

Supplement to “Minimum distance estimators for dynamic games”

(*Quantitative Economics*, Vol. 4, No. 3, November 2013, 549–583)

SORAWOOT SRISUMA
School of Economics, University of Surrey

This supplementary material contains Appendices A and B. Appendix A illustrates the theoretical issue of consistent estimation related to Bajari, Benkard, and Levin’s (2007) inequality approach, as well as providing some remedies for discrete action games. Appendix B contains the proofs of Theorems 1 and 2 in the main paper.

APPENDIX A: CONSISTENT ESTIMATION WITH BBL’S METHODOLOGY

This appendix illustrates a potential problem with the inequality approach of Bajari, Benkard, and Levin’ (2007, hereafter BBL). We provide two examples in Section A.1, each showing a scenario where the inequality restrictions imposed by the equilibrium are satisfied by a unique element in the parameter space and the uniqueness can be lost when a strict subclass of inequalities is considered. The first example has no conditioning variables so as to emphasize that the source of information loss here differs from the instrumental variable model in Domínguez and Lobato (2004). The second example corresponds to Design 2 of the simulation study in Section 5. In Section A.2, we provide a class of inequalities that retains the identifying information of the (identified) parameters of some discrete action games. We conclude with a brief discussion in Section A.3.

A.1 *Mathematical examples*

Single agent problem Consider a simple optimization problem where an economic agent maximizes the payoff function

$$u_{\theta}(a, \varepsilon) = -a^2 + 2\theta a\varepsilon.$$

Here a and ε denote the action and state variables, respectively, and θ belongs to Θ , some positive subset of \mathbb{R} . The model is generated from some distribution of ε_n that is absolutely continuous with respect to the Lebesgue measure and has finite second moment. Notice that the current setup satisfies conditions in Section 2.3 as a special case of a single agent static decision problem ($\beta = 0$ and $I = 1$). Since the payoff function

Sorawoot Srisuma: s.srisuma@surrey.ac.uk

Copyright © 2013 Sorawoot Srisuma. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](https://creativecommons.org/licenses/by-nc/3.0/). Available at <http://www.qeconomics.org>.

DOI: 10.3982/QE266

is concave, the optimal strategy follows from the first order condition

$$\alpha_\theta(\varepsilon_n) = \theta\varepsilon_n \quad \text{a.s. for all } \theta \in \Theta.$$

It is clear that the distribution of $\alpha_\theta(\varepsilon_n)$ is identified. Let θ_0 denote the true parameter and suppose we observe a random sample $\{a_n\}_{n=1}^N$, where $a_n = \alpha_{\theta_0}(\varepsilon_n)$ for each n .

The inequality approach of BBL defines an estimator for θ_0 that satisfies the system of moment inequalities in the limit,

$$E[u_\theta(\alpha_{\theta_0}(\varepsilon_n), \varepsilon_n)] \geq E[u_\theta(\tilde{\alpha}(\varepsilon_n), \varepsilon_n)] \quad \text{for all } \tilde{\alpha} \in \mathfrak{A}_0, \quad (\text{SA1})$$

where \mathfrak{A}_0 is some user-chosen class of functions (of alternative strategies). We first consider a popular class of strategies based on additive perturbations and show that it cannot be used to identify θ_0 . Formally, let \mathbb{S} be some subset of \mathbb{R} . Then define $\mathfrak{A}_0(\mathbb{S}) = \{\tilde{\alpha}(\cdot; \eta) \text{ for } \eta \in \mathbb{S} : \tilde{\alpha}(\varepsilon; \eta) = \alpha_{\theta_0}(\varepsilon) + \eta \text{ for all } \varepsilon \in \mathcal{E}\}$.¹ It follows from some simple algebra that, for any η ,

$$E[u_\theta(\alpha_{\theta_0}(\varepsilon_n), \varepsilon_n)] - E[u_\theta(\tilde{\alpha}(\varepsilon_n; \eta), \varepsilon_n)] = \eta^2 + 2\eta(\theta_0 - \theta)E[\varepsilon_n].$$

When ε_n has mean zero, $\mathfrak{A}_0(\mathbb{S})$ has no identifying information for θ_0 in the sense that, for all $\theta \in \Theta$,

$$E[u_\theta(\alpha_{\theta_0}(\varepsilon_n), \varepsilon_n)] \geq E[u_\theta(\tilde{\alpha}(\varepsilon_n), \varepsilon_n)] \quad \text{for all } \tilde{\alpha} \in \mathfrak{A}_0(\mathbb{S}),$$

even if $\mathbb{S} = \mathbb{R}$. Therefore, $\mathfrak{A}_0(\mathbb{S})$ cannot be used to consistently estimate θ_0 .

