

Minimum distance estimators for dynamic games

SORAWOOT SRISUMA

School of Economics, University of Surrey

We develop a minimum distance estimator for dynamic games of incomplete information. We take a two-step approach, following Hotz and Miller (1993), based on the pseudo-model that does not solve the dynamic equilibrium so as to circumvent the potential indeterminacy issues associated with multiple equilibria. The class of games estimable by our methodology includes the familiar discrete unordered action games as well as games where players' actions are monotone (discrete, continuous, or mixed) in their private values. We also provide conditions for the existence of pure strategy Markov perfect equilibria in monotone action games under increasing differences condition.

KEYWORDS. Dynamic games, Markov perfect equilibrium, semiparametric estimation with nonsmooth objective functions.

JEL CLASSIFICATION. C13, C14, C15, C51.

1. INTRODUCTION

We propose a new estimator for a class of dynamic games of incomplete information that builds on the Markov discrete decision framework reviewed in Rust (1994). Our estimator adds to a growing list of methodologies to analyze empirical games discussed in the surveys of Akerberg et al. (2005) and Aguirregabiria and Mira (2010). Two well known obstacles to structural estimation of dynamic games arise from multiple equilibria and the computational value functions that represent future expected returns. More specifically, for each structural parameter, the model may have nonunique equilibria that predict different distributions of actions and even when there are no issues of equilibrium selection, it is numerically demanding to evaluate the value functions that are defined as fixed points of some nonlinear functional equations. We take a two-step approach that does not solve out the full dynamic optimization problem and is designed to circumvent these issues.

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

This paper was previously circulated under the title “Minimum Distance Estimation for a Class of Markov Decision Processes.” It is a significant revision of the second chapter of my doctoral thesis. I am most grateful to my advisor, Oliver Linton, for his encouragement and guidance. I am thankful for many comments and suggestions from a co-editor and three anonymous referees who greatly helped improve the paper. I also wish to thank Guilherme Carmona, Joachim Groeger, Emmanuel Guerre, James Heckman, Javier Hidalgo, Stefan Hoderlein, Tatiana Komarova, Arthur Lewbel, Martin Pesendorfer, Carlos Santos, Marcia Schafgans, Philipp Schmidt-Dengler, Myung Hwan Seo, and seminar participants at numerous conferences and universities for valuable comments and suggestions. This research is supported by the ERC and ESRC.

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

DOI: [10.3982/QE266](https://doi.org/10.3982/QE266)

We begin with an assumption that pure strategy Markov perfect equilibria exist and data are generated from a single equilibrium. Most two-step estimators in the literature, following Hotz and Miller's (1993) work in a single agent discrete choice problem, consider the pseudo-model where the intractable value functions are replaced by easy to compute policy value functions that can be constructed using beliefs observed from the data. Each player's pseudo-decision problem can then be interpreted as playing a single stage game against nature. When the pseudo-decision problem has a unique solution almost surely, each player's best response is a pure strategy so that any candidate structural parameter is mapped into an implied distribution function that defines a complete pseudo-model (as opposed to *incomplete models*, for instance, see Tamer (2003)). Conditions for the existence of Markov perfect equilibria, as well as the uniqueness of the solution to pseudo-decision problems, have been established for games where players actions are modeled to be (unordered) discrete and players' private values enter the payoff functions additively; see Aguirregabiria and Mira (2007, hereafter AM), Bajari et al. (2009), and Pesendorfer and Schmidt-Dengler (2008, hereafter PSD).¹ In an independent work, Schrimpf (2011) also recently proposed an estimator for continuous action games. While the aforementioned papers make use of the pseudo-decision problem and focus on games with a single type of actions, Bajari et al. (2007, hereafter BBL) took a different approach, using forward simulation, that can handle models with both discrete and/or continuous decisions. BBL's methodology is versatile; in particular, it has been applied to model games where players' actions are monotone in their private values; for some examples, see Gowrisankaran et al. (2010), Ryan (2012), and Santos (2010).

