transform the reduced-form parameters according to

$$s = \hat{J}_n^{1/2} D_n(\phi - \hat{\phi}_n) \tag{15}$$

and use $P^s_{Y^n}$ to denote the posterior distribution of s. The matrix D_n is deterministic with elements that are diverging as $n \longrightarrow \infty$. It is chosen to ensure that $\hat{J}_n$ is convergent. In models with *iid* data D_n is typically diagonal with elements $\sqrt{n}$. However, in time series models that contain trending regressors often a more general form of D_n is required, see for instance, Phillips and Ploberger (1998) and Kim (1998).

Assumption 1 *Let $Y^n(\omega)$ be a sequence of random vectors defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ (i) The sequence of maximum likelihood estimators is convergent: $\hat{\phi}_n \stackrel{\mathbb{P}}{\longrightarrow} \phi_0$. The matrix $\|D_n\| \uparrow \infty$. The likelihood function $l_n(\phi)$ is twice continuously differentiable w.p.a. 1, such that $\hat{J}_n$ is well defined. Moreover, the Hessian of the log-likelihood function has a positive definite probability limit: $\hat{J}_n \Longrightarrow J_0 > 0$ and $\hat{J}_n^{-1} \Longrightarrow J_0^{-1}$. (ii) The posterior distribution of ϕ is asymptotically normal:*

$$\|P^s_{Y^n} - N(0, I)\| \stackrel{\mathbb{P}}{\longrightarrow} 0,$$

where $s = \hat{J}_n^{1/2} D_n(\phi - \hat{\phi}_n)$.

Assumption 1(ii) states the asymptotic normality in terms of L_1 convergence of probability measures. Conditions for L_1 convergence are provided in van der Vaart (1998) and Gosh and Ramamoorthi (2003). Not all of the assumptions stated in Part (i) need to be satisfied in order to obtain a posterior distribution that is asymptotically normal. For instance, it could be the case that the maximum likelihood estimator is not convergent under $\mathbb{P}$, but the posterior distribution is still asymptotically normal around $\hat{\phi}_n$. However, the low-level assumptions under which the asymptotic normality is proved, typically are consistent with the assumptions in Part (i). Assuming the convergence of $\hat{\phi}_n$ allows us to state conditions on the conditional prior P^θ_ϕ locally, in a neighborhood of ϕ_0. While in many models the sequence of

Hessian matrices $\hat{J}_n$ converges in probability to a constant matrix, non-stationary time-series models provide an important class of models for which $\hat{J}_n$ converges to a stochastic limit.

In view of the large literature on extremum estimators based on non-differentiable objective functions, the assumption of a twice differentiable likelihood function might appear overly restrictive. While it rules out, for instance, likelihood functions constructed from densities with jumps, densities with singularities, or change-point problems, we note that most of the applied Bayesian analysis tends to be based on differentiable likelihood functions. When the likelihood function is only first-order stochastically differentiable, Bayesian posterior asymptotics require an alternative definition of $\hat{J}_n$ and are provided, for example, in van der Vaart (1998). Further large-sample approximations for Bayesian posterior in irregular models can be found, for instance, in Ghosh, Ghosal, and Samanta (1994, 1995). In these models one can often establish that the posterior of a properly centered and scaled parameter converges to the limit of a suitably defined likelihood ratio process, which might lead to non-Gaussian approximations.

3.2 Second Approximation: $p(\theta|Y^n)$

The object of interest in this paper is the posterior distribution of the structural parameter θ given by (14). We begin our analysis by replacing the posterior distribution $P^{\phi}_{Y^n}$ in (14) with its asymptotic normal distribution $N(0,I)$. For any measurable set $A \subseteq \Theta$ let

$$P^{\theta}_{N,Y^n}(A) = \int_{\mathbb{R}^m} P^{\theta}_{\hat{\phi}_n + D_n^{-1}\hat{J}_n^{-1/2}s}(A)\, dN\big(0,I\big)(s). \tag{16}$$

It can be easily verified that P^{θ}_{N,Y^n} provides a large sample approximation of $P^{\theta}_{Y^n}$. For any measurable real valued function $f(\theta)$ such that $|f| \leq 1$, one can define the function

$$g(s) = \int f(\theta) dP^{\theta}_{\hat{\phi}_n + D_n^{-1}\hat{J}_n^{-1/2}s}$$

where $|g(s)| \leq 1$. Thus,

$$\left| \int f(\theta) dP^{\theta}_{Y^n} - \int f(\theta) dP^{\theta}_{N,Y^n} \right| = \left| \int g(s) dP^{s}_{Y^n} - \int g(s) dN\big(0,I\big)(s) \right| \tag{17}$$

According to Assumption 1(ii) and the definition of L_1 convergence, the right-hand-side of (17) converges in probability to zero for every function $|g(s)| \leq 1$ and we can deduce that

$$\|P^{\theta}_{Y^n} - P^{\theta}_{N,Y^n}\| \xrightarrow{\mathbb{P}} 0. \tag{18}$$

While the distribution P^{θ}_{N,Y^n} provides a valid large sample approximation of the posterior, we shall construct a second approximation that is more insightful in regard to comparisons between Bayesian credible sets and frequentist confidence sets. This second approximation requires an additional assumption, which we will discuss in more detail after stating the main result.

Assumption 2 *Let $N_\delta(\phi_0) = \{\phi \in \Phi \mid \|\phi - \phi_0\| < \delta\}$. For $\phi_1, \phi_2 \in N_\delta(\phi_0)$ there exists a constant $M(\phi_0, \delta)$ such that $\|P^{\theta}_{\phi_1} - P^{\theta}_{\phi_2}\| \leq M(\phi_0, \delta)\|\phi_1 - \phi_2\|$.*

Theorem 1 *Let $s = \hat{J}_n^{1/2} D_n(\phi - \hat{\phi}_n)$. (i) If Assumption 1 is satisfied*

$$\|P^{\theta}_{Y^n} - P^{\theta}_{N,Y^n}\| \xrightarrow{\mathbb{P}} 0.$$

(ii) If Assumptions 1 and 2 are satisfied, then

$$\|P^{\theta}_{Y^n} - P^{\theta}_{\hat{\phi}_n}\| \xrightarrow{\mathbb{P}} 0 \quad \text{and} \quad \|P^{\theta}_{Y^n} - P^{\theta}_{\phi_0}\| \xrightarrow{\mathbb{P}} 0,$$

where $P^{\theta}_{\hat{\phi}_n}$ ($P^{\theta}_{\phi_0}$) denotes the conditional prior P^{θ}_{ϕ} evaluated at $\phi = \hat{\phi}_n$ ($\phi = \phi_0$).

The proof of Part (i) was provided above. Part (ii) states that the posterior distribution of θ can be approximated either by the conditional prior distribution P^{θ}_{ϕ} evaluated at the maximum likelihood estimate $\hat{\phi}_n$ or at its probability limit ϕ_0. A proof is provided in the Appendix.

Remark 1: Consider an interval-identified location with $m = 1$, $k = 1$, $\Theta(\phi) = [\phi, \phi + \lambda]$ and a prior density of the form

$$p(\theta|\phi) = \frac{1}{\lambda} f\left(\frac{\theta - \phi}{\lambda}\right), \tag{19}$$

where $\int f(x)dx = 1$ and $f(x) = 0$ outside of the unit interval. Here $\lambda > 0$ is the known length of the identified set. Moreover, assume that $f(x)$ is differentiable with uniformly bounded first derivative $|df/dx| < M_1$ for $0 \leq x \leq 1$. Our assumptions imply that the function $f(x)$ can be bounded by a constant M_0. This model provides a slight generalization of the example studied in Section 2. We begin by verifying Assumption 2. Without loss of generality, let $\phi_1 \leq \phi_2$. For $\phi_2 \geq \phi_1 + \lambda$ we obtain the bound

$$\int |p(\theta|\phi_1) - p(\theta|\phi_2)|d\theta = 2 \leq 2\frac{\|\phi_1 - \phi_2\|}{\lambda}. \tag{20}$$

For $\phi_2 < \phi_1 + \lambda$ we obtain

$$|p(\theta|\phi_1) - p(\theta|\phi_2)| = \begin{cases} p(\theta|\phi_1) & \text{if } \phi_1 \leq \theta \leq \phi_2 \\ |p(\theta|\phi_1) - p(\theta|\phi_2)| & \text{if } \phi_2 \leq \theta \leq \phi_1 + \lambda \\ p(\theta|\phi_2) & \text{if } \phi_1 + \lambda \leq \theta \leq \phi_2 + \lambda \\ 0 & \text{otherwise} \end{cases}$$

Both $p(\theta|\phi_1)$ and $p(\theta|\phi_2)$ are bounded by M_0. On the set $\phi_2 \leq \theta \leq \phi_1 + \lambda$ a first-order Taylor expansion of the prior density implies that

$$|p(\theta|\phi_1) - p(\theta|\phi_2)| \leq \frac{M_1}{\lambda^2}\|\phi_1 - \phi_2\|.$$

Thus, integrating over θ we deduce

$$\int |p(\theta|\phi_1) - p(\theta|\phi_2)|d\theta \leq 2M_0\|\phi_1 - \phi_2\| + \frac{M_1}{\lambda}\|\phi_1 - \phi_2\|. \tag{21}$$

To obtain a bound for the integral over the range $\phi_2 \leq \theta \leq \phi_1 + \lambda$ we used the fact that $\phi_1 \leq \phi_2$, which implies that $(\phi_1 + \lambda - \phi_2)/\lambda \leq 1$. Combining (20) and (21) by setting $M = \max\{2/\lambda,\ 2M_0 + M_1/\lambda\}$ provides a verification of Assumption 2. ■

Remark 2: If the length λ of the identified set is zero or the continuous prior distribution on the identified set with length $\lambda > 0$ is replaced by a point mass, the Lipschitz condition will fail. Let P_ϕ^θ be the probability measure that assigns mass one to $\theta = \phi + \tau$, where $0 \leq \tau \leq \lambda$. Then for every $\phi_1 \neq \phi_2$

$$\|P_{\phi_1}^\theta - P_{\phi_2}^\theta\| \geq |P_{\phi_1}^\theta(\{\phi_1 + \tau\}) - P_{\phi_2}^\theta(\{\phi_1 + \tau\})| = 1.$$

Alternatively, if one considers a sequence of models with $\Theta_n(\phi) = [\phi,\ \phi + \lambda_n]$ and priors

$$p_n(\theta|\phi) = \frac{1}{\lambda_n} f\left(\frac{\theta - \phi}{\lambda_n}\right)$$

where the length of the identified $\lambda_n \longrightarrow 0$ fast enough such that $D_n\lambda_n \longrightarrow 0$, it can be verified by direct calculation (see Appendix) that

$$\|P^{\theta}_{Y^n} - N(\hat{\phi}_n, D_n^{-2}\hat{J}_n^{-1})\| \stackrel{\mathbb{P}}{\longrightarrow} 0. \tag{22}$$

The approximation in Theorem 1(ii) implies that the uncertainty about the reduced-form parameter ϕ is asymptotically negligible and the shape of the posterior is determined by the conditional prior $P^{\theta}_{\hat{\phi}_n}$. However, this conclusion only holds, if the identified set is large relative to the uncertainty about its location. If the size of the identified set is small, the uncertainty about ϕ dominates the posterior and the shape of P^{θ}_{ϕ} is essentially irrelevant as in (22). This result is consistent with the simulations presented in Section 2. According to Table 1 for $\lambda = 0.01$ the length of the credible interval for $n = 100$ ($n = 500$) is 0.26 (0.12), which corresponds to $2z_{0.1}/\sqrt{n}$, that is an interval constructed from a $N(\hat{\phi}_n, 1/n)$ distribution. For $\lambda = 10$, on the other hand, the posterior credible interval has length 8.0 and corresponds to an interval with 80% probability under a $\mathcal{U}[\hat{\phi}_n, \hat{\phi}_n + 10]$ distribution, which is $P^{\theta}_{\hat{\phi}_n}$. ■

Remark 3: In many applications, including the entry game discussed in Section 4, the identified set $\Theta(\phi)$ lies in a lower dimensional subspace of Θ. This tends to lead to a violation of Assumption 2 with respect to the joint distribution of all elements of the θ vector, but the Lipschitz condition might still be satisfied for various marginal distributions as the following example illustrates. Suppose the interval-identified location model is modified as follows. We maintain $m = 1$ but increase the dimension of the structural parameter to $k = 2$. The identified set is given by

$$\Theta(\phi) = \left\{(\theta_1, \theta_2) \in \mathbb{R}^2 \mid \phi \leq \theta_1 \leq \phi + \lambda,\ \theta_2 = \theta_1 - \phi\right\}.$$

By construction the intersection of $\Theta(\phi_1)$ and $\Theta(\phi_2)$ for $\phi_1 \neq \phi_2$ is empty. Thus, the L_1 distance between measures $P^{\theta}_{\phi_1}$ and $P^{\theta}_{\phi_2}$ is 2 if $\phi_1 \neq \phi_2$ and the Lipschitz condition is violated.

In order to specify a prior distribution P^θ_ϕ it suffices to specify a distribution $P^{\theta_1}_\phi$. Given the shape of the identified set the distribution of $P^{\theta_2}_{(\phi,\theta_1)}$ is simply a point mass at $\theta_2 = \theta_1 - \phi$. Suppose we use the density (19) to characterize $P^{\theta_1}_\phi$ and thereby P^θ_ϕ. While the Lipschitz condition is violated for P^θ_ϕ, our previous calculations imply that the condition is satisfied for the marginal distribution of θ_1 given ϕ. This marginal has support on the projection of the identified set $\Theta(\phi)$ onto the domain of θ_1, which is given by $\Theta_1(\phi) = [\phi,\, \phi+\lambda]$. A simple change of variables $\theta_2 = \theta_1 - \phi$ can be used to confirm that the Lipschitz condition is also satisfied for $P^{\theta_2}_\phi$, which has support on $\Theta_2(\phi) = [0,\, \lambda]$. ■

3.3 Highest-Posterior-Density Sets for θ

We now turn to the construction of HPD sets. We shall assume that the prior distributions P^θ_ϕ, the posterior distribution $P^\theta_{Y^n}$, and its large sample approximations P^θ_{N,Y^n} and $P^\theta_{\hat{\phi}_n}$ have densities $p(\theta|\phi)$, $p(\theta|Y^n)$, $p_N(\theta|Y^n)$, and $p(\theta|\hat{\phi}_n)$ with respect to Lebesgue measure in $\mathbb{R}^k$. Define the finite-sample HPD set as

$$CS^\theta_{HPD}(Y^n) = \{\theta \mid p(\theta|Y^n) \geq \kappa_{Y^n}\},$$

where the cut-off κ_{Y^n} is chosen to ensure that $P^\theta_{Y^n}(CS^\theta_{HPD}(Y^n)) = 1 - \tau$. We also define a highest-prior-density set[5] conditional on ϕ as

$$CS^\theta_{HPD}(\phi) = \{\theta \mid p(\theta|\phi) \geq \kappa_\phi\}$$

with $P^\theta_\phi(CS^\theta_{HPD}(\phi)) = 1 - \tau$. If the object of inference is a subvector of θ, say θ_1 then the joint densities $p(\theta|Y^n)$ and $p(\theta|\phi)$ can simply be replaced by marginal densities $p(\theta_1|Y^n)$ and $p(\theta_1|\phi)$, respectively.