However, the set of inequalities that considers *all* alternative strategies can actually identify θ_0 . To see this, we begin by calculating the difference between the expected returns from α_{θ_0} and a generic alternative strategy $\tilde{\alpha}$:

$$E[u_\theta(\alpha_{\theta_0}(\varepsilon_n), \varepsilon_n)] - E[u_\theta(\tilde{\alpha}(\varepsilon_n), \varepsilon_n)] = -(\theta - \theta_0)^2 E[\varepsilon_n^2] + E[(\theta\varepsilon_n - \tilde{\alpha}(\varepsilon_n))^2].$$

If we consider an inequality based on multiplicative perturbation, say $\mathfrak{A}_1(\mathbb{S}) = \{\tilde{\alpha}(\cdot; \eta) \text{ for } \eta \in \mathbb{S} : \tilde{\alpha}(\varepsilon; \eta) = \eta\alpha_{\theta_0}(\varepsilon) \text{ for all } \varepsilon \in \mathcal{E}\}$, then by choosing $\tilde{\alpha}$ from $\mathfrak{A}_1(\mathbb{S})$, the difference above simplifies to $((\theta - \eta\theta_0)^2 - (\theta - \theta_0)^2)E[\varepsilon_n^2]$. It is easy to see that whenever $\theta \neq \theta_0$, the inequality in (SA1) will be violated for some range of values of η sufficiently close to 1: more precisely, if $\theta > \theta_0$, then violation occurs for $\eta \in (1, \theta/\theta_0)$; otherwise take $\eta \in (\theta/\theta_0, 1)$. Therefore, the class of multiplicative perturbations has sufficient identifying power for θ_0 in the sense that when \mathbb{S} contains any open ball centered at 1, then

$$E[u_\theta(\alpha_{\theta_0}(\varepsilon_n), \varepsilon_n)] \geq E[u_\theta(\tilde{\alpha}(\varepsilon_n), \varepsilon_n)] \quad \text{for all } \tilde{\alpha} \in \mathfrak{A}_1(\mathbb{S}) \text{ if and only if } \theta = \theta_0.$$

¹In an application of the BBL methodology, the user puts a distribution on η that has support \mathbb{S} . A random sequence from this distribution is then drawn to construct the objective function; for instance, if $\mathbb{S} = \mathbb{R}$, then η can be drawn from a normal distribution.

Cournot game Consider the setup of Design 2 in Section 5. Here we give a slightly more informal argument for why inequalities based on additive perturbation lose some identifying information on the data generating parameter while multiplicative perturbations can preserve it.

Consider player 1. For any given a_2, x, ε_1 , $u_{1,\theta}(a_1, a_2, x, \varepsilon_1)$ is concave in a_1 since $\theta_1 > 0$. Taking the first derivative gives

$$\frac{\partial}{\partial a_1} u_{1,\theta}(a_1, a_2, x, \varepsilon_1) = x - \theta_2 \varepsilon_1 - \theta_1 x (2a_1 + a_2).$$

Since a_1 and a_2 enter the first derivative linearly and separately, the expected (symmetric) optimal action, which we denote by γ_θ , can be obtained by finding the zero to solve the first order condition

$$\gamma_\theta(x_n) = \arg \text{zero}_{a \in A} E \left[\frac{\partial}{\partial a_1} u_{1,\theta}(a_1, a_2, x_n, \varepsilon_{1n}) | x_n \right] \Big|_{a_1=a_2=a}.$$

Given that ε_{1n} is a random variable with mean 0 and variance 1, it then follows that $\gamma_\theta(x_n) = \frac{1}{3\theta_1}$. Therefore, for any x, ε_1 , player 1's optimal choice, $\alpha_\theta(x, \varepsilon_1)$, can be characterized by the zero of $\frac{\partial}{\partial a_1} u_{1,\theta}(a_1, \gamma_\theta(x), x, \varepsilon_1)$ that is equal to $\frac{1}{3\theta_1} - \frac{\theta_2 \varepsilon_1}{2x\theta_1}$. It is clear that the distribution of $\alpha_\theta(x, \varepsilon_n)$ is identified.