The main contribution of this paper is to provide an alternative estimator for a large class of games that includes the models considered in BBL and their subsequent applications. A distinctive feature of BBL's methodology is the use of inequality restrictions to construct objective functions. Since little guidance on how to select inequalities exists, we show that some popular classes of inequalities can lead to objective functions that do not have unique (minimizing) solutions as the sample size tends to infinity, even when the underlying model is actually point-identified. Our estimator is obtained by minimizing the distance between distributions of actions observed from the data and predicted by the pseudo-model. We provide a set of conditions to ensure our estimator is consistent and asymptotically normal.

We also contribute by providing important foundations for the modeling of games where players play monotone strategies. The existence of pure strategy Markov equilibria is often assumed in dynamic games where players employ monotone strategies with respect to their private information; for examples, see BBL, Gowrisankaran et al. (2010) (ordered discrete action), and Schrimpf (2011) (continuous action). We provide primitive conditions based on increasing differences that ensure monotone pure strategy Markov equilibria exist for dynamic games when the action variable can be discrete, continuous, or a mixture of both. We also show that the same conditions are sufficient for each player's best response to the pseudo-decision problem to be a pure strategy almost surely. Therefore, the pseudo-model can bypass the issues associated with multiple equilibria for this class of games.

¹Bajari et al. (2009) also considered a one-step estimator.

BBL defined their estimator using a system of moment inequality restrictions implied by the equilibrium condition. Their estimator satisfies a necessary condition of an equilibrium that the implied expected return from the optimal strategy is at least as large as the returns from employing alternative strategies, where each alternative strategy is represented by an inequality. To give an intuition of why inequality selection may have a nontrivial implication, suppose the parameter of interest is uniquely identified by the inequality restrictions implied by the equilibrium. However, the equilibrium imposes that inequalities must hold for *all alternative strategies*. If we restrict our attention to certain classes of inequalities, for example, additive or multiplicative perturbations, these inequalities may not be able to identify the parameter of interest in the sense that there are other elements in the parameter space that also satisfy these less restrictive sets of inequality restrictions. Our comment is closely related to the general issue of consistent estimation in conditional moment models. Particularly, in a familiar instrumental variable framework, Domínguez and Lobato (2004) provided explicit examples when there is a unique value in the parameter space that satisfies a conditional moment (equality) restriction, but the uniqueness is lost when the conditional moment is converted into a finite number of unconditional moments. Domínguez and Lobato (2004) and Khan and Tamer (2009) also showed how to construct objective functions that preserve the identifying information content of conditional moment models commonly used in economics, with equality and inequality restrictions, respectively. However, their techniques are not applicable to BBL's estimation methodology. We show that the loss of identifying information associated with BBL's inequality selection problem can occur even without any conditioning variable.

Our estimator is motivated by a characterization of a Markov perfect equilibrium as fixed points of an operator that maps beliefs into distributions of best responses. Thus, our construction of the pseudo-model can be seen as a generalization of AM and PSD, who provided analogous characterizations for unordered discrete games that also play central roles in their estimation methodologies. We show that the game they considered is included in our general setup. We define a class of minimum distance estimators from the characterization of the equilibrium. Our estimation methodology proceeds in two stages. In the first stage, we use the distributions of actions from the data as the nonparametric beliefs to simulate the distributions of the pseudo-model implied best responses. We then compare the simulated distributions with the nonparametric distributions in the second stage by minimizing some L^2 distance.

We prove our equilibrium existence results by closely following the arguments in Athey (2001), who showed that pure strategy equilibria exist for static games of incomplete information under single crossing conditions. Athey's results are amenable to the dynamic games we consider once we restrict ourselves to players playing stationary Markov strategies. The existence of Markov equilibria in other related games can be found in AM and PSD for a class of unordered discrete action games, and in Doraszelski and Satterthwaite (2010) for games with entry/exit decisions with investment decisions.