The following theorem states that the posterior probability of the symmetric difference between the sets $CS^\theta_{HPD}(Y^n)$ and $CS^\theta_{HPD}(\phi_0)$ as well as the difference between $CS^\theta_{HPD}(Y^n)$ and $CS^\theta_{HPD}(\hat{\phi}_n)$ converges to zero. We adopt this notion of convergence of a sequence

[5] In slight abuse of notation we use the abbreviation HPD for both highest-posterior-density and highest-prior-density sets.

of finite sample credible sets from Severini (1991, Page 613). In order to establish the convergence result we assume that for $\phi = \phi_0$ the conditional prior distribution P^θ_ϕ allows the construction of a unique HPD set and that the conditional prior density has no flat regions at the threshold level κ_ϕ. This assumption rules out, for instance, a conditional prior distribution that is uniform on the identified set because HPD sets are not uniquely defined for uniform distributions.

Assumption 3 *At $\phi = \phi_0$ the prior density $p(\theta|\phi)$ leads to a unique $1-\tau$ highest-prior-density set with a threshold $\kappa_\phi < \infty$ and $\int I\{\, p(\theta|\phi) = \kappa_\phi\} p(\theta|\phi) d\theta = 0$.*

Theorem 2 *Suppose Assumptions 1, 2, and 3 are satisfied. Then (i):*

$$\int \left| I\{\theta \in CS^\theta_{HPD}(Y^n)\} - I\{\theta \in CS^\theta_{HPD}(\phi_0)\} \right| dP^\theta_{Y^n} \xrightarrow{\mathbb{P}} 0$$

and (ii)

$$\int \left| I\{\theta \in CS^\theta_{HPD}(Y^n)\} - I\{\theta \in CS^\theta_{HPD}(\hat{\phi}_n)\} \right| dP^\theta_{Y^n} \xrightarrow{\mathbb{P}} 0.$$

A proof of the theorem can be found in the appendix. In view of Theorem 1 the distribution $P^\theta_{Y^n}$ can also be replaced by $P^\theta_{\hat{\phi}_n}$ or $P^\theta_{\phi_0}$. Theorem 2 generalizes the results obtained in the context of the simple example in Section 2. By construction $CS^\theta_{HPD}(\hat{\phi}_n) \subset \Theta(\hat{\phi}_n)$. Thus, according to Part (ii) the Bayesian HPD set is located inside of the estimated identified set eventually. Moreover, since $CS^\theta_{HPD}(\phi_0) \subset \Theta(\phi_0)$ Part (i) implies that there exists a subset of $\Theta(\phi_0)$ that is excluded from the Bayesian credible set with probability approaching one. Thus, the credible set does not provide an asymptotically valid confidence set.

3.4 Comparison to Frequentist Confidence Sets

Since we model partial identification through the identified set correspondence $\Theta(\phi)$, we can relate the structural parameter of interest θ and the identified reduced form parameter ϕ via nuisance parameters, say α, $\phi = G(\theta, \alpha)$, where $\alpha \in \mathcal{A}_\theta$:

$$\theta \in \Theta(\phi) \quad \text{iff} \quad \exists\, \alpha \in \mathcal{A}_\theta \text{ such that } \phi = G(\theta, \alpha). \tag{23}$$

Suppose that $\hat{\phi}_n$ is a consistent and asymptotically normal estimator of ϕ. In this case a natural objective function for frequentist inference is the minimum distance criterion function

$$Q_n(\theta; \hat{\phi}_n) = \min_{\alpha \in \mathcal{A}_\theta} \; n \big\| \hat{\phi}_n - G(\theta, \alpha) \big\|_{W_n} \; \geq 0, \tag{24}$$

where $\{W_n\}$ is a sequence of positive definite weight matrices. Now suppose that we consider frequentist confidence sets of the form

$$CS^{\theta}_{F,\tau}(Y^n) = \left\{ \theta \;\middle|\; Q_n(\theta; \hat{\phi}_n) \leq c^2_{\tau,n}(\theta) \right\}. \tag{25}$$

For $CS^{\theta}_{F,\tau}$ to be a confidence set that is asymptotically valid the following condition has to be satisfied

$$\lim_{n \longrightarrow \infty} \inf_{\phi \in \Phi} \inf_{\theta \in \Theta(\phi)} P^{Y^n}_{\phi} \{ Q_n(\theta; \hat{\phi}_n) \leq c^2_{\tau,n}(\theta) \} \geq 1 - \tau. \tag{26}$$

Constructing sequences of critical value functions such that (26) holds is the subject of a rapidly growing literature, e.g., Imbens and Manski (2004), Chernozhukov, Hong, and Tamer (2007), Rosen (2008), Andrews and Guggenberger (2009), Stoye (2009), Andrews and Soares (2010), and Bugni (2010). The goal of this literature is to sharpen the confidence set to achieve the desired coverage probability in (26). However, regardless of the construction of $c_{\tau,n}(\theta)$, a key property of the frequentist objective function is that $Q_n(\theta; \hat{\phi}_n) = 0$ if $\theta \in \Theta(\hat{\phi}_n)$, which follows from (23). Therefore, we can deduce that $\Theta(\hat{\phi}_n)$ is contained in a frequentist confidence set that takes the form of (25):

$$\Theta(\hat{\phi}_n) \subseteq CS^{\theta}_{F,\tau}. \tag{27}$$

In many instances, the weak inclusion can be replaced by a strict inclusion. In view of Theorem 2 we conclude that the frequentist confidence set $CS^{\theta}_{F,\tau}$ is too large from a Bayesian perspective.

4 Illustration: A Two-Player Entry Game

To provide a second numerical illustration of Bayesian credible sets and frequentist confidence sets, we consider an example that has received a lot of attention in the microeconometric literature on partially identified models: a two-player entry game, see for instance Bresnahan and Reiss (1991), Berry (1994), Tamer (2003), and Ciliberto and Tamer (2009). We will focus on a fairly simple version of the entry game without regressors. Depending on the entry decision of the competing firm, Firm j either does not enter market i, operates as monopolist, or operates as duopolist. Potential monopoly (M) and duopoly (D) profits are given by

$$\pi_{i,j}^M = \beta_j + \epsilon_{i,j}, \quad \pi_{i,j}^D = \beta_j - \gamma_j + \epsilon_{i,j}, \quad j = 1, 2, \quad i = 1, \dots, n. \tag{28}$$

The $\epsilon_{i,j}$'s capture latent profit components that are known to the two firms but unobserved by the econometrician. The econometrician simply observes which firm(s) enter each of the n markets. Thus $y_i \in \{(1,1),\ (0,0),\ (1,0),\ (0,1)\}$ and we use n_{11}, n_{00}, n_{10}, n_{01} to denote the frequency with which the four possible market configurations are observed in the sample. We assume that the outcome of the entry game in each market is a pure strategy Nash equilibrium. It is straightforward to verify that the Nash equilibrium is unique, except if both firms are profitable as monopolist but not as duopolist. In the latter case, the model is silent about which firm actually enters the market. As a consequence, the model only delivers bounds for the probability of observing a particular monopoly.

Suppose that $\epsilon_{i,j} \sim iidN(0,1)$ and let $\theta = [\beta_1, \gamma_1, \beta_2, \gamma_2]'$. Using (28) it is straightforward to calculate probabilities that firm j is profitable as monopolist (duopolist) in market i. We denote these probabilities by $m_j(\theta) = \Phi_N(\beta_j)$ and $d_j(\theta) = \Phi_N(\beta_j - \gamma_j)$, respectively. Moreover, we use $\phi = [\phi_{11},\ \phi_{00},\ \phi_{10}]'$ to denote the non-redundant reduced form probabilities of observing a monopoly, no entry, or the entry of Firm 1. It can be verified that the relationship between the reduced form probabilities ϕ and the structural parameters θ is given by the following set of equalities and inequalities which defines the identified set $\Theta(\phi)$.

$$\phi_{11} = d_1(\theta)d_2(\theta) \tag{29}$$

$$\phi_{00} = (1 - m_1(\theta))(1 - m_2(\theta)) \tag{30}$$

$$\phi_{10} \leq m_1(\theta)(1 - d_2(\theta)) \tag{31}$$

$$\phi_{10} \geq m_1(\theta)(1 - m_2(\theta)) + d_1(\theta)(m_2(\theta) - d_2(\theta)). \tag{32}$$

The relationships between the reduced form probabilities ϕ_{11} and ϕ_{00} and the structural parameter θ imply that β_2 and γ_2 are uniquely determined conditional on ϕ, β_1, and γ_1. Thus, the identified set lies in a 2-dimensional subspace of $\mathbb{R}^4$, which resembles the case discussed in Remark 3 in Section 3.2.

The remainder of this section is organized as follows. In Section 4.1 we characterize the identified set in terms of a functional relationship between ϕ, θ, and an auxiliary parameter and describe some of its properties. In Section 4.2 we describe the priors that are used for Bayesian inference as well as the objective function that is used for the construction of frequentist confidence sets. The numerical results are presented in Section 4.3. We generate two samples, one of size $n = 100$ and one of size $n = 1,000$ and compute 60% and 90% credible and confidence sets as well as estimates of the identified set. Computational details are relegated to a Technical Appendix.

4.1 Identified Set Correspondence and Data Generation

The correspondence between θ and ϕ can be represented by the relationship $\phi = \tilde{G}(\theta, \psi)$ where $\psi \in [0,1]$ and the two inequalities (31) and (32) are replaced by (omitting the θ-arguments):

$$\phi_{10} = m_1(1 - m_2) + d_1(m_2 - d_2) + \psi(m_1 - d_1)(m_2 - d_2). \tag{33}$$

The second term, which is pre-multiplied by ψ, is the probability that both firms are profitable as monopolists but not as duopolists. Consequentially, the auxiliary parameter ψ can be interpreted as the probability of a sunspot shock that selects Firm 1 if the Nash equilibrium is not unique. This representation is convenient for data generation. Parameter values of our data generating process for θ and ψ are reported in the second column of Table 2. According to our parameterization, Firm 1 is slightly more profitable than Firm 2, that is $\beta_1 > \beta_2$, and more likely to enter a market that can only sustain a single monopolist, $\psi > 0.5$. Using (29), (30), and (33) we calculated the associated reduced form probabilities

ϕ, which are also reported in the table. The probabilities of a a Firm 1 monopoly, and Firm 2 monopoly, and a duopoly are 48%, 33%, and 12%, respectively.

In order to implement frequentist inference the following characterization of the relationship between ϕ and θ in (33) is more convenient:

$$\phi_{10} = m_1(1-d_2) - (1-\psi)(m_1-d_1)(m_2-d_2) = m_1(1-d_2) - \alpha, \tag{34}$$

where $0 \leq \alpha \leq \bar{\alpha}(\theta)$ and $\bar{\alpha}(\theta) = (m_1-d_1)(m_2-d_2)$. In turn, we can define the non-negative function

$$Q(\theta;\phi) = \min_{0\leq\alpha\leq\bar{\alpha}(\theta)} \big\|\phi - G(\theta,\alpha)\big\|, \tag{35}$$

where

$$G(\theta,\alpha) = \big[d_1d_2,\ (1-m_1)(1-m_2),\ m_1(1-d_2)\big]' - \big[0,\ 0,\ \alpha\big]'.$$

It is straightforward to verify that $\theta \in \Theta(\phi)$ if and only if $Q(\theta;\phi) = 0$. The re-parameterization in terms of α makes the relationship between ϕ and θ additively separable with respect to the auxiliary parameter and simplifies the evaluation of $Q(\theta;\phi)$. Since the identified set $\Theta(\phi)$ is 2-dimensional and it is challenging to graphically depict credible and confidence sets in more than two dimensions, we will restrict inference to the sub-vector θ_1 that contains the profit function parameters for Firm 1: $\theta_1 = [\beta_1, \gamma_1]'$. The identified set of θ_1 is given by the projection $\Theta_1(\phi)$ defined as

$$\Theta_1(\phi) = \left\{ \theta_1 \,\middle|\, \exists\, \theta_2 \text{ s.t. } Q([\theta_1', \theta_2']';\phi) = 0 \right\}. \tag{36}$$

4.2 Bayesian and Frequentist Inference

Bayesian inference requires the specification of a prior distribution. We consider two priors specified on the (θ,ψ) space as well as one prior directly specified on the (θ_1,ϕ) space. The third column of Table 2 provides information on Priors 1 and 2. Both priors are based on the same distribution for θ but differ with respect to the distribution of ψ. The priors are constructed as products of marginal distributions:

$$\text{Priors 1, 2:} \qquad p_{(\theta,\psi)}(\theta,\psi) = p(\beta_1)p(\gamma_1)p(\beta_2)p(\gamma_2)p(\psi). \tag{37}$$

We use fairly diffuse Gaussian priors for the elements of the θ vector. The distributions of γ_1 and γ_2 are truncated at zero to ensure that duopoly profits are less than monopoly profits. The auxiliary parameter ψ has support on the unit interval. Under Prior 1 the mean of ψ, which can be interpreted as an equilibrium selection probability, is 0.5. Under Prior 2 the mean of ψ is set to 0.8, which reflects the belief that it is more likely that Firm 1 enters if a monopoly of either firm is profitable. By evaluating the function $\tilde{G}(\theta, \psi)$ at random draws from the prior distribution of (θ, ψ) we obtain draws from the prior distribution of ϕ. Means and standard deviation are reported in the last four rows of Table 2 for Prior 1.