Suppose the data are generated from a random sample of $\{a_{1n}, a_{2n}, x\}_{n=1}^N$, where $a_{in} = \alpha_{\theta_0}(x_n, \varepsilon_{in})$ for $i = 1, 2$ and every n , for some $\theta_0 = (\theta_{01}, \theta_{02}) \in \mathbb{R}^+ \times \mathbb{R}^+$. To study whether additive perturbations can be used to construct objective functions that identify θ_0 , we consider $u_{1,\theta}(a_1 + \eta, \gamma_{\theta_0}(x), x, \varepsilon_1)$ for some η . Through some tedious algebra, it can be shown that

$$\begin{aligned} & u_{1,\theta}(a_1 + \eta, \gamma_{\theta_0}(x), x, \varepsilon_1) \\ &= u_{1,\theta}(a_1, \gamma_{\theta_0}(x), x, \varepsilon_1) + \eta(x - \theta_1 x \gamma_{\theta_0}(x) - \theta_2 \varepsilon) - \theta_1 x (2a_1 \eta + \eta^2). \end{aligned}$$

Comparing the expected returns from using the optimal strategy and a perturbed one gives

$$\begin{aligned} & E[u_{1,\theta}(\alpha_\theta(s_{1n}), \gamma_{\theta_0}(x_n), s_{1n}) | x_n = x] - E[u_{1,\theta}(\alpha_{\theta_0}(s_{1n}) + \eta, \gamma_{\theta_0}(x_n), s_{1n}) | x_n = x] \\ &= -\eta(x - \theta_1 x \gamma_{\theta_0}(x)) + \theta_1 x (2\gamma_{\theta_0}(x) \eta + \eta^2) \\ &= -\eta x \left(1 - \frac{\theta_1}{\theta_{01}} \right) + \theta_1 x \eta^2. \end{aligned}$$

Clearly, $\theta' = (\theta_{01}, \theta'_2)$ satisfies the necessary condition implied by the equilibrium for all values of θ'_2 . Therefore, the objective functions constructed using additive perturbations cannot identify θ_{02} in the limit. Next, we consider the multiplicative perturbation. For the calculations, it is convenient to write the multiplicative factor as $(1 + \eta)$. Then it can be shown that

$$\begin{aligned} & u_{1,\theta}(a_1(1 + \eta), \gamma_{\theta_0}(x), s_1) \\ &= u_{1,\theta}(a_1, \gamma_{\theta_0}(x), s_1) + \eta a_1 (x - \theta_1 x \gamma_{\theta_0}(x) - \theta_2 \varepsilon) - \theta_1 x (2\eta + \eta^2) a_1^2. \end{aligned}$$

Taking conditional expectation and comparing the expected returns gives

$$\begin{aligned} & E[u_{1,\theta}(\alpha_{\theta_0}(s_{1n}), \gamma_{\theta_0}(x_n), s_{1n})|x_n = x] \\ & \quad - E[u_{1,\theta}(\alpha_{\theta_0}(s_{1n}) + \eta, \gamma_{\theta_0}(x_n), s_{1n})|x_n = x] \\ & = -\frac{\eta}{\theta_{01}} \left(\frac{x}{3} \delta_1 + \frac{\theta_{02}}{2x} (\delta_1 \theta_{02} + \delta_2) \right) + \theta_1 x \eta^2 \left(\left(\frac{1}{3\theta_{01}} \right)^2 + \left(\frac{\theta_{02}}{2\theta_{01}x} \right)^2 \right), \end{aligned}$$

where $\delta_1 = 1 - \frac{\theta_1}{\theta_{01}}$ and $\delta_2 = \theta_2 - \theta_{02}$. For any $\delta_1, \delta_2 \neq 0$, with a small enough $|\eta|$, the squared (second) term above is of smaller order and the first term will be strictly negative for some state x with either $\eta > 0$ or $\eta < 0$. Therefore, we expect $\{\tilde{\alpha}(\cdot; \eta)\}$ for $\eta \in \mathbb{S}: \tilde{\alpha}(s_i; \eta) = \eta \alpha_{\theta_0}(s_i)$ for all $s_i \in S_i\}$ to be able to preserve the identifying information of θ_0 when \mathbb{S} contains an open ball centered at 1.