Throughout the paper, we treat the transition law of the observed states nonparametrically since the transition law is a model primitive about which we often have little

information. We also maintain a common assumption in this literature that the observable states take finitely many values. Therefore, the estimation problem is semiparametric when the action variable is continuous. The effective rate of convergence of the nonparametric estimator in our methodology is determined by a one-dimensional object, which is consistent with the nature of a simultaneous-move game where each player forms an expectation by conditioning only on her action. Therefore, our proposed estimator does not suffer from the nonparametric curse of dimensionality with respect to the number of players. This is in contrast to extending the forward simulation method of BBL (Step 3, p. 1343) to estimate a semiparametric model, where future states are drawn conditionally on the actions of *all* players.² We note that it is also possible to extend our estimation procedure to allow for continuous states, as illustrated by Srisuma and Linton (2012) when action is discrete, although this may be of limited practical interest when the action is also continuous.

The rest of the paper proceeds as follows. Section 2 introduces the class of games that are estimable by our two-step approach. We provide the details of our methodology in Section 3. A general large sample theory is given in Section 4. Section 5 reports results from Monte Carlo studies, where we also consider the performance of BBL estimators when the objective functions used cannot identify the parameter of interest in the limit. Section 6 concludes. The Appendices are available in a supplementary file on the journal website, <http://qeconomics.org/supp/266/supplement.pdf>. Appendix A concerns the issue of consistent estimation using the BBL methodology; it contains three parts (A.1–A.3). In Appendix A.1, we give two examples where the inequality restrictions imposed by the equilibrium are satisfied by a unique element in the parameter space, but the uniqueness is lost when some well known subclasses of all inequalities are considered. In Appendix A.2, we show that a simple class of inequalities can be used to construct objective functions that preserve the identifying information from the equilibrium in discrete action games where players' best responses are characterized by some cutoff rules, that is, by choosing alternative strategies based on perturbing the cutoff values only in the first period. The suggested inequalities are applicable for unordered and ordered action games. Appendix A.3 provides some additional discussion. Appendix B contains proofs of the theorems.

2. MARKOVIAN GAMES

This section introduces the class of estimable games for our methodology. We begin by describing the elements of the general model and defining the equilibrium concept. We then consider the players' decision problems and show that when players' best responses to any Markovian beliefs are pure strategies almost surely, then the equilibrium can be characterized by a fixed point of an operator that maps beliefs into distributions of best responses. We end the section by providing examples of Markovian games that have been used in the literature. In particular, we study in detail the games where payoffs satisfy an increasing differences condition.

²BBL only considered a fully parametric estimation framework.

2.1 Model

We consider a dynamic game with I players, indexed by $i \in \mathcal{I} = \{1, \dots, I\}$, over an infinite time horizon. The elements of the game in each period are as follows.

ACTIONS. We denote the action variable for player i by $a_{it} \in A_i$. Let $\mathbf{a}_t = (a_{1t}, \dots, a_{It}) \in A = A_1 \times \dots \times A_I$. We will also occasionally abuse notation and write $\mathbf{a}_t = (a_{it}, \mathbf{a}_{-it})$, where $\mathbf{a}_{-it} = (a_{1t}, \dots, a_{i-1t}, a_{i+1t}, \dots, a_{It}) \in A_{-i} = A \setminus A_i$.

STATES. Player i 's information set is represented by the state variables $s_{it} \in S_i$, where $s_{it} = (x_{it}, \varepsilon_{it})$ such that $x_{it} \in X_i$ is common knowledge to all players and $\varepsilon_{it} \in \mathcal{E}_i$ denotes private information only observed by player i . For notational simplicity, we set $x_{it} = x_t$ for all i ; this is without any loss of generality as we can define $x_t = (x_{1t}, \dots, x_{It}) \in X$. We use s_i and (x, ε_i) interchangeably. We define $(\mathbf{s}_t, \mathbf{s}_{-it}, \boldsymbol{\varepsilon}_t, \boldsymbol{\varepsilon}_{-it}, \mathcal{E})$ analogously to $(\mathbf{a}_t, \mathbf{a}_{-it}, A)$ and denote the support of \mathbf{s}_t by $S = X \times \mathcal{E}$.

STATE TRANSITION. Future states are uncertain. Players' actions and states today affect future states. The evolution of the states is summarize by a Markov transition law $P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$.

PER PERIOD PAYOFF FUNCTIONS. Each player has a payoff function, $u_i: A \times S_i \rightarrow \mathbb{R}$, that is time separable. The payoff function for player i can depend generally on $(\mathbf{a}_t, x_t, \varepsilon_{it})$, but not directly on $\boldsymbol{\varepsilon}_{-it}$.