The theoretical analysis in Section 3 highlighted that a crucial ingredient of Bayesian inference in partially identified models is the conditional prior distribution of $p(\theta|\phi)$. Since Priors 1 and 2 are specified on the (θ, ψ) space, the shape of this conditional distribution is not directly evident. We can use (29), (30), and (33) to construct a one-to-one mapping f between (θ_1, ϕ) and (θ, ψ) such that

$$p_{(\theta_1)}(\theta_1|\phi) \propto p_{(\theta_1,\phi)}(\theta_1, \phi) = p_{(\theta,\psi)}(f(\theta_1, \phi)) \cdot |f^{(1)}((\theta_1, \phi))|. \tag{38}$$

The change of variables introduces the Jacobian term $|f^{(1)}((\theta_1, \phi))|$. In addition to Priors 1 and 2 we construct a Prior 3 that is flat with respect to ϕ and uniform on the identified set $\Theta_1(\phi)$ conditional on ϕ. This prior is directly specified on (θ_1, ϕ) space:

$$\text{Prior 3:} \qquad p_{(\theta_1,\phi)}(\theta_1, \phi) \propto \frac{I\{\theta_1 \in \Theta_1(\phi)\}}{\int I\{\theta_1 \in \Theta_1(\phi)\} d\theta_1}. \tag{39}$$

To obtain the posterior distributions, the three prior distributions are combined with the likelihood function $p(Y^n|\phi) = p(Y^n|\tilde{G}(\theta, \psi))$. Draws from the posterior associated with Priors 1 and 2 are generated with a Random Walk Metropolis Algorithm that is described in many textbooks, e.g. Geweke (2005). Posterior draws under Prior 3 are obtained through direct sampling.

Since we stressed the numerical differences between Bayesian credible sets and frequentist confidence sets in Sections 2 and 3, we also construct confidence sets. To do so, we use the minimum-distance approach described in Section 3.4. Let $\hat{\phi}_n$ be the maximum likelihood estimator of ϕ, which has the property that $\sqrt{n}(\hat{\phi}_n - \phi) \Longrightarrow N(0, \Lambda)$. Let $\hat{\Lambda}_n$ be a consistent estimator of Λ and define a sample analogue of $Q(\theta; \phi)$ in (35) as

$$Q_n(\theta; \hat{\phi}_n) = \min_{0 \leq \alpha \leq \bar{\alpha}(\theta)} n \big\| \hat{\phi}_n - G(\theta, \alpha) \big\|_{\hat{\Lambda}_n^{-1}}. \tag{40}$$

For the numerical illustration we use for simplicity a fixed critical value and the projection method to obtain a confidence set for θ_1

$$CS_F^{\theta_1}(Y^n) = \left\{ \theta_1 \;\middle|\; \exists\, \theta_2 \text{ s.t } Q_n([\theta_1', \theta_2']'; \hat{\phi}_n) \leq c_\tau^2 \right\}. \tag{41}$$

It can be shown that a fixed critical value that delivers uniformly valid confidence sets over the entire parameter space is given by the $1-\tau$ quantile of a $\chi^2(df = 3)$ distribution with three degrees of freedom. Three degrees of freedom arise because as (one of) the interaction parameters γ_1 and γ_2 approach zero, both inequalities (31) and (32) become binding. Since the equalities (29) and (30) always need to hold, it can also be shown that any θ-dependent critical value $c_\tau^2(\theta)$ never falls below the $1-\tau$ quantile of a $\chi^2(df = 2)$ distribution. We will use this insight to obtain a lower bound on the frequentist confidence set.

4.3 Numerical Results

In our entry game the maximum likelihood estimator $\hat{\phi}_n$ combined with the sample size n provides a sufficient statistic. To compute the credible and confidence sets we set $\hat{\phi}_n = \phi_0$, where ϕ_0 is given in the second column of Table 2. We consider the sample sizes $n = 100$ and $n = 1,000$ and set $\tau = 0.4$ and 0.1, respectively. Results for $n = 100$ are depicted in Figure 2. The solid ellipsoid-like contour in the six panels indicates the identified set $\Theta_1(\phi)$ conditional on $\phi = \hat{\phi}_n = \phi_0$. The dashed contours in the top panels depict credible sets conditional on $\hat{\phi}_n = \phi_0$ under Prior 1 and Prior 2. The plots highlight that this conditional distribution is not uniform and that it is sensitive to beliefs about the structural parameters θ and ψ. In case of Prior 2 most of the mass concentrates near the upper edge of the identified set. The posterior confidence sets obtained under Priors 1 and 2 are dominated by the uncertainty about the reduced form parameter ϕ, though the shape of $p(\theta_1|\phi)$ has some noticeable effect on inference as well. The bottom left panel shows posterior credible sets under Prior 3, which is uniform on $\Theta_1(\phi)$ conditional on ϕ.

Finally, the frequentist confidence sets are shown in the bottom right panel. Most strikingly, the confidence sets are substantially larger than the credible sets. As discussed above, the frequentist confidence sets are conservative because they are constructed with critical values from a $\chi^2(df = 3)$. These critical values are $c^2_{\tau=0.4} = 2.95$ and $c^2_{\tau=0.1} = 6.25$. The 10% critical values for a $\chi^2(df = 2)$ is 4.61. This implies that a 90% confidence set constructed from the same objective function with θ-dependent critical values cannot be smaller than the 60% confidence set depicted in the figure.

In Figure 3 we show the same credible and confidence sets but now computed based on a sample of $n = 1,000$ observations. As predicted by the large sample theory presented in Section 3 the posterior credible sets closely resemble the credible sets associated with the conditional prior distribution $p(\theta_1|\phi)$. Under Prior 3, the remaining uncertainty about ϕ induces some curvature into the posterior density of θ_1 and the credible sets mimic the contours of the identified set. The increased sample size tightens the frequentist confidence set, but by construction it extends beyond the boundaries of $\Theta_1(\hat{\phi}_n)$ for any sample size.

Remark 1: The entry game is incompletely specified in the sense that it is silent about the equilibrium selection in case both a Firm 1 and Firm 2 monopoly are profitable, but a duopoly is not. The relationship between structural and reduced form parameters can be characterized through the function $\phi = \tilde{G}(\theta, \psi)$, where ψ can be interpreted as the probability that Firm 1 enters the market in case the model does not have a unique pure strategy Nash equilibrium. Since it is an inherent feature of Bayesian analysis to place probability distributions over unknown parameters, the specification of a prior entails placing a probability distribution on the likelihood ψ that Firm 1 instead of Firm 2 enters the market. In Priors 1 and 2 the distribution of ψ is explicitly specified. In our illustration the distribution of ψ is independent of the profit function parameters θ, but it does not have to be. Bajari, Hong, and Ryan (2009), for instance, consider more sophisticated probabilistic equilibrium selection mechanisms. Prior 3 masks the distribution of the equilibrium selection probability because it is specified on (θ_1, ϕ) space. However, a change of variables can easily recover the implicit beliefs about ψ that lead to a prior that is uniform on $\Theta_1(\phi)$ conditional on ϕ. ■

Remark 2: Our illustration is restricted to the case of pure strategy Nash equilibria. The extension of the Bayesian analysis to mixed strategies is conceptually straightforward but computationally more involved. Suppose that $\beta_1 = \beta_2 = 0$. Whenever $[\epsilon_{i,1}, \epsilon_{i,2}] \in [0, \gamma_1] \otimes [0, \gamma_2]$ one has to consider three possibilities: a Firm 1 monopoly, a Firm 2 monopoly, or a mixed strategy equilibrium in which Firms 1 and 2 enter with probabilities $\epsilon_{i,2}/\gamma_2$ and $\epsilon_{i,1}/\gamma_1$, respectively. Linking reduced form parameters and structural parameters now requires a two-dimensional auxiliary parameter ψ that summarizes the selection probabilities for the three possible Nash equilibria. Under the assumption that $\epsilon_{i,j} \sim iidN(0,1)$ the evaluation of the function $\phi = \tilde{G}(\theta, \phi)$ involves the evaluation of normal cdfs (as in the case of pure strategy equilibria) and the calculation of means of truncated normal random variates. Methods to compute the identified set are discussed in Beresteanu, Molchanov, and Molinari (2009). ■

Remark 3: We excluded covariates from the firms' profit function. The introduction of covariates X leads to covariate-specific reduced form and auxiliary parameters, that is, ϕ would have to be replaced by $\phi(X)$ and ψ by $\psi(X)$. If the regressor space is discretized, as it is in many applications in industrial organization, the reduced-form parameter vector remains finite-dimensional and the large sample analysis in Section 3 remains applicable. ■

5 Conclusion

In regular identified models credible sets and confidence sets are numerically approximately identical if the sample size is large, despite different probabilistic underpinnings. We derived a large sample approximation for the posterior distribution of a structural parameter vector in a partially identified model that is characterized by a finite-dimensional vector of reduced form parameters to compare Bayesian credible sets and frequentist confidence sets.

Frequentist confidence sets have to extend beyond the boundaries of the estimated identified set, whereas Bayesian credible sets tend to be located in the interior of the identified set asymptotically. Thus, in large samples frequentist sets are too large from a Bayesian perspective and Bayesian sets are too small from a frequentist perspective. The Bayesian approach induces a probability distribution on the identified set $\Theta(\phi)$ conditional on ϕ. This distribution is not updated through the likelihood function and creates a challenge for the reporting of Bayesian inference. In this regard it is important to report estimates of the identified set $\Theta(\hat{\phi}_n)$ as well as the conditional prior $P^{\theta}_{\hat{\phi}_n}$ along with Bayesian posteriors so that the audience can assess whether due to the choice of prior the posterior concentrates in a small subset of the identified set. In the context of a two-player entry game we were able to conduct inference based on a reference prior that is conditionally uniformly distributed on the (compact) identified set $\Theta(\phi)$. Such a prior might have appeal to a broad audience.

References

Andrews, Donald and Patrik Guggenberger (2009): "Validity of Subsampling and 'Plug-in Asymptotics' Inference for Parameters Defined by Moment Inequalities," *Econometric Theory*, **25**, 669-709.

Andrews, Donald and Gustavo Soares (2010): "Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection," *Econometrica*, **78**, 119-158.

Bajari, Patrick, Lanier Benkard, and Jonathan Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, **75**, 1331-1370.

Bajari, Patrick, Han Hong, and Stephen Ryan (2009): "Identification and Estimation of a Discrete Game of Complete Information," *Econometrica*, forthcoming.

Beresteanu, Arie, Ilya Molchanov, and Francesca Molinari (2009): "Sharp Identification Regions in Models with Convex Predictions: Games, Individual Choice, and Incomplete Data," *CEMMAP Working Paper*, **CWP27/09**, Institute for Fiscal Studies, London.

Bernstein, S.N. (1934): *Theory of Probability*, Gostekhizdat, Moscow.

Berry, Steven (1994): "Estimating Discrete Choice Models of Product Differentiation," *RAND Journal of Economics*, **25**, 242-262.

Bresnahan, Timothy and Peter Reiss (1991): "Empirical Models of Discrete Games," *Journal of Econometrics*, **48**, 57-81.

Bollinger, Christopher and Martijn van Hasselt (2008): "A Bayesian Analysis of Binary Misclassification: Inference in Partially Identified Models," *Manuscript*, University of Kentucky and University of Western Ontario.

Bugni, Federico (2010): "Bootstrap Inference in Partially Identified Models Defined by Moment Inequalities: Coverage of the Identified Set," *Econometrica*, **78**, 735-754.

Canova, Fabio and Gianni De Nicolo (2002): "Monetary Disturbances Matter for Business Cycle Fluctuations in the G-7," *Journal of Monetary Economics*, **49**, 1131-59.

Chernozhukov, Victor, Han Hong, and Elie Tamer (2007): "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, **75**, 1243-1284.

Ciliberto, Federico and Elie Tamer (2009): "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, **77**, 1791-1828.

Geweke, John (2005): *Contemporary Bayesian Econometrics and Statistics*, John Wiley & Sons, Hoboken.

Ghosh, Jayanta K., Subhashia Ghosal, and Tapas Samanta (1994): "Stability and Convergence of the Posterior in Non-Regular Problems," in S.S. Gupta and J.O. Berger (eds.) *Statistical Decision Theory and Related Topics V*, 183-199, Springer Verlag, New York.

Ghosh, Jayanta K., Subhashia Ghosal, and Tapas Samanta (1995): "On the Convergence of Posterior Distributions," *Annals of Statistics*, **23**(6), 2145-2152.

Ghosh, Jayanta K and R.V. Ramamoorthi (2003): "Bayesian Nonparametrics," *Springer Verlag*, New York.

Haile, Phil and Elie Tamer (2003): "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*, **111**, 1-51.

Hartigan, John (1983): *Bayes Theory*, Springer Verlag, New York.

Imbens, Guido and Charles Manski (2004): "Confidence Intervals for Partially Identified Parameters," *Econometrica*, **72**, 1845-57.

Johnson, Richard, (1970): "Asymptotic Expansions of Associated with Posterior Distributions," *Annals of Mathematical Statistics*, **41**, 851-864.

Kadane, Joseph B. (1974): "The Role of Identification in Bayesian Theory," in S.E. Fienberg and A. Zellner (eds.) *Studies in Bayesian Econometrics and Statistics*, 175-191. North Holland, Amsterdam.

Kim, Jae-Young (1998): "Large Sample Properties of Posterior Densities, Bayesian Information Criterion and the Likelihood Principle in Nonstationary Time Series Models,"*Econometrica*, **66**, 359-380.

LeCam, Lucien (1953): "On Some Asymptotic Properties of Maximum Likelihood Estimates and Related Bayes Estimates," *University of California Publications in Statistics*, **1**, 277-330.

Liao, Yuan and Wenxin Jiang (2010): "Bayesian Analysis in Moment Inequality Models," *Annals of Statistics*, **38**, 275-316.

Lubik, Thomas and Frank Schorfheide (2004): "Testing for Indeterminacy: An Application to U.S. Monetary Policy," *American Economic Review*, **94**, 190-217.

Manski, Charles and Elie Tamer (2002): "Inference on Regressions with Interval Data on a Regressor or Outcome," *Econometrica*, **70**, 519-547.