A.2 Perturbations for discrete action games

We first consider a binary action game that satisfy Assumptions M1, M2, M3, and D' in Section 2, where D' is the parameterized version of D that replaces u_i everywhere with $u_{i,\theta}$. To keep the calculation of the expected returns tractable, we only use the class of alternative strategies where players only deviate from the equilibrium action in the first stage; BBL (see p. 1348) also suggested this among other ways to construct inequalities. In particular, we can, therefore, adopt the framework of the pseudo-model constructed in Section 3.1. Suppose the data $\{(a_{in}, \mathbf{a}_{-in}, x_n, x'_n)\}_{n=1}^N$ are generated from a pure strategy Markov equilibrium when $\theta = \theta_0$. In the limit, the pseudo-objective function (see equation (8)) is

$$\begin{aligned} \Lambda_{i,\theta}(a_i, x, \varepsilon_i) & = E[u_{i,\theta}(a_i, \mathbf{a}_{-in}, x_n, \varepsilon_i)|x_n = x] + \beta_i g_{i,\theta}(a_i, x) \\ & = v_{i,\theta}(a_i, x) + \varepsilon_i(a_i), \end{aligned}$$

where $v_{i,\theta}(a_i, x) = E[\pi_{i,\theta}(a_i, \mathbf{a}_{-in}, x_n)|x_n = x] + \beta_i g_{i,\theta}(a_i, x)$; Pesendorfer and Schmidt-Dengler (2008) called $v_{i,\theta}$ the continuation value net of the payoff shocks. Since we only focus on identification, $v_{i,\theta}$ is taken as known; conditions for consistent estimation of $v_{i,\theta}$ and other details can be found in Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008). It is also convenient to define the differences between the choice-specific continuation values and private values. Let $\Delta \Lambda_{i,\theta}(a_i, a'_i, s_i) = \Lambda_{i,\theta}(a_i, s_i) - \Lambda_{i,\theta}(a'_i, s_i)$, and also let $\Delta v_{i,\theta}(x) = v_{i,\theta}(1, x) - v_{i,\theta}(0, x)$ and $\omega_{in} = \varepsilon_{in}(0) - \varepsilon_{in}(1)$. Note that under Assumption D(iii), ω_{in} is absolutely continuous with respect to the Lebesgue measure with support on \mathbb{R} . The pseudo-best response is characterized by a cutoff rule,

$$\alpha_{i,\theta}(s_{in}) = \mathbf{1}[\Delta v_{i,\theta}(x_n) > \omega_{in}] \quad \text{a.s. for all } \theta \in \Theta \text{ and } i = 1, \dots, I.$$

Then $\Delta v_{i,\theta_0}(x)$ is identified from $Q_{\omega_i}^{-1}(P_i(1|x))$, where $P_i(1|x)$ denotes the underlying equilibrium choice probability of choosing action 1 and $Q_{\omega_i}^{-1}$ is the inverse of the distribution function of ω_{in} .

We assume that θ_0 is identified (see Definition 3 in Section 4.1) and we claim that a class of alternative strategies that consists of perturbing the cutoff values has sufficient identifying power for θ_0 . More formally, let $\mathfrak{A}_i^U(\mathbb{S}) = \{\tilde{\alpha}_i(\cdot; \eta)\}$ for $\eta \in \mathbb{S}: \tilde{\alpha}_i(s_i; \eta) =$

$\mathbf{1}[\Delta v_{i,\theta_0}(x) + \eta > \omega_i]$ for all $s_i \in S_i$. Then $\mathfrak{A}_i^U(\mathbb{S})$ has sufficient identifying power for θ_0 in the sense that

$$E[A_{i,\theta}(\alpha_{i,\theta_0}(s_{in}), s_{in})|x_n = x] \geq E[A_{i,\theta}(\tilde{\alpha}_i(s_{in}), s_{in})|x_n = x] \quad (\text{SA2})$$

for all i, x and $\tilde{\alpha}_i \in \mathfrak{A}_i^U(\mathbb{S})$ if and only if $\theta = \theta_0$,

for some appropriate \mathbb{S} . To see this, we first show that whenever $\theta \neq \theta_0$, we can find some i, s_i , and η such that $\Delta \Lambda_{i,\theta}(\alpha_{i,\theta_0}(s_i), \tilde{\alpha}_i(s_i; \eta), s_i) < 0$.