DISCOUNTING FACTOR. Future period's payoffs are discounted at the rate $\beta_i \in (0, 1)$ for each player.

Every period, all players observe their state variables and then they choose their actions simultaneously. We consider a Markovian framework where players' behavior is stationary across time and players are assumed to play pure strategies. More specifically, for some $\alpha_i: S_i \rightarrow A_i$, $a_{it} = \alpha_i(s_{it})$ for all i, t , so that whenever $s_{it} = s_{i\tau}$, then $\alpha_i(s_{it}) = \alpha_i(s_{i\tau})$ for any τ . Next, we introduce three modeling assumptions that are assumed to hold throughout the paper.

ASSUMPTION M1 (Conditional Independence). *The transitional distribution of the states has the factorization $P(x_{t+1}, \boldsymbol{\varepsilon}_{t+1} | x_t, \boldsymbol{\varepsilon}_t, \mathbf{a}_t) = Q(\boldsymbol{\varepsilon}_{t+1})G(x_{t+1} | x_t, \mathbf{a}_t)$, where Q is the cumulative distribution function of $\boldsymbol{\varepsilon}_t$ and G denotes the transition law of x_{t+1} conditioning on \mathbf{a}_t and x_t .*

ASSUMPTION M2 (Independent Private Values). *The private information is independently distributed across players, that is, $Q(\boldsymbol{\varepsilon}) = \prod_{i=1}^I Q_i(\varepsilon_i)$, where Q_i denotes the cumulative distribution function of ε_{it} .*

ASSUMPTION M3 (Discrete Public Values). *The support of x_t is finite such that $X = \{x^1, \dots, x^J\}$ for some $J < \infty$.*

Assumptions **M1** and **M2** generalize Rust's (1987) conditional independence framework to dynamic games. They are the key restrictions commonly imposed on the class of games in this literature. Assumption **M1** implies that ε_t is independent of x_t and all variables in the past, and ε_t is only correlated to x_{t+1} through the choice variables \mathbf{a}_t . It is conceptually straightforward to relax the former condition and allow for ε_t to be conditionally independent of the past given x_t , although this is rarely done in practice. Assumption **M2** rules out games with correlated private values. Assumption **M3** is a simplifying assumption that has an important practical implication, however it is not necessary for a general estimation methodology; for examples, see [Bajari et al. \(2009\)](#) and [Srisuma and Linton \(2012\)](#).

Under **M1** and **M2**, player i 's beliefs, which we denote by σ_i , are a stationary distribution of $\mathbf{a}_t = (\alpha_1(s_{1t}), \dots, \alpha_I(s_{It}))$ conditional on x_t for some pure Markov strategies $(\alpha_1, \dots, \alpha_I)$. Then following [Maskin and Tirole \(2001\)](#), we define the equilibrium concept as follows.

DEFINITION 1 (Markov Perfect Equilibrium). A collection $(\boldsymbol{\alpha}, \boldsymbol{\sigma}) = (\alpha_1, \dots, \alpha_I, \sigma_1, \dots, \sigma_I)$ is a Markov perfect equilibrium if the following statements hold:

- (i) For all i , α_i is a best response to $\boldsymbol{\alpha}_{-i}$ given the beliefs σ_i at almost all states x .
- (ii) All players use Markov strategies.
- (iii) For all i , the beliefs σ_i are consistent with the strategies $\boldsymbol{\alpha}$.