Manski, Charles (2003): *Partial Identification of Probability Distributions,* New York, Springer Verlag.

Moon, Hyungsik Roger and Frank Schorfheide (2009): "Bayesian and Frequentist Inference in Partially Identified Models," *NBER Working Paper*, **14882**.

Moon, Hyungsik Roger, Frank Schorfheide, Eleonora Granziera, and Mihye Lee (2009): "Inference for VARs Identified with Sign Restrictions," *Manuscript.* University of Southern California and University of Pennsylvania.

Phillips, Peter C.B. and Werner Ploberger (1996): "An Asymptotic Theory of Bayesian Inference for Time Series," *Econometrica*, **64**, 381-412.

Poirier, Dale (1998): "Revising Beliefs in Nonidentified Models," *Econometric Theory*, **14**, 483-509.

Romano, Joseph P. and Azeem M. Shaikh (2010): "Inference for Partially Identified Set in Partially Identified Econometric Models," *Econometrica*, **78**, 169-212.

Rosen, Adam (2008): "Confidence Sets for Partially Identified Parameters that Satisfy a Finite Number of Moment Inequalities," *Journal of Econometrics*, **146**, 107-117.

Severini, Thomas (1991): "On the Relationship between Bayesian and Non-Bayesian Interval Estimates," *Journal of the Royal Statistical Society. Series B*, **53** 611-618.

Sims, Christopher and Harald Uhlig (1991): "Understanding Unit Rooters: A Helicopter Tour," *Econometrica*, **59**, 1591-1599.

Stoye, Jörg (2009): "More on Confidence Intervals for Partially Identified Parameters," *Econometrica*, **77**, 1299-1315.

Uhlig, Harald (2005): "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, **52**, 381-419.

van der Vaart, Aad (1998): *Asymptotic Statistics*, Cambridge University Press, Cambridge.

von Mises, Richard (1965): *Mathematical Theory of Probability and Statistics,* Academic Press, New York.

Figure 1: Inference in the Inequality Condition Model, Known Length

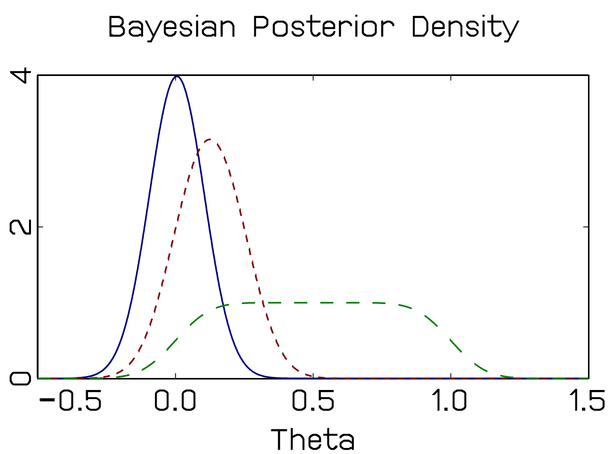

Notes: The figure depicts posterior pdfs of θ for λ equal to 0.01, 0.25, and 1.00. The sample size is $n = 100$ and $\hat{\phi}_n = 0$.

Table 1: 80% Frequentist Confidence Intervals and Credible Sets

		Length of $\Theta(\phi)$			
n		0.01	0.25	1.00	10
10	$P^{\theta}_{Y^n}\{\theta \in CS^{\theta}_F\}$	0.80	0.82	0.93	0.99
	Length of CS^{θ}_F	0.81	0.87	1.53	10.5
	Length of CS^{θ}_B	0.81	0.83	1.11	8.00
50	$P^{\theta}_{Y^n}\{\theta \in CS^{\theta}_F\}$	0.80	0.88	0.97	1.00
	Length of CS^{θ}_F	0.36	0.49	1.24	10.2
	Length of CS^{θ}_B	0.36	0.41	0.86	8.00
100	$P^{\theta}_{Y^n}\{\theta \in CS^{\theta}_F\}$	0.80	0.91	0.98	1.00
	Length of CS^{θ}_F	0.26	0.42	1.17	10.2
	Length of CS^{θ}_B	0.26	0.32	0.82	8.00
500	$P^{\theta}_{Y^n}\{\theta \in CS^{\theta}_F\}$	0.80	0.96	0.99	1.00
	Length of CS^{θ}_F	0.12	0.33	1.08	10.1
	Length of CS^{θ}_B	0.12	0.22	0.80	8.00

Notes: CS^{θ}_F is the frequentist confidence interval, CS^{θ}_B is the Bayesian credible interval, and $P^{\theta}_{Y^n}\{\theta \in CS^{\theta}_F\}$ is the posterior probability that θ lies in the frequentist confidence interval.

Table 2: Entry Game: "True" Parameters and Prior

Parameter	True Value	Prior Distribution
	Structural Parameters θ	
β_1	0.7	$N(0, 4^2)$
γ_1	1.0	$N_+(0, 4^2)$
β_2	0.5	$N(0, 4^2)$
γ_2	1.0	$N_+(0, 4^2)$
	Auxiliary Parameter ψ	
ψ	0.7	Prior 1: $\mathcal{B}(0.5, 0.2^2)$
	0.7	Prior 2: $\mathcal{B}(0.8, 0.1^2)$
	Implied Reduced Form Parameters ϕ (Prior 1)	
ϕ_{11}	0.12	$\mu_{11} = 0.13$, $\sigma_{11} = 0.28$
ϕ_{00}	0.07	$\mu_{00} = 0.25$, $\sigma_{00} = 0.37$
ϕ_{10}	0.48	$\mu_{10} = 0.31$, $\sigma_{10} = 0.40$

Notes: for the prior distribution of the reduced form parameters we report means μ and standard deviations σ under $\alpha \sim \mathcal{B}(0.5, 0.2^2)$. $N(\nu, \sigma^2)$ and $\mathcal{B}(\mu, \sigma^2)$ refer to normal and Beta distributions with mean μ and variance σ^2.

Figure 2: Posterior Credible Sets and Frequentist Confidence Sets, $n = 100$

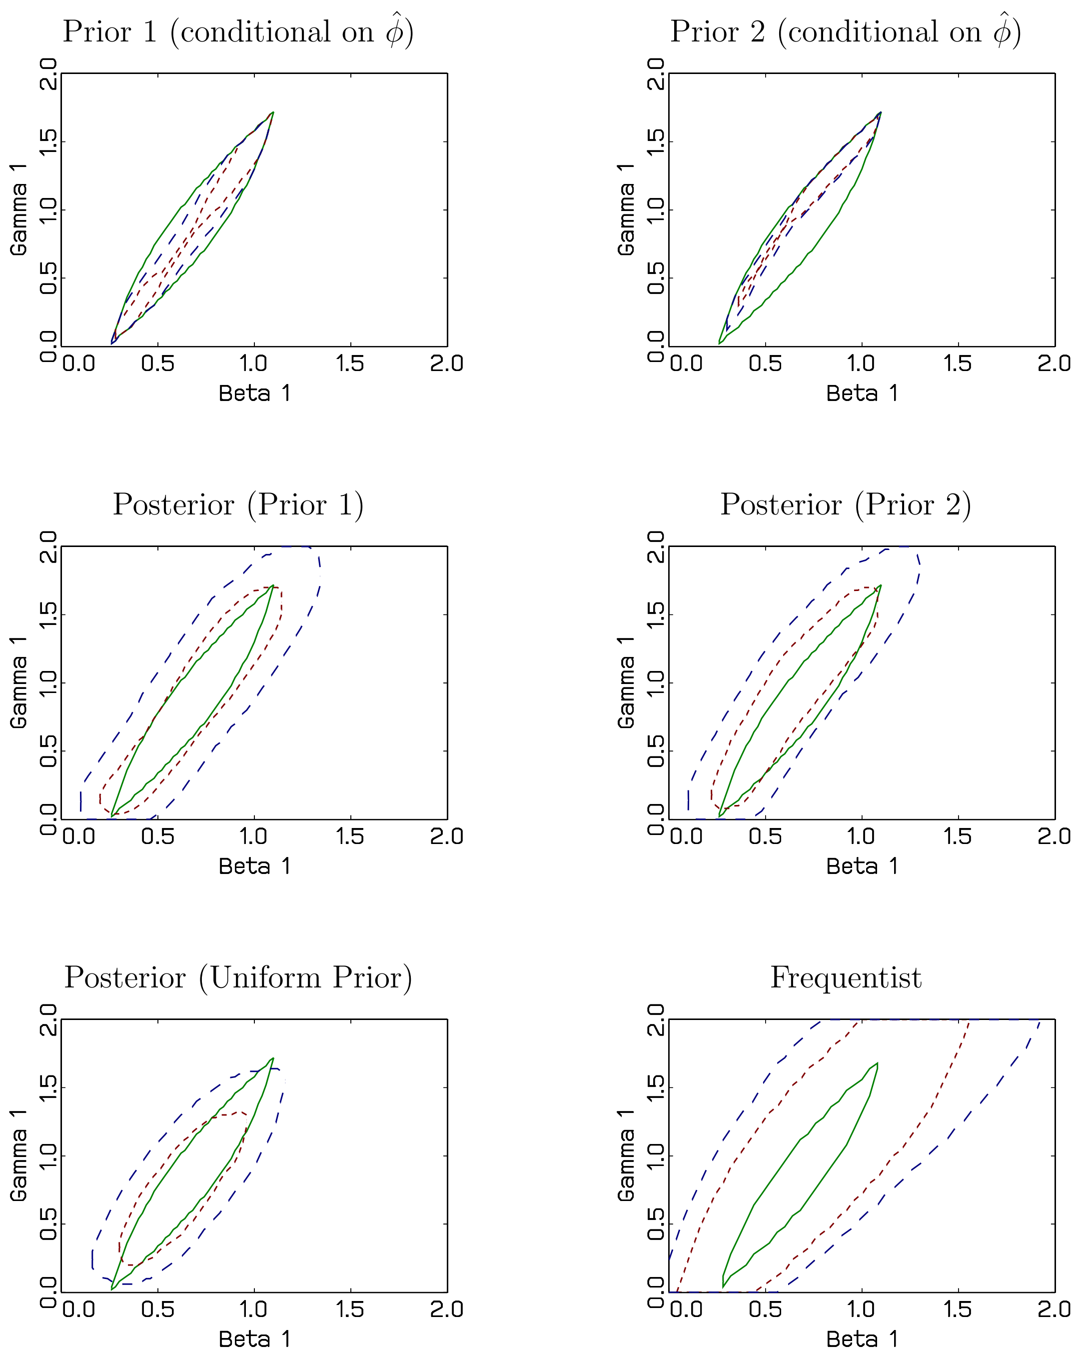

Notes: Figure depicts identified sets (solid), 90% credible (confidence) sets (long dashes), and 60% credible (confidence) sets (short dashes).

Figure 3: Posterior Credible Sets and Frequentist Confidence Sets, $n = 1,000$

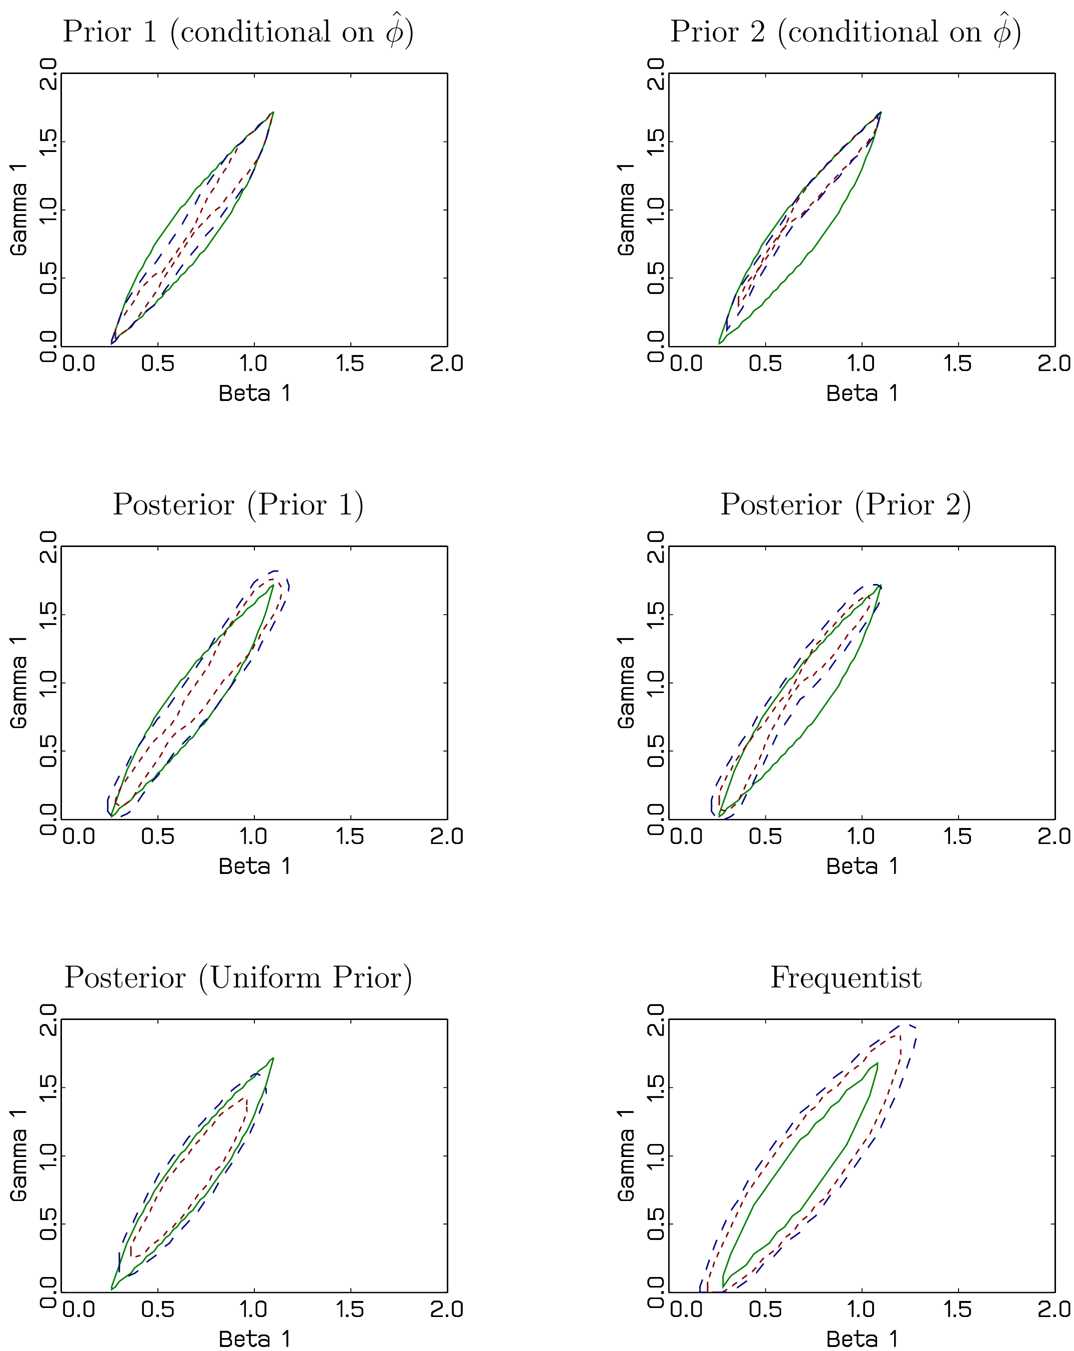

Notes: Figure depicts identified sets (solid), 90% credible (confidence) sets (long dashes), and 60% credible (confidence) sets (short dashes).