Since θ_0 is identified, for any $\theta \neq \theta_0$, there exists some i, x and $\xi \neq 0$ such that $\Delta v_{i,\theta}(x) = \Delta v_{i,\theta_0}(x) + \xi$. Suppose $\xi > 0$. Then any $\eta \in (0, \xi)$ implies

$$\begin{aligned} & \Delta \Lambda_{i,\theta}(\alpha_{i,\theta_0}(s_i), \tilde{\alpha}_i(s_i; \eta), s_i) \\ &= \begin{cases} -(\Delta v_{i,\theta}(x) - \omega_i) < 0, & \text{for } \omega_i \in (\Delta v_{i,\theta_0}(x), \Delta v_{i,\theta_0}(x) + \eta), \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

By an analogous argument, when $\xi < 0$, choosing any $\eta \in (\xi, 0)$ implies that $\Delta \Lambda_{i,\theta}(\alpha_{i,\theta_0}(s_i), \tilde{\alpha}_i(s_i; \eta), s_i)$ takes strictly negative values for all $\omega_i \in (\Delta v_{i,\theta_0}(x) + \eta, \Delta v_{i,\theta_0}(x))$ and is 0 otherwise. Since ω_{in} has a continuous distribution on \mathbb{R} , $E[\Delta \Lambda_{i,\theta}(\alpha_{i,\theta_0}(s_{in}), \tilde{\alpha}_i(s_{in}; \eta), s_{in})|x_n = x] < 0$ for all η on either $(-\xi, 0)$ or $(0, \xi)$ with small enough $\xi > 0$. Therefore, the class of perturbations at the cutoff value has sufficient identifying power for θ_0 if \mathbb{S} contains any open ball that is centered at 0.

Although we do not provide any formal details, due to nontrivial additional notational complexity, an analogous idea can be used for multinomial action games. Suppose $K_i = K$ for all i . Then the optimality condition for the $(K + 1)$ choice problem can be characterized, for each player and state, by K inequality constraints that partition \mathbb{R}^K —the support of the normalized private values. The role of a cutoff value is then replaced by a locus point in \mathbb{R}^K , which is uniquely identified by the inversion result of Hotz and Miller (1993) subject to the choice of a normalization action. Then analogous alternative strategies can be constructed by additively perturbing the locus point using a K -dimensional variable whose support includes a ball in \mathbb{R}^K that contains the origin.

The intuition used in the unordered binary action game can also be applied to the class of discrete monotone action games. Specifically, we now assume M1, M2, M3, S1', S2, and S3', and let the data $\{(a_{in}, \mathbf{a}_{-in}, x_n)\}_{n=1}^N$ be generated from a pure strategy Markov equilibrium when $\theta = \theta_0$. Recall that $\alpha_{i,\theta}(x, \cdot)$ is a nondecreasing function on \mathcal{E}_i (by the arguments of Lemmas 1 and 2). For notational simplicity, suppose that $A_i = \{0, 1\}$ for all i . Then the pseudo-best response is uniquely characterized by a cutoff rule

$$\alpha_{i,\theta}(s_{in}) = \mathbf{1}[\mathcal{C}_{i,\theta}(x_n) \geq \varepsilon_{in}] \quad \text{a.s. for all } \theta \in \Theta \text{ and } i = 1, \dots, I,$$

for some $\mathcal{C}_{i,\theta}$ such that $\underline{\varepsilon}_i \leq \mathcal{C}_{i,\theta}(x) \leq \bar{\varepsilon}_i$ for all i, x, θ . In particular, when $\Pr[\alpha_{i,\theta}(s_{in}) = 1|x_n = x] = 0$, set $\mathcal{C}_{i,\theta}(x) = \underline{\varepsilon}_i$, and when $\Pr[\alpha_{i,\theta}(s_{in}) = 1|x_n = x] = 1$, set $\mathcal{C}_{i,\theta}(x) = \bar{\varepsilon}_i$. As seen previously, $\mathcal{C}_{i,\theta_0}(x)$ is identified by $Q_i^{-1}(P_i(1|x))$, where Q_i^{-1} denotes the inverse of the distribution function of ε_{in} . If θ_0 is identified, then the class of alternative strategies $\mathfrak{A}_i^O = \{\tilde{\alpha}_i(\cdot; \eta) \text{ for } \eta \in \mathbb{S}: \tilde{\alpha}_i(s_i; \eta) = \mathbf{1}[\mathcal{C}_{i,\theta_0}(x) + \eta > \varepsilon_i] \text{ for all } s_i \in S_i\}$ has sufficient identifying power for θ_0 in the sense described in equation (SA2). When there are more than

two actions, suppose $K_i = K$ for all i , then the data generating best response is generally characterized by $K - 1$ boundary points on \mathcal{E}_i for each player and state. These boundary points can be identified from F_i and Q_i . Since $\mathcal{E}_i \subseteq \mathbb{R}$, a simple way to apply the same technique used in binary action games above is to choose the set of alternative strategies that perturb only one of the boundary points at a time and leave all other boundary points the same as those identified by the data.