2.2 Players' decision problems

To characterize the players' optimal behaviors, we consider the decision problem faced by player i for a given σ_i : for all s_i ,

$$\max_{a_i \in A_i} \left\{ E_{\sigma_i} [u_i(a_{it}, \mathbf{a}_{-it}, s_i) | s_{it} = s_i, a_{it} = a_i] + \beta_i E_{\sigma_i} [V_i(s_{it+1}; \sigma_i) | s_{it} = s_i, a_{it} = a_i] \right\}, \quad (1)$$

$$\text{where } V_i(s_i; \sigma_i) = \sum_{\tau=t}^{\infty} \beta^{\tau-t} E_{\sigma_i} [u_i(\mathbf{a}_\tau, s_{i\tau}) | s_{it} = s_i].$$

The subscript σ_i on the expectation operator makes explicit that present and future actions are integrated out with respect to the beliefs σ_i ; in particular, player i forms an expectation for all players' future actions including herself and for today's actions of opposing players. The function V_i is a policy value function since the expected discounted return need not be an optimal value from an optimization problem since σ_i can be any beliefs, not necessarily equilibrium beliefs. Note that the transition law for future states is completely determined by the primitives and the beliefs.³ Thus, we can interpret each player's decision problem in (1) as a single stage game against nature that is determined

³First, note that the use of Markovian beliefs imply that $\mathcal{I}(\mathbf{s}_{t+\tau}, \mathbf{a}_{t+\tau}) = \mathcal{I}(\mathbf{s}_{t+\tau})$ and $\mathcal{I}(s_{it+\tau}, a_{it+\tau}) = \mathcal{I}(s_{it+\tau})$, where $\mathcal{I}(\cdot)$ denotes the information set of (\cdot) . For some random vectors X and Y , let $f_{X,Y}$ and $f_{X|Y}$ denote the joint density of (X, Y) and X given Y , respectively (components of X and Y can be either

by Markov beliefs. Clearly, any strategy profile that solves the decision problems for all i , and is consistent with the beliefs satisfies conditions in Definition 1 and is an equilibrium strategy. To avoid multiple predictions of best responses, the class of games estimable by our methodology requires (1) to have a unique solution almost surely. In this subsection, we show that Markov equilibrium can be represented by a fixed point of a particular mapping when the solution to the decision problem exists and is unique.

First we simplify the objective function of the decision problem by incorporating our modeling assumptions. It is convenient to write V_i recursively as

$$V_i(s_i; \sigma_i) = E_{\sigma_i}[u_i(\mathbf{a}_t, s_{it}) | s_{it} = s_i] + \beta_i E_{\sigma_i}[V_i(s_{it+1}; \sigma_i) | s_{it} = s_i].$$

The ex ante value function can be obtained by taking the conditional expectation of V_i with respect to x_t :

$$E_{\sigma_i}[V_i(s_{it}; \sigma_i) | x_t] = E_{\sigma_i}[u_i(\mathbf{a}_t, s_{it}) | x_t] + \beta_i E_{\sigma_i}[V_i(s_{it+1}; \sigma_i) | x_t].$$

Under M1, by the law of iterated expectation, $E_{\sigma_i}[V_i(s_{it+1}; \sigma_i) | x_t] = E_{\sigma_i}[E_{\sigma_i}[V_i(s_{it+1}; \sigma_i) | x_{t+1}] | x_t]$, so that the ex ante value can be written as a solution to the linear equation

$$m_i(\sigma_i) = r_i(\sigma_i) + \mathcal{L}_{i, \sigma_i} m_i(\sigma_i),$$

where $m_i(\sigma_i) = E_{\sigma_i}[V_i(s_{it}; \sigma_i) | x_t = \cdot]$, $r_i(\sigma_i) = E_{\sigma_i}[u_i(\mathbf{a}_t, s_{it}) | x_t = \cdot]$, and $\mathcal{L}_{i, \sigma_i}$ is a conditional expectation operator so that $\mathcal{L}_{i, \sigma_i} \phi = \beta_i E_{\sigma_i}[\phi(x_{t+1}) | x_t = \cdot]$ for any $\phi: X \rightarrow \mathbb{R}$. Note that $m_i(\sigma_i)$ exists and is unique under great generality since $\mathcal{L}_{i, \sigma_i}$ is typically a contraction map.⁴ Also, under M1, the choice-specific expected future return under beliefs σ_i satisfies $E_{\sigma_i}[V_i(s_{it+1}; \sigma_i) | s_{it}, a_{it}] = E_{\sigma_i}[E_{\sigma_i}[V_i(s_{it+1}; \sigma_i) | x_{t+1}] | x_t, a_{it}]$, which can be represented by $g_i(\sigma_i)$ so that