A Derivations for Section 2

Direct calculation of the posterior density of θ:

$$\begin{aligned} p(\theta|Y^n) &= \frac{1}{\sqrt{2\pi/n}} \int_{-\infty}^{\infty} \frac{1}{\lambda} I\{\phi \leq \theta \leq \phi + \lambda\} \exp\left\{-\frac{n}{2}(\phi - \hat{\phi}_n)^2\right\} d\phi \\ &= \frac{1}{\lambda}\frac{1}{\sqrt{2\pi}} \int_{\sqrt{n}(\theta-\hat{\phi}_n-\lambda)}^{\sqrt{n}(\theta-\hat{\phi}_n)} \exp\left\{-\frac{s^2}{2}\right\} ds \\ &= \frac{1}{\lambda}\left[\Phi_N(\sqrt{n}(\theta - \hat{\phi}_n)) - \Phi_N(\sqrt{n}(\theta - \hat{\phi}_n - \lambda))\right]. \end{aligned}$$

The second equality follows from re-arranging the inequalities in the indicator function and the change of variables $s = \sqrt{n}(\phi - \hat{\phi}_n)$. It is straightforward to verify that $p(\theta|Y^n)$ is has a single mode at $\theta = \hat{\phi}_n + \lambda/2$ and is symmetric around the mode. ■

B Proofs and Derivations for Section 3

Proof of Theorem 1(ii): Since the L_1 distance satisfies the triangle inequality

$$\|P^{\theta}_{Y^n} - P^{\theta}_{\hat{\phi}_n}\| \leq \|P^{\theta}_{Y^n} - P^{\theta}_{N,Y^n}\| + \|P^{\theta}_{N,Y^n} - P^{\theta}_{\hat{\phi}_n}\|$$

it suffices to show that $\|P^\theta_{N,Y^n} - P^\theta_{\hat\phi_n}\| \stackrel{\mathbb{P}}{\longrightarrow} 0$. For any function $f(\theta)$ with $|f(\theta)| \le 1$ we can use the bounds

$$
\begin{aligned}
&\left| \int f(\theta) dP^\theta_{N,Y^n} - \int f(\theta) dP^\theta_{\hat\phi_n} \right| \\
\le \;& \int_{\mathbb{R}^m} \left| \int f(\theta) dP^\theta_{\hat\phi_n + \hat J_n^{-1/2} D_n^{-1} s} - \int f(\theta) dP^\theta_{\hat\phi_n} \right| dN\big(0,I\big)(s) \\
\le \;& \int_{\mathbb{R}^m} \left\| P^\theta_{\hat\phi_n + \hat J_n^{-1/2} D_n^{-1} s} - P^\theta_{\hat\phi_n} \right\| dN\big(0,I\big)(s) \\
\le \;& \int_{\mathbb{R}^m} I\{\|\hat\phi_n - \phi_0\| < \delta\} I\{\|\hat\phi_n - \phi_0 + \hat J_n^{-1/2} D_n^{-1} s\| < \delta\} \left\| P^\theta_{\hat\phi_n + \hat J_n^{-1/2} D_n^{-1} s} - P^\theta_{\hat\phi_n} \right\| dN\big(0,I\big)(s) \\
& + 2 I\{\|\hat\phi_n - \phi_0\| \ge \delta\} + 2 \int_{\mathbb{R}^m} I\{\|\hat\phi_n - \phi_0 + \hat J_n^{-1/2} D_n^{-1} s\| \ge \delta\} dN\big(0,I\big)(s) \\
\le \;& \int_{\mathbb{R}^m} M(\phi_0, \delta) \|\hat J_n^{-1/2} D_n^{-1} s\| dN\big(0,I\big)(s) + 2 I\{\|\hat\phi_n - \phi_0\| \ge \delta\} \\
& + 2 I\{\|\hat\phi_n - \phi_0\| \ge \delta/2\} + 2 \int_{\mathbb{R}^m} I\{\|\hat J_n^{-1/2} D_n^{-1} s\| \ge \delta/2\} dN\big(0,I\big)(s) \\
\le \;& M(\phi_0, \delta) \|\hat J_n^{-1/2}\| \|D_n^{-1}\| \int_{\mathbb{R}^m} \|s\| dN\big(0,I\big)(s) + o_p(1) \stackrel{\mathbb{P}}{\longrightarrow} 0.
\end{aligned}
$$

For the third inequality we bound the L_1 distance $\|P^\theta_{\hat\phi_n + \hat J_n^{-1/2} D_n^{-1} s} - P^\theta_{\hat\phi_n}\|$ by 2 if either $\hat\phi_n$ or $\hat\phi_n + \hat J_n^{-1/2} D_n^{-1} s$ lie outside of the $N_\delta(\phi_0)$ neighborhood. For the fourth inequality we use the Lipschitz bound of Assumption 2 and the inequality $I\{\|x+y\| \ge \delta\} \le I\{\|x\| \ge \delta/2\} + I\{\|y\| \ge \delta/2\}$. The last line follows from Assumption 1 that $\hat\phi_n$ converges in probability to ϕ_0 , $\|D_n\| \uparrow \infty$, and $\hat J_n^{-1/2} = O_p(1)$. A similar argument can be used to establish the convergence of $P^\theta_{Y^n}$ to $P^\theta_{\phi_0}$. ■

Direct Calculations to Verify Equation (22): We begin with the change of variable $s = \hat J_n^{1/2} D_n (\theta - \hat\phi_n + \tilde s)$, which leads to

$$
\begin{aligned}
p(\theta|Y^n) &= p_N(\theta|Y^n) \\
&= \frac{1}{\lambda_n} \int f\left(\frac{\theta - \hat\phi_n - \hat J_n^{-1/2} D_n^{-1} s}{\lambda_n} \right) \varphi_N(s) ds \\
&= \frac{1}{\lambda_n} |\hat J_n^{1/2} D_n| \int_{\tilde s = -\lambda_n}^{0} f(-\lambda_n^{-1} \tilde s) \varphi_N(\hat J_n^{1/2} D_n (\theta - \hat\phi_n + \tilde s)) d\tilde s.
\end{aligned}
$$

The second equality makes use of the assumption that $f(x) = 0$ outside of the unit interval.

The L_1 distance can be bounded as follows:

$$\begin{aligned}
&\int_\theta \left| p_N(\theta|Y^n) - |\hat{J}_n^{1/2} D_n| \varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n)) \right| d\theta \\
&= |\hat{J}_n^{1/2} D_n| \int_\theta \left| \int_{\tilde{s}=-\lambda_n}^0 \frac{1}{\lambda_n} f(-\lambda_n^{-1}\tilde{s}) \left[\varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n + \tilde{s})) - \varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n)) \right] d\tilde{s} \right| d\theta \\
&\leq |\hat{J}_n^{1/2} D_n| \int_{\tilde{s}=-\lambda_n}^0 \int_\theta \frac{1}{\lambda_n} f(-\lambda_n^{-1}\tilde{s}) \left| \varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n + \tilde{s})) - \varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n)) \right| d\theta d\tilde{s} \\
&\leq \int_{\tilde{s}=-\lambda_n}^0 \frac{1}{\lambda_n} f(-\lambda_n^{-1}\tilde{s}) \int_{\tilde{\theta}} \left| \varphi_N(\tilde{\theta} + \hat{J}_n^{1/2} D_n \tilde{s}) - \varphi_N(\tilde{\theta}) \right| d\tilde{\theta} d\tilde{s}. \qquad \text{(A.1)}
\end{aligned}$$

The first equality follows because $\int_0^1 f(x)dx = 1$ and $\varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n))$ does not depend on $\tilde{s}$. The last inequality is based on the change of variables $\tilde{\theta} = \hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n)$.

Now consider the difference $\varphi_N(\tilde{\theta} + h) - \varphi_N(\tilde{\theta})$ for $-\bar{h} \leq h \leq 0$. By direct calculation we obtain

$$\begin{aligned}
|\varphi_N(\tilde{\theta} + h) - \varphi_N(\tilde{\theta})| &= \left| (2\pi)^{-1/2} \exp\left\{ -\frac{1}{2}(\tilde{\theta} + h)^2 \right\} - \varphi_N(\tilde{\theta}) \right| \\
&= \left| \exp\left\{ -\frac{1}{2}(2\tilde{\theta}h + h^2) \right\} - 1 \right| \varphi_N(\tilde{\theta}).
\end{aligned}$$

A first-order Taylor series expansion around $h = 0$ yields

$$\exp\left\{ -\frac{1}{2}(2\tilde{\theta}h + h^2) \right\} - 1 = -(\tilde{\theta} + h_*(\tilde{\theta})) \exp\{-\tilde{\theta} h_*(\tilde{\theta})\} \exp\{-h_*^2(\tilde{\theta})/2\} h,$$

where $-\bar{h} \leq h_*(\tilde{\theta}) \leq 0$. Thus, on the interval $-\bar{h} \leq h \leq 0$ we obtain the bound

$$\left| \exp\left\{ -\frac{1}{2}(2\tilde{\theta}h + h^2) \right\} - 1 \right| \varphi_N(\tilde{\theta}) \leq (|\tilde{\theta}| + \bar{h}) \exp\{-\tilde{\theta}\bar{h} I\{\tilde{\theta} \leq 0\}\} \bar{h} \varphi_N(\tilde{\theta}). \qquad \text{(A.2)}$$

Replacing $\bar{h}$ by $\hat{J}_n^{1/2} D_n \lambda_n$ in (A.2) and combining (A.1) with (A.2) leads to

$$\begin{aligned}
&\int_\theta \left| p_N(\theta|Y^n) - |\hat{J}_n^{1/2} D_n| \varphi_N(\hat{J}_n^{1/2} D_n(\theta - \hat{\phi}_n)) \right| d\theta \\
&\leq \hat{J}_n^{1/2} D_n \lambda_n \int_{\tilde{\theta}} (|\tilde{\theta}| + \hat{J}_n^{1/2} D_n \lambda_n) \exp\{-\tilde{\theta} \hat{J}_n^{1/2} D_n \lambda_n I\{\tilde{\theta} \leq 0\}\} \varphi_N(\tilde{\theta}) d\tilde{\theta} = o_p(1).
\end{aligned}$$

The $o_p(1)$ statement follows because $D_n \lambda_n \longrightarrow 0$ and we can find a finite constant M and an N_M such that for $n > N_M$

$$\int_{\tilde{\theta}} (|\tilde{\theta}| + \hat{J}_n^{1/2} D_n \lambda_n) \exp\{-\tilde{\theta} \hat{J}_n^{1/2} D_n \lambda_n I\{\tilde{\theta} \leq 0\}\} \varphi_N(\tilde{\theta}) d\tilde{\theta} \leq M$$

with probability approaching one. ■

The following Lemma is needed for the subsequent proof of Theorem 2. To simplify the notation let $p_Y(\theta) = p(\theta|Y^n)$, and $p_0(\theta) = p(\theta|\phi_0)$. Similarly, we abbreviate the thresholds κ_{Y^n} and κ_{ϕ_0} by κ_Y and κ_0.

Lemma B.1 *Suppose that* $\int |p_Y(\theta) - p_0(\theta)|d\theta = o_p(1)$ *and* $\int I\{p_0(\theta) = \kappa_0\}p_0(\theta)d\theta = 0$, *where* $\kappa_0 < \infty$. *Then*

$$\int \bigg| I\{p_Y(\theta) \geq \kappa_0\} - I\{p_0(\theta) \geq \kappa_0\} \bigg| p_Y(\theta)d\theta = o_p(1).$$

Proof of Lemma B.1: Write

$$\begin{aligned} &\int \bigg| I\{p_Y(\theta) \geq \kappa_0\} - I\{p_0(\theta) \geq \kappa_0\} \bigg| p_Y(\theta)d\theta \\ &= \int \{\theta \mid p_Y(\theta) \geq \kappa_0,\ p_0(\theta) < \kappa_0\} p_Y(\theta)d\theta + \int \{\theta \mid p_Y(\theta) < \kappa_0,\ p_0(\theta) \geq \kappa_0\} p_Y(\theta)d\theta \\ &= \int_{\theta \in A_n} p_Y(\theta)d\theta + \int_{\theta \in B_n} p_Y(\theta)d\theta = I + II, \end{aligned}$$

say. We will subsequently construct $o_p(1)$ bounds for terms I and II.