A.3 A discussion

The inequality moment restrictions imposed by the equilibrium condition considered in BBL is indexed by a class of functions of alternative strategies. Our examples in Section A.1 illustrate a general point that some alternative strategies may have no identifying information for a subset of the parameter of interest (or the entire parameter space in some cases). In contrast to the examples in Domínguez and Lobato (2004), objective functions constructed from certain classes of alternative strategies not only lack global identification (i.e. do not have a unique optimum), they cannot even distinguish between different parameters locally. We only provide an example when the inequality approach suggested by BBL can fail for a point-identified model (most known applications of their methodology proceed under this assumption). Although BBL also suggested a set estimator for partially identified models, it is intuitively clear that their set estimation approach is exposed to the same criticism as above, in which case some classes of inequalities may only be able to identify a strict superset of the identified set.

We consider dynamic games in Section A.2. We focus on alternative strategies where each player only deviates in the first stage since it provides a more tractable starting point to study identification. It enables us to show that when the parameter is identified in binary action games, inequalities generated from additively perturbing the cutoff values preserve the identifying information. We also explain how such technique can be applied to multinomial choice games as well as discrete action games where players play monotone strategies. However, it is clearly impractical to extend the suggested perturbation method for discrete action games to a continuous action game.

Finally, all of our analytical arguments above only apply to the limiting case where equilibrium and alternative strategies are perfectly known and there are no simulation errors. As the Monte Carlo study in Section 5 shows, it is always possible to obtain an estimate in finite samples, even when the objective function cannot identify the parameter of interest in the limit. Our main message is the choice of alternative strategies, which can be viewed as tuning parameters, is very important since it affects not only efficiency, but also consistency. It remains an interesting issue to find some sufficiency theory for choosing inequalities in a continuous action game.

APPENDIX B: PROOFS OF THEOREMS

Since the first stage estimators are defined implicitly in our objective function $\widehat{M}_N(\theta)$, it suffices to show that Assumptions A1 and A2 imply some familiar conditions from large sample theorems for parametric estimators. For Theorem 1, we make use of a

well known consistency result for extremum estimators; for instance, see Theorem 2.1 of Newey and McFadden (1994). For Theorem 2, we show that A1 and A2 are sufficient for the conditions of Theorem 7.1 of Newey and McFadden (1994), who provided a high level condition for the asymptotic normality of an extremum estimator that maximizes a nonsmooth objective function.

PROOF OF THEOREM 1. Under A1(i), Θ is compact. Assumption A1(ii)–(iv) ensure that $M(\theta)$ has a well separated minimum at θ_0 . Next, we show that the sample objective function converges uniformly in probability to its limit. By the triangle inequality,

$$\begin{aligned} |\widehat{M}_N(\theta) - M(\theta)| &\leq 4 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} |\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x)| \mu_{i,x}(da_i) \\ &\quad + 4 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} |\widehat{F}_{i,\theta}(a_i|x) - F_{i,\theta}(a_i|x)| \mu_{i,x}(da_i) \\ &\quad + 4 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} |\widehat{F}_i(a_i|x) - F_i(a_i|x)| \mu_{i,x}(da_i) \end{aligned}$$

asymptotically since distribution functions are bounded above by 1 and $\widetilde{F}_{i,\theta}, \widehat{F}_{i,\theta}$ are uniformly consistent under A1(v)–(vii). Under A1(iv), the measures are finite, hence $\sup_{\theta \in \Theta} |\widehat{M}_N(\theta) - M(\theta)| = o_p(1)$ by A1(v)–(vii). Consistency then follows by a standard argument. \square

PROOF OF THEOREM 2. Conditions (i)–(iii) of Newey and McFadden (1994, Theorem 7.1) are trivially satisfied by the definition of our estimator and condition A2(i) and (ii). It remains to show that there exists a sequence C_N that has an asymptotic normal distribution at the root- N rate, which satisfies the (*stochastic differentiability*) condition

$$\sup_{\|\theta - \theta_0\| < \delta_N} \left| \frac{\mathcal{D}_N(\theta)}{1 + \sqrt{N} \|\theta - \theta_0\|} \right| = o_p(1)$$

for any positive sequence $\delta_N = o(1)$, where

$$\mathcal{D}_N(\theta) = \sqrt{N} \frac{\widehat{M}_N(\theta) - \widehat{M}_N(\theta_0) - (M(\theta) - M(\theta_0)) - (\theta - \theta_0)^\top C_N}{\|\theta - \theta_0\|}.$$