$$g_i(\sigma_i) = \mathcal{H}_{i, \sigma_i} m_i(\sigma_i),$$

continuous, discrete, or a mixture). Then, for a one-step-ahead transition, by M1,

$$\begin{aligned} f_{\mathbf{s}_{t+1} | s_{it}, a_{it}} &= f_{x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_{it}, a_{it}} \\ &= f_{\epsilon_{t+1}} f_{x_{t+1} | x_t, a_{it}}, \end{aligned}$$

where $f_{\epsilon_{t+1}}$ and $f_{x_{t+1} | x_t, a_{it}}$ can be deduced from the model primitives given any beliefs σ_i . For two periods ahead, note that $f_{\mathbf{s}_{t+2} | s_{it}, a_{it}} = \int f_{\mathbf{s}_{t+2}, \mathbf{s}_{t+1} | s_{it}, a_{it}} d\mathbf{s}_{t+1}$, using the same line of arguments as above:

$$\begin{aligned} f_{\mathbf{s}_{t+2}, \mathbf{s}_{t+1} | s_{it}, a_{it}} &= f_{\mathbf{s}_{t+2} | \mathbf{s}_{t+1}, s_{it}, a_{it}} f_{\mathbf{s}_{t+1} | s_{it}, a_{it}} \\ &= f_{\epsilon_{t+2}} f_{x_{t+2} | x_{t+1}, \mathbf{a}_{t+1}} f_{\epsilon_{t+1}} f_{x_{t+1} | x_t, a_{it}}. \end{aligned}$$

Similar arguments can be applied recursively for any future periods.

⁴Let X be some compact subset of \mathbb{R}^{L_X} and let B be a space of bounded real-valued functions defined on X . Consider a Banach space $(B, \|\cdot\|)$ equipped with the sup-norm, that is, $\|\phi\| = \sup_{x \in X} |\phi(x)|$ for any $\phi \in B$. For any $x \in X$, $\mathcal{L}_{i, \sigma_i} \phi(x) = \beta_i E_{\sigma_i}[\phi(x_{t+1}) | x_t = x]$, then it follows that $|\mathcal{L}_{i, \sigma_i} \phi(x)| \leq \beta_i \sup_{x \in X} |\phi(x)|$. In other words, $\|\mathcal{L}_{i, \sigma_i} \phi\| \leq \beta_i \|\phi\|$, hence the operator norm $\|\mathcal{L}_{i, \sigma_i}\|$ is bounded above by β_i . Since $\beta_i \in (0, 1)$, $\mathcal{L}_{i, \sigma_i}$ is a contraction. Therefore, the inverse of $I - \mathcal{L}_{i, \sigma_i}$ exists. Furthermore, it is a linear bounded operator and admits a Neumann series representation $\sum_{\tau=0}^{\infty} \mathcal{L}_{i, \sigma_i}^{\tau}$ (see Kreyszig (1989)).

where \mathcal{H}_{i,σ_i} is a conditional expectation operator so that $\mathcal{H}_{i,\sigma_i}\phi = E_{\sigma_i}[\phi(x_{t+1})|x_t = \cdot, a_{it} = \cdot]$ for any $\phi: X \rightarrow \mathbb{R}$. Since, under **M1** and **M2**, a_{it} and ε_{it} has no additional information on \mathbf{a}_{-it} given x_t , then the objective function in (1), which we henceforth denote by Λ_i , can be written as

$$\Lambda_i(a_i, s_i; \sigma) = E_{\sigma_i}[u_i(a_i, \mathbf{a}_{-it}, x_t, \varepsilon_i)|x_t = x] + \beta_i g_i(a_i, x; \sigma).$$

The corresponding set of best responses is defined as

$$\text{BR}_i(s_i; \sigma_i) = \{a_i \in A_i : \Lambda_i(a_i, s_i; \sigma_i) \geq \Lambda_i(a'_i, s_i; \sigma_i) \text{ for all } a'_i \in A_i\}.$$