Bound for I: We deduce from the L_1 convergence assumption of $p_Y(\theta)$ to $p_0(\theta)$ that

$$I = \int_{\theta \in A_n} p_Y(\theta)d\theta = \int_{\theta \in A_n} p_0(\theta)d\theta + o_p(1) = Ia + o_p(1).$$

Thus, it suffices construct an $o_p(1)$ bound for Ia. Define the function

$$f_n(\theta) = p_Y(\theta) - p_0(\theta)$$

and notice that $f_n(\theta) > 0$ for $\theta \in A_n$. With this definition,

$$\begin{aligned} \int_{A_n} f_n(\theta)p_0(\theta)d\theta &= \int_{A_n} |p_Y(\theta) - p_0(\theta)|p_0(\theta)d\theta \\ &\leq \kappa_0 \int_{A_n} |p_Y(\theta) - p_0(\theta)|d\theta = o_p(1). \end{aligned} \tag{A.3}$$

The inequality follows from $p_0(\theta) < \kappa_0$ on the set A_n. The $o_p(1)$ statement is a consequence of the assumptions that $p_Y(\theta)$ converges to $p_0(\theta)$ in L_1 and that κ_0 is finite.

Now notice that

$$I\{\theta \in A_n\} = I\Big\{ I\{\theta \in A_n\} f_n(\theta) > 0 \Big\}. \tag{A.4}$$

If $\theta \in A_n$ then $f_n(\theta) > 0$, which means that $I\{\theta \in A_n\} f_n(\theta) > 0$. Moreover, for any $\eta > 0$ we obtain the inequality

$$I\Big\{ I\{\theta \in A_n\} f_n(\theta) > \eta \Big\} \leq \frac{1}{\eta} I\{\theta \in A_n\} f_n(\theta). \tag{A.5}$$

Thus,

$$\begin{aligned}
Ia &= \int I\Big\{ I\{\theta \in A_n\} f_n(\theta) > 0 \Big\} p_0(\theta) d\theta \\
&\leq \int I\Big\{ I\{\theta \in A_n\} f_n(\theta) > 0 \Big\} p_0(\theta) d\theta - \int I\Big\{ I\{\theta \in A_n\} f_n(\theta) > \eta \Big\} p_0(\theta) d\theta \\
&\quad + \frac{1}{\eta} \int_{A_n} f_n(\theta) p_0(\theta) d\theta \\
&= \int I\Big\{ 0 < I\{\theta \in A_n\} f_n(\theta) \leq \eta \Big\} p_0(\theta) d\theta + \frac{1}{\eta} \int_{A_n} f_n(\theta) p_0(\theta) d\theta \\
&= Ib + Ic,
\end{aligned}$$

say. The first equality follows from (A.4). The inequality is a consequence of (A.5).

To bound Ib notice that

$$I\Big\{ 0 < I\{\theta \in A_n\} f_n(\theta) \leq \eta \Big\} \leq I\Big\{ \kappa_0 - \eta \leq p_0(\theta) \leq \kappa_0 + \eta \Big\}.$$

For the indicator function on the left-hand-side to be one, it has to be the case that $\theta \in A_n$ and $f_n(\theta) \leq \eta$. On the set A_n $p_Y(\theta) \geq \kappa_0$ which leads to

$$\kappa_0 \leq p_Y(\theta) = p_0(\theta) + f_n(\theta) \leq p_0(\theta) + \eta,$$

that is,

$$\kappa_0 - \eta \leq p_0(\theta).$$

Moreover, $p_0(\theta) < \kappa_0 \leq \kappa_0 + \eta$ and therefore the following inequality is satisfied:

$$\kappa_0 - \eta \leq p_0(\theta) \leq \kappa_0 + \eta.$$

Thus,

$$Ib \leq \int I\Big\{\kappa_0 - \eta \leq p_0(\theta) \leq \kappa_0 + \eta\Big\}p_0(\theta)d\theta.$$

Based on the Dominated Convergence Theorem and the assumption $\int I\{p_0(\theta) = \kappa_0\}p_0(\theta) = 0$ we deduce that

$$\lim_{\eta \longrightarrow 0} \int I\Big\{\kappa_0 - \eta \leq p_0(\theta) \leq \kappa_0 + \eta\Big\}p_0(\theta)d\theta = \int I\Big\{p_0(\theta) = \kappa_0\Big\}p_0(\theta) = 0. \qquad \text{(A.6)}$$

Notice that our bound for Ib is deterministic.

To establish that $Ia \stackrel{\mathbb{P}}{\longrightarrow} 0$ it suffices to show that for every $\epsilon > 0$ and $\delta > 0$ there exists an $N(\epsilon, \delta)$ such that for $n \geq N(\epsilon, \delta)$

$$\mathbb{P}\{Ia > \epsilon\} \leq \mathbb{P}\{Ib > \epsilon/2\} + \mathbb{P}\{Ic > \epsilon/2\} < \delta.$$

Based on (A.6) we can find an $\eta(\epsilon) > 0$ such that $\mathbb{P}\{Ib > \epsilon/2\} = 0$. To obtain a bound for Ic define $Z_n = \int_{A_n} f_n(\theta)p_0(\theta)d\theta$ such that $Ic = Z_n/\eta$. According to (A.3), $Z_n = o_p(1)$. Thus, we can find an $N(\epsilon, \delta)$ such that

$$\mathbb{P}\left\{|Z_n| > \eta(\epsilon)\frac{\epsilon}{2}\right\} < \delta$$

whenever $n \geq N(\epsilon, \delta)$, which shows that $Ia = o_p(1)$.

Bound for II: This bound can be obtained following the same steps. Change the definition of $f_n(\theta)$ to

$$f_n(\theta) = p_0(\theta) - p_Y(\theta).$$

Using this definition we obtain that

$$\begin{aligned}
\int_{\theta \in B_n} f_n(\theta)p_Y(\theta)d\theta &= \int_{\theta \in B_n} (p_0(\theta) - p_Y(\theta))p_Y(\theta)d\theta \\
&\leq \kappa_0 \int_{\theta \in B_n} \big|p_0(\theta) - p_Y(\theta)\big|d\theta = o_p(1)
\end{aligned}$$

because on the set B_n the density $p_Y(\theta)$ is bounded by κ_0. Now consider

$$\begin{aligned} II &= \int_{B_n} p_Y(\theta)d\theta = \int I\Big\{ I\{\theta \in B_n\} f_n(\theta) > 0 \Big\} p_Y(\theta)d\theta \\ &\leq \int I\Big\{ I\{\theta \in B_n\} f_n(\theta) > 0 \Big\} p_Y(\theta)d\theta - \int I\Big\{ I\{\theta \in B_n\} f_n(\theta) > \eta \Big\} p_Y(\theta)d\theta \\ &\quad + \frac{1}{\eta}\int_{B_n} f_n(\theta)p_Y(\theta)d\theta \\ &= \int I\Big\{ 0 < I\{\theta \in B_n\} f_n(\theta) \leq \eta \Big\} p_0(\theta)d\theta + \frac{1}{\eta}\int_{B_n} f_n(\theta)p_Y(\theta)d\theta + o_p(1) \\ &= IIb + IIc + o_p(1). \end{aligned}$$

In the last line we used the L_1 convergence to replace $p_Y(\theta)$ by $p_0(\theta)$ in the definition of term IIb which introduces an additional $o_p(1)$ term.

To bound IIb notice that

$$I\Big\{ 0 < I\{\theta \in B_n\} f_n(\theta) \leq \eta \Big\} \leq I\Big\{ \kappa_0 - \eta \leq p_n(\theta) \leq \kappa_0 + \eta \Big\}.$$

For the indicator function on the left-hand-side to be one, it has to be the case that $\theta \in B_n$ and $f_n(\theta) \leq \eta$. On the set B_n $p_Y(\theta) < \kappa_0$ which leads to

$$\kappa_0 > p_Y(\theta) = p_0(\theta) - f_n(\theta) \geq p_0(\theta) - \eta.$$

that is,

$$\kappa_0 + \eta \geq p_0(\theta).$$

Moreover, $p_0(\theta) \geq \kappa_0 \geq \kappa_0 - \eta$ and therefore the following inequality is satisfied:

$$\kappa_0 - \eta \leq p_0(\theta) \leq \kappa_0 + \eta.$$

Thus,

$$IIb \leq \int I\Big\{ \kappa_0 \leq p_0(\theta) < \kappa_0 + \eta \Big\} p_0(\theta)d\theta.$$

Dominated convergence implies that the bound converges to zero as $\eta \longrightarrow 0$. The remaining steps needed to establish that $II = o_p(1)$ are identical to the steps followed for term I. ■

Proof of Theorem 2: Part (i): To simplify the notation let $p_Y(\theta) = p(\theta|Y^n)$ and $p_0(\theta) = p(\theta|\phi_0)$. Similarly, we abbreviate the thresholds κ_{Y^n} and κ_{ϕ_0} by κ_Y and κ_0. Write

$$\begin{aligned}
&\int \Big| I\{p_Y(\theta) \geq \kappa_Y\} - I\{p_0(\theta) \geq \kappa_0\} \Big| p_Y(\theta) d\theta \\
&= \int \Big| I\{p_Y(\theta) \geq \kappa_Y\} - I\{p_Y(\theta) \geq \kappa_0\} \Big| p_Y(\theta) d\theta \\
&\quad + \int \Big| I\{p_Y(\theta) \geq \kappa_0\} - I\{p_0(\theta) \geq \kappa_0\} \Big| p_Y(\theta) d\theta \\
&= I + II,
\end{aligned}$$

say. In view of our assumptions Lemma B.1 provides an $o_p(1)$ bound for term II. Now consider term I. Since by construction

$$\int I\{p_Y(\theta) \geq \kappa_Y\} p_Y(\theta) d\theta = 1 - \tau,$$

we can write term I as

$$\begin{aligned}
I &= \int I\Big\{p_Y(\theta) \geq \min\{\kappa_0, \kappa_Y\}\Big\} p_Y(\theta) d\theta - \int I\Big\{p_Y(\theta) \geq \max\{\kappa_0, \kappa_Y\}\Big\} p_Y(\theta) d\theta \\
&= I\{\kappa_0 \geq \kappa_Y\} \left[(1-\tau) - \int I\{p_Y(\theta) \geq \kappa_0\} p_Y(\theta) d\theta \right] \\
&\quad + I\{\kappa_0 < \kappa_Y\} \left[\int I\{p_Y(\theta) \geq \kappa_0\} p_Y(\theta) d\theta - (1-\tau) \right] \\
&= \left| \int I\{p_Y(\theta) \geq \kappa_0\} p_Y(\theta) d\theta - (1-\tau) \right|.
\end{aligned}$$

In order to show that $I = o_p(1)$ we add and subtract $\int I\{p_0(\theta) \geq \kappa_0\} p_Y(\theta) d\theta$ and using the triangle inequality:

$$\begin{aligned}
I &\leq \left| \int I\{p_Y(\theta) \geq \kappa_0\} p_Y(\theta) d\theta - \int I\{p_0(\theta) \geq \kappa_0\} p_Y(\theta) d\theta \right| \\
&\quad + \left| \int I\{p_0(\theta) \geq \kappa_0\} p_Y(\theta) d\theta - (1-\tau) \right| \\
&= \left| \int I\{p_Y(\theta) \geq \kappa_0\} p_Y(\theta) d\theta - \int I\{p_0(\theta) \geq \kappa_0\} p_Y(\theta) d\theta \right| \\
&\quad + \left| \int I\{p_0(\theta) \geq \kappa_0\} p_Y(\theta) d\theta - \int I\{p_0(\theta) \geq \kappa_0\} p_0(\theta) d\theta \right| \\
&\leq \int \Big| I\{p_Y(\theta) \geq \kappa_0\} - I\{p_0(\theta) \geq \kappa_0\} \Big| p_Y(\theta) d\theta \\
&\quad + \int I\{p_0(\theta) \geq \kappa_0\} \big| p_Y(\theta) - p_0(\theta) \big| d\theta = o_p(1).
\end{aligned}$$

The first equality holds because $\int I\{p_0(\theta) \geq \kappa_0\} p_0(\theta) d\theta = 1 - \tau$. The final $o_p(1)$ result follows from Lemma B.1 and the L_1 convergence of the posterior densities established in Theorem 1.

Part (ii): The triangle inequality implies that

$$\|P^\theta_{\hat{\phi}_n} - P^\theta_{\phi_0}\| \leq \|P^\theta_{Y^n} - P^\theta_{\hat{\phi}_n}\| + \|P^\theta_{Y^n} - P^\theta_{\phi_0}\| \xrightarrow{\mathbb{P}} 0$$

by Theorem 1(ii). Let $p_n(\theta) = p(\theta|\hat{\phi}_n)$ and $\kappa_n = \kappa_{\hat{\phi}_n}$. Then using the same argument as for Part (i), replacing $p_Y(\theta)$ by $p_n(\theta)$ and κ_Y by κ_n we can easily establish that

$$\int \left| I\{\theta \in CS^\theta_{HPD}(\hat{\phi}_n)\} - I\{\theta \in CS^\theta_{HPD}(\phi_0)\} \right| dP^\theta_{Y^n} \xrightarrow{\mathbb{P}} 0. \tag{A.7}$$

Now consider the following inequality

$$\begin{aligned} \big|I\{\theta \in A\} - I\{\theta \in B\}\big| &\leq \big|I\{\theta \in A\} - I\{\theta \in C\}\big| + \big|I\{\theta \in B\} - I\{\theta \in C\}\big| \qquad \text{(A.8)} \\ &= I + II. \end{aligned}$$

If the left-hand side of (A.8) is zero, then the inequality is trivially satisfied. The left-hand side of (A.8) is one if $\theta \in A$ and $\theta \notin B$ or if $\theta \notin A$ and $\theta \in B$. Since the statement of the inequality is symmetric in A and B we focus on the first case. If $\theta \in A$, $\theta \notin B$, and $\theta \in C$, then $I = |1 - 1| = 0$ and $II = |0 - 1| = 1$. If $\theta \in A$, $\theta \notin B$, and $\theta \notin C$, then $I = |1 - 0| = 1$ and $II = |0 + 0| = 0$. We deduce that whenever the left-hand side of (A.8) is equal to one, the right-hand side is equal to one as well, which confirms the inequality. Now let

$$A = CS^\theta_{HPD}(Y^n), \quad B = CS^\theta_{HPD}(\hat{\phi}_n), \quad \text{and} \quad C = CS^\theta_{HPD}(\phi_0).$$

Integrating both sides of (A.8) yields

$$\begin{aligned} &\int \big|I\{\theta \in A\} - I\{\theta \in B\}\big| p_Y(\theta) d\theta \\ &\leq \int \big|I\{\theta \in A\} - I\{\theta \in C\}\big| p_Y(\theta) d\theta + \int \big|I\{\theta \in B\} - I\{\theta \in C\}\big| p_Y(\theta) d\theta = o_p(1). \end{aligned}$$

The $o_p(1)$ statement follows from Part (i) and (A.7). ■

Technical Appendix for Section 4 - Not Intended for Publication

The probabilities that firm i is profitable as monopolist and duopolist are

$$m_i = \Phi_N(\beta_i) \quad \text{and} \quad d_i = \Phi_N(\beta_i - \gamma_i). \tag{B.1}$$

The relationship between the reduced-form entry probabilities and m_i and d_i, $i = 1, 2$ is given by

$$\begin{aligned} \phi_{11} &= d_1 d_2 && \text{(B.2)} \\ \phi_{00} &= (1 - m_1)(1 - m_2) && \text{(B.3)} \\ \phi_{10} &= m_1(1 - m_2) + d_1(m_2 - d_2) + \psi(m_1 - d_1)(m_2 - d_2) && \text{(B.4)} \\ &= m_1(1 - d_2) - (1 - \psi)(m_1 - d_1)(m_2 - d_2), \end{aligned}$$

where $\psi \in [0, 1]$. The vector of non-redundant reduced form parameters is given by $\phi = [\phi_{11}, \phi_{00}, \phi_{10}]'$ and the structural parameters are $\theta = [\beta_1, \gamma_1, \beta_2, \gamma_2]'$. In addition, there is an auxiliary parameter ψ.