We show that

$$\begin{aligned} &\widehat{M}_N(\theta) - \widehat{M}_N(\theta_0) - (\widehat{M}(\theta) - \widehat{M}(\theta_0)) - (\theta - \theta_0)^\top C_N \\ &= o_p \left(\|\theta - \theta_0\|^2 + \frac{\|\theta - \theta_0\|}{\sqrt{N}} + \frac{1}{N} \right) \end{aligned} \tag{SA3}$$

holds uniformly for $\|\theta - \theta_0\| \leq \delta_N$. The additional $o_p(N^{-1})$ term added in (SA3) does not affect Newey and McFadden's results as it is the rate that our estimator (approximately) minimizes the objective function, which coincides with condition (i) of their theorem.

For θ in a neighborhood of θ_0 , we write $\widehat{M}_N(\theta) - \widehat{M}_N(\theta_0) - (M(\theta) - M(\theta_0))$ as a sum, $E_1(\theta) + E_2(\theta)$, where

$$E_1(\theta) = M_N(\theta) - M_N(\theta_0) - (M(\theta) - M(\theta_0)),$$

$$E_2(\theta) = \widehat{M}_N(\theta) - \widehat{M}_N(\theta_0) - (M_N(\theta) - M_N(\theta_0)).$$

Under A2(ii), M_N and M are twice continuously differentiable in a neighborhood of θ_0 . By Taylor's theorem,

$$E_1(\theta) = (\theta - \theta_0)^\top \frac{\partial}{\partial \theta} M_N(\theta_0) + \frac{1}{2}(\theta - \theta_0)^\top \frac{\partial^2}{\partial \theta \partial \theta^\top} (M_N(\bar{\theta}) - M(\bar{\theta}'))(\theta - \theta_0)$$

for some mean value functions $\bar{\theta}, \bar{\theta}'$ that depend on (i, a_i, x) . Note that $\frac{\partial}{\partial \theta} M(\theta)$ vanishes when $\theta = \theta_0$ under A1(ii). For $\frac{\partial}{\partial \theta} M_N(\theta_0)$, we have

$$\begin{aligned} & \frac{\partial}{\partial \theta} M_N(\theta_0) \\ &= 2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial}{\partial \theta} \widehat{F}_{i, \theta_0}(a_i | x) (\widehat{F}_{i, \theta_0}(a_i | x) - \widehat{F}_i(a_i | x)) \mu_{i, x}(da_i) \\ &= 2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial}{\partial \theta} F_{i, \theta_0}(a_i | x) (\widehat{F}_{i, \theta_0}(a_i | x) - \widehat{F}_i(a_i | x)) \mu_{i, x}(da_i) + o_p\left(\frac{1}{\sqrt{N}}\right), \end{aligned}$$

where the second equality follows from the finiteness of $\{\mu_{i, x}\}_{i \in \mathcal{I}, x \in X}$ and A2(ii), (iv), and (ix). Importantly, by A2(ix) and the continuous mapping theorem, $\sqrt{N} \frac{\partial}{\partial \theta} M_N(\theta_0) \Rightarrow N(0, \mathcal{V})$, where \mathcal{V} is defined in (17). For the Hessians of M_N and M ,

$$\begin{aligned} \frac{\partial^2}{\partial \theta \partial \theta^\top} M_N(\theta) &= \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial^2}{\partial \theta \partial \theta^\top} \widehat{F}_{i, \theta}(a_i | x) (\widehat{F}_{i, \theta}(a_i | x) - \widehat{F}_i(a_i | x)) \mu_{i, x}(da_i) \\ &\quad + \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial}{\partial \theta} \widehat{F}_{i, \theta}(a_i | x) \frac{\partial}{\partial \theta^\top} \widehat{F}_{i, \theta}(a_i | x) \mu_{i, x}(da_i), \\ \frac{\partial^2}{\partial \theta \partial \theta^\top} M(\theta) &= \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial^2}{\partial \theta \partial \theta^\top} F_{i, \theta}(a_i | x) (F_{i, \theta}(a_i | x) - F_i(a_i | x)) \mu_{i, x}(da_i) \\ &\quad + \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial}{\partial \theta} F_{i, \theta}(a_i | x) \frac{\partial}{\partial \theta^\top} F_{i, \theta}(a_i | x) \mu_{i, x}(da_i). \end{aligned}$$