A pure strategy best response is a particular selection from the best response that satisfies $\alpha_i(\cdot; \sigma_i) \in \text{BR}_i(\cdot; \sigma_i)$, that is, for all s_i ,

$$\Lambda_i(\alpha_i(s_i; \sigma_i), s_i; \sigma_i) \geq \Lambda_i(a'_i, s_i; \sigma_i) \quad \text{for } a'_i \in A_i. \quad (2)$$

Since we assume that $\text{BR}_i(s_i; \sigma_i)$ is a singleton for all s_i, σ_i , there is no need for a selection from the best response set. Thus, there is a single-valued map Ψ_i such that

$$F_i = \Psi_i(\sigma_i), \quad \text{where } F_i(a_i|x; \sigma_i) = \Pr[\alpha_i(s_{it}; \sigma_i) \leq a_i|x_t = x] \text{ for all } a_i, x. \quad (3)$$

Under independence (Assumption **M2**), information on all marginal distributions of actions provides equivalent information for the joint distribution of actions, so that any equilibrium beliefs must satisfy condition (2) and the beliefs are consistent with the actions according to (3), where each σ_i can be represented by $\prod_{l=1}^I F_l = \prod_{l=1}^I \Psi_l(\sigma_l)$ for all i . We can, therefore, summarize the necessary condition that the equilibrium beliefs must satisfy by a fixed point of a map Ψ that takes any vector $\mathbf{F} = (F_1, \dots, F_I)$ into $\Psi(\mathbf{F}) = (\Psi_1(\prod_{l=1}^I F_l), \dots, \Psi_I(\prod_{l=1}^I F_l))$, that is, the condition

$$\mathbf{F} = \Psi(\mathbf{F}). \quad (4)$$

The fixed point of Ψ fully characterizes the equilibrium since any \mathbf{F} that satisfies equation (4) can be extended to construct a Markov perfect equilibrium, as $\alpha_i(s_i; \prod_{l=1}^I F_l) = \arg \max_{a_i \in A_i} \Lambda_i(a_i, s_i; \prod_{l=1}^I F_l)$ is the best response that is consistent with the beliefs by construction.

Equation (4) forms the basis of our minimum distance estimator, where, in Section 3, we look to minimize the distance between the distribution of actions from the data and the implied distribution generated by the empirical version of $\Psi(\mathbf{F})$. The characterization of an equilibrium as a fixed point to equation (4) is very similar to the approach taken by AM (Representation Lemma) and PSD (Proposition 1), who considered a particular class of unordered discrete choice game (see Assumption **D** below).⁵

⁵Equation (4) can also be useful for proving the existence of a Markov perfect equilibrium when Ψ is known to satisfy regularity conditions to ensure that a fixed point exists, as well as providing an alternative numerical calculation of equilibrium probabilities; see [Pesendorfer and Schmidt-Dengler \(2008\)](#) for further discussions.

2.3 Games under increasing differences

In many economic applications it is natural to model players' best responses to be monotone in their private values. The action space can be finite, for example, in investment models where firms purchase or rent goods in discrete units, or the action variable can have a continuous contribution (with or without a discrete component), as in the traditional investment and pricing models. The source of the monotonicity can often be derived from an intuitive restriction imposed on the interim payoff differences when player i chooses action a_i over a'_i , which we denote by $\Delta u_i(a_i, a'_i, \mathbf{a}_{-i}, x, \varepsilon_i) = u_i(a_i, \mathbf{a}_{-i}, x, \varepsilon_i) - u_i(a'_i, \mathbf{a}_{-i}, x, \varepsilon_i)$, that increases with ε_i . Increasing differences have numerous applications in economics; see the monograph by Topkis (1998) for examples. We consider games that satisfy the following conditions.