Identified Set

We will now provide a characterization of the identified set $\Theta(\phi)$. Define

$$G(\theta, \alpha) = \begin{bmatrix} G_1(\theta) \\ G_2(\theta) \end{bmatrix} - \begin{bmatrix} 0_{2\times 1} \\ \alpha \end{bmatrix}, \tag{B.5}$$

where

$$G_1(\theta) = \begin{bmatrix} d_1 d_2 \\ (1 - m_1)(1 - m_2) \end{bmatrix}, \quad G_2(\theta) = m_1(1 - d_2).$$

and

$$\alpha = (1 - \psi)(m_1 - d_1)(m_2 - d_2).$$

Moreover, let

$$\bar{\alpha}(\theta) = (m_1 - d_1)(m_2 - d_2) \tag{B.6}$$

and

$$Q(\theta;\phi) = \min_{0 \le \alpha \le \bar{\alpha}(\theta)} \Big\| \phi - G(\theta,\alpha) \Big\|. \tag{B.7}$$

Notice that by construction $Q(\theta;\phi) \ge 0$. In view of (B.2) to (B.4) and (B.5) it is straightforward to verify that the identified set can be characterized as follows:

$$\theta \in \Theta(\phi) \quad \text{iff} \quad Q(\theta;\phi) = 0.$$

Suppose we partition θ into $\theta = [\theta_1', \theta_2']'$. (B.2) and (B.3) imply that conditional on ϕ and θ_1 the subvector θ_2 is uniquely determined. Thus, the dimension of the identified set $\Theta(\phi)$ is 2. Since the entry game is symmetric with respect to Firm 1 and Firm 2, our illustration focuses on inference for θ_1. We denote the identified set for this subvector by $\Theta_1(\phi)$ and it can be characterized by the projection

$$\Theta_1(\phi) = \left\{ \theta_1 \;\middle|\; \exists\, \theta_2 \text{ s.t. } Q([\theta_1', \theta_2']';\phi) = 0 \right\}.$$

Frequentist Inference

Starting point of the frequentist inference is a large sample approximation of the sampling distribution of $\hat{\phi}_n$, defined as

$$\hat{\phi}_n = \left[\frac{n_{11}}{n}, \; \frac{n_{00}}{n}, \frac{n_{10}}{n} \right]', \tag{B.8}$$

where n_{11} is the number of markets with a duopoly, n_{00} is the number of markets without entry, and n_{10} is the number of markets with a Firm 1 monopoly. We assume that

$$\sqrt{n}(\hat{\phi}_n - \phi) \Longrightarrow N(0, \Lambda(\phi)) \tag{B.9}$$

uniformly in ϕ, where $\Lambda(\phi)$ can be consistently estimated by $\hat{\Lambda}$. Now define

$$Q_n(\theta;\hat{\phi}_n) = \min_{0 \le \alpha \le \bar{\alpha}(\theta)} n \Big\| \hat{\phi}_n - G(\theta,\alpha) \Big\|_{\hat{\Lambda}^{-1}}. \tag{B.10}$$

We shall construct a confidence set for θ as a level set of $Q_n(\theta;\hat{\phi}_n)$. To do so, we examine the sampling distribution of $Q_n(\theta;\hat{\phi}_n)$ for $\theta\in\Theta(\phi)$.

We partition $\hat{\phi}_n$ into $\hat{\phi}_{1,n}$ and $\hat{\phi}_{2,n}$ where the partitions conform with $G_1(\theta)$ and $G_2(\theta)$. Moreover, define

$$\hat{H}_1(\theta)=\hat{\phi}_{1,n}-G_1(\theta),\quad \hat{H}_2(\theta)=\hat{\phi}_{2,n}-G_2(\theta),$$

and partition $\hat{\Lambda}$ accordingly. In addition, let

$$\hat{H}_{2.11}(\theta)=\hat{H}_2(\theta)-\hat{\Lambda}_{21}\hat{\Lambda}_{11}^{-1}\hat{H}_1(\theta),\quad \hat{\Lambda}_{2.11}=\hat{\Lambda}_{22}-\hat{\Lambda}_{21}\hat{\Lambda}_{11}^{-1}\hat{\Lambda}_{12}.$$

Using the formula for factorizing a joint normal density into a marginal and a conditional density we can re-write the objective function as

$$Q_n(\theta;\hat{\phi}_n)=\min_{0\le\alpha\le\bar{\alpha}(\theta)} n\left(\big\|\hat{H}_1(\theta)\big\|_{\hat{\Lambda}_{11}^{-1}}+\big\|\hat{H}_{2.11}(\theta)+\alpha\big\|_{\hat{\Lambda}_{2.11}^{-1}}\right). \tag{B.11}$$

The minimizing value of α which we denote by $\hat{\alpha}(\theta)$ is given by

$$\hat{\alpha}(\theta)=\begin{cases}0 & \text{if } 0\le\hat{H}_{2.11}(\theta)\\ -\hat{H}_{2.11}(\theta) & \text{if } -\bar{\alpha}(\theta)\le\hat{H}_{2.11}(\theta)<0\\ \bar{\alpha}(\theta) & \text{otherwise}\end{cases}. \tag{B.12}$$

In turn, the objective function becomes

$$Q_n(\theta;\hat{\phi}_n)=\begin{cases}n\big\|\hat{H}_1(\theta)\big\|_{\hat{\Lambda}_{11}^{-1}}+n\big\|\hat{H}_{2.11}(\theta)\big\|_{\hat{\Lambda}_{2.11}^{-1}} & \text{if } 0\le\hat{H}_{2.11}(\theta)\\ n\big\|\hat{H}_1(\theta)\big\|_{\hat{\Lambda}_{11}^{-1}} & \text{if } -\bar{\alpha}(\theta)\le\hat{H}_{2.11}(\theta)<0\\ n\big\|\hat{H}_1(\theta)\big\|_{\hat{\Lambda}_{11}^{-1}}+n\big\|\hat{H}_{2.11}(\theta)+\bar{\alpha}(\theta)\big\|_{\hat{\Lambda}_{2.11}^{-1}} & \text{otherwise}\end{cases}. \tag{B.13}$$

As shown in Andrews and Guggenberger (2009), critical values for the construction of uniformly valid confidence sets can be obtained by considering the behavior of the objective function $Q_n(\cdot)$ under sequences of parameters. To do so, suppose data are generated based on $\phi_n=G(\theta_n,\alpha_n)$. To approximate the distribution of $Q_n(\theta_n;\hat{\phi}_n)$, notice that

$$\begin{aligned}\hat{H}_1(\theta_n) &= \hat{\phi}_{1,n}-G_1(\theta_n)\\ &= \hat{\phi}_{1,n}-\phi_{1,n}\\ \hat{H}_{2.11}(\theta_n) &= \hat{\phi}_{2,n}-G_2(\theta_n)-\hat{\Lambda}_{21}\hat{\Lambda}_{11}^{-1}[\hat{\phi}_{1,n}-G_1(\theta_n)]\\ &= \hat{\phi}_{2,n}-\phi_{2,n}-\alpha_n-\hat{\Lambda}_{21}\hat{\Lambda}_{11}^{-1}(\hat{\phi}_{1,n}-\phi_{1,n}).\end{aligned}$$

Let

$$Z_{1,n} = \sqrt{n}\hat{\Lambda}_{11}^{-1/2}(\hat{\phi}_{1,n} - \phi_{1,n}), \quad Z_{2.11,n} = \sqrt{n}\hat{\Lambda}_{2.11}^{-1/2}[\hat{\phi}_{2,n} - \phi_{2,n} - \hat{\Lambda}_{21}\hat{\Lambda}_{11}^{-1}(\hat{\phi}_{1,n} - \phi_{1,n})].$$

Using this notation, we can rewrite the objective function as

$$Q_n(\theta_n; \hat{\phi}_n) = \begin{cases} \|Z_{1,n}\| + \|Z_{2.11,n} - \sqrt{n}\hat{\Lambda}_{2.11}^{-1/2}\alpha_n\| & \text{if } \sqrt{n}\hat{\Lambda}_{2.11}^{-1/2}\alpha_n \leq Z_{2.11,n} \\ \|Z_{1,n}\| + \|Z_{2.11,n} + \sqrt{n}\hat{\Lambda}_{2.11}^{-1/2}(\bar{\alpha}(\theta_n) - \alpha_n)\| & \text{if } Z_{2.11,n} < -\sqrt{n}\hat{\Lambda}_{2.11}^{-1/2}(\bar{\alpha}(\theta_n) - \alpha_n) \\ \|Z_{1,n}\| & \text{otherwise} \end{cases}. \tag{B.14}$$

Now suppose that $\sqrt{n}\Lambda_{2.11}^{-1/2}\alpha_n \longrightarrow a$, $\sqrt{n}\Lambda_{2.11}^{-1/2}(\bar{\alpha}(\theta_n) - \alpha_n) \longrightarrow \bar{a}$, where $a \in \mathbb{R}^+ \cup \infty$ and $\bar{a} \in \mathbb{R}^+ \cup \infty$. Thus,

$$Q_n(\theta_n; \hat{\phi}_n) \Longrightarrow \begin{cases} \|Z_1\| + \|Z_{2.11} - a\| & \text{if } a \leq Z_{2.11} \\ \|Z_1\| + \|Z_{2.11} + \bar{a}\| & \text{if } Z_{2.11} < -\bar{a} \\ \|Z_1\| & \text{otherwise} \end{cases}, \tag{B.15}$$

where $Z_1 \sim N(0, I_2)$ and $Z_{2.11} \sim N(0,1)$ and Z_1 and $Z_{2.11}$ are independent. We have to distinguish three cases. First,

$$Q_n(\theta_n; \hat{\phi}_n) \Longrightarrow \|Z_1\| \leq \|Z_1\| + \|Z_{2.11}\| I\{Z_{2.11} \geq 0\} \quad \text{if} \quad a = \infty,\ \bar{a} = \infty.$$

Second,

$$Q_n(\theta_n; \hat{\phi}_n) \Longrightarrow \|Z_1\| + \|Z_{2.11} - a\| I\{Z_{2.11} \geq a\} \leq \|Z_1\| + \|Z_{2.11}\| I\{Z_{2.11} \geq 0\} \quad \text{if} \quad a < \infty,\ \bar{a} = \infty.$$

Third,

$$\begin{aligned} Q_n(\theta_n; \hat{\phi}_n) \quad &\Longrightarrow \quad \|Z_1\| + \|Z_{2.11} - a\| I\{Z_{2.11} \geq a\} + \|Z_{2.11} + \bar{a}\| I\{Z_{2.11} < -\bar{a}\} \quad \text{if} \quad a < \infty,\ \bar{a} < \infty \\ &\leq \quad \|Z_1\| + \|Z_{2.11}\|. \end{aligned}$$

The bound for this last case is weaker than the bounds for the first two cases. The case $\bar{a} < 0$ arises only if $\bar{\alpha}(\theta_n) \longrightarrow 0$ sufficiently fast, meaning that θ_n approaches an area of the parameter space in which the model is point identified. From the definition of $\bar{\alpha}(\theta)$ in (B.6) it follows that the third case arises if one of the interaction parameters is close to zero. In our numerical illustration we use a conservative fixed critical value obtained from the $1-\tau$ quantile of a $\chi^2(df=3)$.

A frequentist confidence set for the 4-dimensional parameter vector θ can then be defined as the level set

$$CS_F^{\theta}(Y^n) = \{\theta \mid Q_n(\theta; \hat{\phi}_n) \leq c_\tau^2\}. \tag{B.16}$$

We are restricting our attention to confidence sets constructed from fixed (rather than sample-size and θ dependent) critical values. In principle, one can construct the set $CS_F^{\theta}(Y^n)$ by evaluating the objective function $Q_n(\theta\ \hat{\phi}_n)$ on a 4-dimensional grid. However, since the identified set $\Theta(\phi)$ lies in a 2-dimensional subspace the specification of a suitable grid is difficult. Moreover, our goal is to construct a confidence set for the subvector θ_1. Thus, we let

$$\underline{Q}_n(\theta_1; \hat{\phi}_n) = \min_{\theta_2}\ Q_n([\theta_1', \theta_2]'; \hat{\phi}_n)$$

and define

$$CS_F^{\theta_1}(Y^n) = \{\theta \mid \underline{Q}_n(\theta_1; \hat{\phi}_n) \leq c_\tau^2\}. \tag{B.17}$$

The confidence set $CS_F^{\theta_1}(Y^n)$ is the projection of $CS_F^{\theta}(Y^n)$ onto the domain of θ_1. To compute the projection-based confidence set we specify a 2-dimensional grid for θ_1 and evaluate the objective function $\underline{Q}_n(\theta_1; \hat{\phi}_n)$ for each grid point. A parameter value is included in the confidence set if $\underline{Q}_n(\theta_1; \hat{\phi}_n) \leq c_\tau^2$.