By repeated applications of the triangle inequality, and making use of A2(ii), (iv), and (v), it is straightforward to show that $|\frac{\partial^2}{\partial \theta_l \partial \theta_{l'}} (M_N(\bar{\theta}) - M(\bar{\theta}'))| = o_p(1)$ for all (l, l') as $\|\theta - \theta_0\| \rightarrow 0$. Therefore, we have

$$E_1(\theta) = (\theta - \theta_0)^\top \frac{\partial}{\partial \theta} M_N(\theta_0) + o_p(\|\theta - \theta_0\|^2).$$

Let $\xi(\theta) = \widehat{M}_N(\theta) - M_N(\theta)$, so that $E_2(\theta) = \xi(\theta) - \xi(\theta_0)$. From the definitions of \widehat{M}_N and M_N ,

$$\begin{aligned} \xi(\theta) &= \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x)) \\ &\quad \times (\widetilde{F}_{i,\theta}(a_i|x) + \widehat{F}_{i,\theta}(a_i|x) - 2\widehat{F}_i(a_i|x)) \mu_{i,x}(da_i). \end{aligned}$$

By repeatedly adding nulls, we can write

$$\xi(\theta) = \xi_1(\theta) + \xi_2(\theta) + \xi_3(\theta) + \xi_4(\theta), \quad \text{where}$$

$$\xi_1(\theta) = \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x))^2 \mu_{i,x}(da_i),$$

$$\xi_2(\theta) = 2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x)) (\widehat{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta_0}(a_i|x)) \mu_{i,x}(da_i),$$

$$\xi_3(\theta) = 2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x)) (\widehat{F}_{i,\theta_0}(a_i|x) - F_{i,\theta_0}(a_i|x)) \mu_{i,x}(da_i),$$

$$\xi_4(\theta) = -2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x)) (\widehat{F}_i(a_i|x) - F_i(a_i|x)) \mu_{i,x}(da_i).$$

In sum, $\xi(\theta)$ is $o_p(N^{-1/2}\|\theta - \theta_0\| + N^{-1})$ since $\xi_1(\theta)$ is $o_p(N^{-1})$ by A2(vi); $\xi_2(\theta)$ is $o_p(N^{-1/2}\|\theta - \theta_0\|)$, using a mean value expansion in θ and then applying A2(ii), (iv), and (vi); $\xi_3(\theta)$ is $o_p(N^{-1})$ by A2(iv) and (vii); $\xi_4(\theta)$ is $o_p(N^{-1})$ by A2(vi) and (viii). Therefore, $E_2(\theta) = o_p(N^{-1/2}\|\theta - \theta_0\| + N^{-1})$. Thus, condition (SA3) is satisfied uniformly for $\|\theta - \theta_0\| \leq \delta_N$ with $C_N = \frac{\partial}{\partial \theta} M_N(\theta_0)$. Since $\frac{\partial^2}{\partial \theta \partial \theta} M(\theta_0)$ equals \mathcal{W} (defined in equation (18)), the desired limiting distribution of $\widehat{\theta}$ follows from applying Theorem 7.1 of Newey and McFadden (1994). \square

REFERENCES

- Aguirregabiria, V. and P. Mira (2007), “Sequential estimation of dynamic discrete games.” *Econometrica*, 75 (1), 1–53. [4]
- Bajari, P., C. L. Benkard, and J. Levin (2007), “Estimating dynamic models of imperfect competition.” *Econometrica*, 75, 1331–1370. [1]
- Domínguez, M. A. and I. N. Lobato (2004), “Consistent estimation of models defined by conditional moment restrictions.” *Econometrica*, 72, 1601–1615. [1, 6]
- Hotz, V. and R. A. Miller (1993), “Conditional choice probabilities and the estimation of dynamic models.” *Review of Economic Studies*, 60, 497–531. [5]
- Newey, W. K. and D. McFadden (1994), “Large sample estimation and hypothesis testing.” In *Handbook of Econometrics*, Vol. 4 (R. F. Engle and D. McFadden, eds.), North-Holland, Amsterdam. [7, 9]

Pesendorfer, M. and P. Schmidt-Dengler (2008), “Asymptotic least squares estimators for dynamic games.” *Review of Economic Studies*, 75, 901–928. [4]

Submitted April, 2012. Final version accepted October, 2012.