Bayesian Inference – Draws from Conditional Prior

Prior 1 and Prior 2 are specified on the θ - ψ space through densities $p(\theta, \psi)$. These priors induce a prior distribution on the reduced form parameters ϕ. As explained in the main text, the conditional prior of θ given ϕ will not get updated through the likelihood function and the posterior will converge to $p(\theta|\hat{\phi}_n)$. In order to characterize the conditional prior $p(\theta_1|\phi)$ we conduct the following change of variables. Let

$$Z = [\beta_1, \gamma_1, \beta_2, \gamma_2, \psi]' \tag{B.18}$$

and

$$X = [\beta_1, \gamma_1, \phi_{11}, \phi_{00}, \phi_{10}]'. \tag{B.19}$$

To convert a prior density for $Z = f(X)$ into a prior for X, we can use

$$p_X(X) = p_Z(f(X))|f'(X)|. \tag{B.20}$$

Once we have derived $p_X(X)$ we can proceed as follows. Notice that

$$p(\theta_1|\phi) \propto p(\theta_1, \phi). \tag{B.21}$$

We use a Random-Walk Metropolis Algorithm to generate draws from $p(\theta_1|\phi)$. For this algorithm it is sufficient to be able evaluate the joint density $p(\theta_1, \phi)$ numerically. Descriptions of the algorithm can be found in many textbooks, e.g., Geweke (2005). Our proposal density is multivariate Gaussian with a covariance matrix that equals a suitably scaled identity matrix.

We shall proceed by characterizing the function $f(X)$, component by component and then derive the Jacobian $f'(X)$. The following functional relationships will be useful:

$$m_1 = \Phi_N(\beta_1), \quad m_2 = \Phi_N(\beta_2), \quad d_1 = \Phi_N(\beta_1 - \gamma_1), \quad d_2 = \Phi_N(\beta_2 - \gamma_2).$$

Since we will have to solve for β_2 and γ_2, notice that

$$\beta_2 = \Phi_N^{-1}(m_2), \quad \gamma_2 = \Phi_N^{-1}(m_2) - \Phi_N^{-1}(d_2).$$

The Nash equilibrium conditions imply that

$$\begin{aligned}
\phi_{00} &= (1 - m_1)(1 - m_2) \\
\phi_{11} &= d_1 d_2 \\
\phi_{10} &= m_1(1 - m_2) + d_1(m_2 - d_2) + \psi(m_1 - d_1)(m_2 - d_2).
\end{aligned}$$

We can use these conditions to solve for m_2, d_2, and ψ:

$$\begin{aligned}
m_2 &= 1 - \frac{\phi_{00}}{1 - m_1} \\
d_2 &= \frac{\phi_{11}}{d_1} \\
\psi &= \frac{\phi_{10} - m_1(1 - m_2) - d_1(m_2 - d_2)}{(m_1 - d_1)(m_2 - d_2)}.
\end{aligned}$$

The expression for ψ can be simplified by replacing m_2 and d_2:

$$\begin{aligned}
\psi &= \frac{\phi_{10} - m_1(1-m_2) - d_1(m_2-d_2)}{(m_1-d_1)(m_2-d_2)} \\
&= \frac{\phi_{10} - \phi_{00}\frac{m_1}{1-m_1} - d_1\left(1 - \frac{\phi_{00}}{1-m_1} - \frac{\phi_{11}}{d_1}\right)}{(m_1-d_1)\left(1 - \frac{\phi_{00}}{1-m_1} - \frac{\phi_{11}}{d_1}\right)} \\
&= \frac{\phi_{10}(1-m_1) - \phi_{00}m_1 - d_1\left(1 - m_1 - \phi_{00} - \frac{\phi_{11}(1-m_1)}{d_1}\right)}{(m_1-d_1)\left(1 - m_1 - \phi_{00} - \frac{\phi_{11}(1-m_1)}{d_1}\right)} \\
&= \frac{\phi_{10}(1-m_1) - \phi_{00}m_1 - d_1 g(X)}{(m_1-d_1)g(X)},
\end{aligned}$$

where

$$g(X) = \left(1 - m_1 - \phi_{00} - \frac{\phi_{11}(1-m_1)}{d_1}\right).$$

Combining terms, we obtain the following expressions for the components of $f(X)$:

$$\begin{aligned}
f_1(X) &= \beta_1 \\
f_2(X) &= \gamma_1 \\
f_3(X) &= \Phi_N^{-1}\left(1 - \frac{\phi_{00}}{1-\Phi_N(\beta_1)}\right) \\
f_4(X) &= f_3(X) - \Phi_N^{-1}\left(\frac{\phi_{11}}{\Phi_N(\beta_1-\gamma_1)}\right) \\
f_5(X) &= \frac{A_5(X)}{B_5(X)} = \frac{\phi_{10}(1-\Phi_N(\beta_1)) - \phi_{00}\Phi_N(\beta_1) - \Phi_N(\beta_1-\gamma_1)g(X)}{(\Phi_N(\beta_1) - \Phi_N(\beta_1-\gamma_1))g(X)}
\end{aligned}$$

where

$$g(X) = \left(1 - \Phi_N(\beta_1) - \phi_{00} - \frac{\phi_{11}(1-\Phi_N(\beta_1))}{\Phi_N(\beta_1-\gamma_1)}\right).$$

Now we can calculate the derivatives for the jacobian matrix. For this define

$$\psi(z) = \frac{\partial \Phi_N^{-1}(z)}{\partial z} = \frac{1}{\phi_N(\Phi_N^{-1}(z))}.$$

Term $f_1(X)$:

$$\frac{\partial f_1(X)}{\partial \beta_1} = 1.$$

Term $f_2(X)$:

$$\frac{\partial f_2(X)}{\partial \gamma_1} = 1.$$

Term $f_3(X)$:

$$\begin{aligned}
\frac{\partial f_3(X)}{\partial \beta_1} &= -\psi \left(1 - \frac{\phi_{00}}{1-\Phi_N(\beta_1)}\right) \frac{\phi_{00}}{[1-\Phi_N(\beta_1)]^2} \phi_N(\beta_1) \\
\frac{\partial f_3(X)}{\partial \phi_{00}} &= -\psi \left(1 - \frac{\phi_{00}}{1-\Phi_N(\beta_1)}\right) \frac{1}{1-\Phi_N(\beta_1)}.
\end{aligned}$$

Term $f_4(X)$:

$$\begin{aligned}
\frac{\partial f_4(X)}{\partial \beta_1} &= \frac{\partial f_3(X)}{\partial \beta_1} + \psi \left(\frac{\phi_{11}}{\Phi_N(\beta_1-\gamma_1)}\right) \frac{\phi_{11}\phi_N(\beta_1-\gamma_1)}{\Phi_N^2(\beta_1-\gamma_1)} \\
\frac{\partial f_4(X)}{\partial \gamma_1} &= -\psi \left(\frac{\phi_{11}}{\Phi_N(\beta_1-\gamma_1)}\right) \frac{\phi_{11}\phi_N(\beta_1-\gamma_1)}{\Phi_N^2(\beta_1-\gamma_1)} \\
\frac{\partial f_4(X)}{\partial \phi_{11}} &= -\psi \left(\frac{\phi_{11}}{\Phi_N(\beta_1-\gamma_1)}\right) \frac{1}{\Phi_N(\beta_1-\gamma_1)} \\
\frac{\partial f_4(X)}{\partial \phi_{00}} &= \frac{\partial f_3(X)}{\partial \phi_{00}}.
\end{aligned}$$

Term $f_5(X)$:

$$\frac{\partial f_5(X)}{\partial x} = \frac{\frac{\partial A(X)}{\partial x} B(X) - A(X)\frac{\partial B(X)}{\partial x}}{B(X)^2}.$$

Term $A(X)$:

$$\begin{aligned}
\frac{\partial A(X)}{\partial \beta_1} &= -(\phi_{10}+\phi_{00})\phi_N(\beta_1) - \phi_N(\beta_1-\gamma_1)g(X) - \Phi_N(\beta_1-\gamma_1)\frac{\partial g(X)}{\partial \beta_1} \\
\frac{\partial A(X)}{\partial \gamma_1} &= \phi_N(\beta_1-\gamma_1)g(X) - \Phi_N(\beta_1-\gamma_1)\frac{\partial g(X)}{\partial \gamma_1} \\
\frac{\partial A(X)}{\partial \phi_{11}} &= -\Phi_N(\beta_1-\gamma_1)\frac{\partial g(X)}{\partial \phi_{11}} \\
\frac{\partial A(X)}{\partial \phi_{00}} &= -\Phi_N(\beta_1) - \Phi_N(\beta_1-\gamma_1)\frac{\partial g(X)}{\partial \phi_{00}} \\
\frac{\partial A(X)}{\partial \phi_{10}} &= (1-\Phi_N(\beta_1)) - \Phi_N(\beta_1-\gamma_1)\frac{\partial g(X)}{\partial \phi_{10}}.
\end{aligned}$$

Term $B(X)$:

$$\begin{aligned}
\frac{\partial B(X)}{\partial \beta_1} &= (\phi_N(\beta_1) - \phi_N(\beta_1 - \gamma_1))g(X) + (\Phi_N(\beta_1) - \Phi_N(\beta_1 - \gamma_1))\frac{\partial g(X)}{\partial \beta_1} \\
\frac{\partial B(X)}{\partial \gamma_1} &= \phi_N(\beta_1 - \gamma_1)g(X) + (\Phi_N(\beta_1) - \Phi_N(\beta_1 - \gamma_1))\frac{\partial g(X)}{\partial \gamma_1} \\
\frac{\partial B(X)}{\partial \phi_{11}} &= (\Phi_N(\beta_1) - \Phi_N(\beta_1 - \gamma_1))\frac{\partial g(X)}{\partial \phi_{11}} \\
\frac{\partial B(X)}{\partial \phi_{00}} &= (\Phi_N(\beta_1) - \Phi_N(\beta_1 - \gamma_1))\frac{\partial g(X)}{\partial \phi_{00}}.
\end{aligned}$$

Term $g(X)$:

$$\begin{aligned}
\frac{\partial g(X)}{\partial \beta_1} &= -\phi_N(\beta_1) + \frac{\phi_{11}\phi_N(\beta_1)}{\Phi_N(\beta_1 - \gamma_1)} + \frac{\phi_{11}(1 - \Phi_N(\beta_1))\phi_N(\beta_1 - \gamma_1)}{\Phi_N^2(\beta_1 - \gamma_1)} \\
\frac{\partial g(X)}{\partial \gamma_1} &= -\frac{\phi_{11}(1 - \Phi_N(\beta_1))\phi_N(\beta_1 - \gamma_1)}{\Phi_N^2(\beta_1 - \gamma_1)} \\
\frac{\partial g(X)}{\partial \phi_{11}} &= -\frac{1 - \Phi_N(\beta_1)}{\Phi_N(\beta_1 - \gamma_1)} \\
\frac{\partial g(X)}{\partial \phi_{00}} &= -1.
\end{aligned}$$

Bayesian Inference – Draws from Posterior

According to Equations (B.1) to (B.4) we can express the reduced form probabilities as functions of θ and ψ. Thus, the likelihood function is given by

$$p(Y^n|\theta, \psi) = \phi_{11}^{n_{11}}(\theta, \psi)\phi_{00}^{n_{00}}(\theta, \psi)\phi_{10}^{n_{10}}(\theta, \psi)\phi_{01}^{n_{01}}(\theta, \psi). \tag{B.22}$$

If this prior distribution is combined with a prior specified on the θ - ψ space, then the posterior is given by

$$p(\theta, \psi|Y^n) \propto p(Y^n|\theta, \psi)p(\theta, \psi) \tag{B.23}$$

and draws can be generated with a Random Walk Metropolis Algorithm.

In addition to Priors 1 and 2 we consider a prior that is flat with respect to the reduced form parameters. Conditional on ϕ, the prior for θ_1 is uniform on the identified set $\Theta_1(\phi)$. In order to obtain draws from the posterior distribution of θ_1 we sample from (i) $p(\phi|Y^n)$ and

(ii) from $p(\theta_1|\phi)$. For Step (i) notice that under the flat prior for ϕ, the posterior distribution $P^{\phi}_{Y^n}$ takes the form of a Dirichlet distribution

$$[\phi_{11}, \phi_{00}, \phi_{10}, \phi_{01}]' \sim \text{ Dirichlet } (n_{11}+1, n_{00}+1, n_{10}+1, n_{01}).$$

A draw from this Dirichlet distribution can be generated as follows: Let $a_j \sim \mathcal{G}(n_j+1, 1)$, where $j \in \{11, 00, 10, 01\}$ and $\mathcal{G}(\alpha, 1)$ denotes a Gamma distribution with shape parameter α and scale parameter 1. Then set

$$\phi = [a_{11}, a_{00}, a_{10}, a_{01}]'/(a_{11}+a_{00}+a_{10}+a_{01}).$$

For Step (ii) we specify a two-dimensional grid for θ_1 in order construct projections of the identified set $\Theta_1(\phi)$ onto the β_1 and γ_1 ordinates. Let these projections be delimited by $\underline{\beta}_1$, $\bar{\beta}_1$, $\underline{\gamma}_1$, and $\bar{\gamma}_1$. We then use an acceptance sampler with a proposal density that is uniform on $[\underline{\beta}_1, \bar{\beta}_1] \otimes [\underline{\gamma}_1, \bar{\gamma}_1]$ to obtain a draw of θ_1 conditional on ϕ.

Bayesian Inference – Credible Sets

Credible sets are computed according to the following steps:

1. Construct two independent sequences $\{\theta^{(1)}_{1,s}\}_{s=1}^S$ and $\{\theta^{(2)}_{1,s}\}_{s=1}^S$ of draws from the distribution of θ_1.

2. Use the $\{\theta^{(1)}_{1,s}\}_{s=1}^S$ draws to construct Kernel density estimates $\hat{p}(\theta^{(2)}_{1,s})$ for each $\theta^{(2)}_{1,s}$, $s = 1, \dots, S$.

3. Find a cutoff κ such that $(1-\tau)S$ of the density estimates $\hat{p}(\theta^{(2)}_{1,s})$ are greater or equal than κ.

4. Use the $\{\theta^{(1)}_{1,s}\}_{s=1}^S$ draws to construct Kernel density estimates $\hat{p}(\theta_1)$ for values of θ_1 on a 2-dimensional grid. Include a particular grid point into the credible set if $\hat{p}(\theta_1) \geq \kappa$.