ASSUMPTION S1 (Increasing Differences). *For any $a_i > a'_i$ and $\varepsilon_i > \varepsilon'_i$, $\Delta u_i(a_i, a'_i, \mathbf{a}_{-i}, x, \varepsilon_i) > \Delta u_i(a_i, a'_i, \mathbf{a}_{-i}, x, \varepsilon'_i)$ for all i, \mathbf{a}_{-i}, x .*

ASSUMPTION S2. *The distribution of ε_{it} is absolutely continuous with respect to the Lebesgue measure with a bounded density on $\mathcal{E}_i = [\underline{\varepsilon}_i, \bar{\varepsilon}_i] \subset \mathbb{R}$ for all i .*

Assumptions S1 and S2 are versions of the conditions used in Athey (2001) to study the equilibrium properties in static games. Importantly, increasing differences of u_i in (a_i, ε_i) imply that the incremental return satisfies the single crossing condition in (a_i, ε_i) (see Definition 1 in Athey). Our increasing differences condition is strict and holds uniformly over (\mathbf{a}_{-i}, x) , which, although generally not necessary for pure strategy equilibria to exist, will be convenient for modeling games where players employ pure strategies almost surely. When u_i is differentiable in (a_i, ε_i) , the increasing differences condition has a simple characterization: $\frac{\partial^2}{\partial a_i \partial \varepsilon_i} u_i(a_i, \mathbf{a}_{-i}, x, \varepsilon_i) > 0$ for all \mathbf{a}_{-i}, x . We also comment that compactness of \mathcal{E}_i is assumed here only for the purpose of establishing the existence of equilibria. In an econometric application, \mathcal{E}_i can have full support on \mathbb{R} . Next, we show that existence theorems for equilibria in static games under the single crossing condition of Athey (2001) can be applied to our dynamic games.

For the first case, we restrict the support of the action variable to be discrete and impose an integrability condition.

ASSUMPTION S3. *The variable A_i is finite for all i and $\int |u_i(a_i, \mathbf{a}_{-i}, x, \varepsilon_i)| dQ_i(\varepsilon_i) < \infty$ for all $i, a_i, \mathbf{a}_{-i}, x$.*

Under Assumptions M1, M2, M3, S2, and S3, all expected returns, particularly Λ_i , exist and $\text{BR}_i(s_i; \sigma_i)$ is nonempty by the finiteness of A_i for all s_i, σ_i . Let $\Delta \Lambda_i(a_i, a'_i, s_i; \sigma_i) = \Lambda_i(a_i, s_i; \sigma_i) - \Lambda_i(a'_i, s_i; \sigma_i)$. Then we have the following results.

LEMMA 1 (Increasing Differences in Expected Returns). *Under M1, M2, M3, S1, S2, and S3, for any $a_i > a'_i$ and $\varepsilon_i > \varepsilon'_i$, $\Delta \Lambda_i(a_i, a'_i, x, \varepsilon_i; \sigma_i) > \Delta \Lambda_i(a_i, a'_i, x, \varepsilon'_i; \sigma_i)$ for all i, x, σ_i .*

PROOF. Under **M1** and **M2**, $g_i(\sigma_i)$ does not depend on ε_i . Therefore, we have

$$\begin{aligned} & \Delta\Lambda_i(a_i, a'_i, x, \varepsilon_i; \sigma_i) - \Delta\Lambda_i(a_i, a'_i, x, \varepsilon'_i; \sigma_i) \\ &= E_{\sigma_i}[\Delta u_i(a_i, a'_i, \mathbf{a}_{-it}, x_t, \varepsilon_i) - \Delta u_i(a_i, a'_i, \mathbf{a}_{-it}, x_t, \varepsilon'_i) | x_t = x] \\ &> 0, \end{aligned}$$

where the inequality follows from Assumption **S1**. □

LEMMA 2 (Pure Strategy Best Response). *Under **M1**, **M2**, **M3**, **S1**, **S2**, and **S3**, $\text{BR}_i(s_{it}; \sigma_i)$ is a singleton set almost surely for all i , σ_i .*

PROOF. For any σ_i , let $\alpha_i(\cdot; \sigma_i)$ and $\alpha'_i(\cdot; \sigma_i)$ denote distinct selections from $\text{BR}_i(\cdot; \sigma_i)$ so that for some x , there exists $\varepsilon_i > \varepsilon'_i$ such that (without any loss of generality) $\alpha_i(x, \varepsilon'_i; \sigma_i) > \alpha'_i(x, \varepsilon_i; \sigma_i)$. By definition of a best response, $\Delta\Lambda_i(\alpha_i(x, \varepsilon'_i; \sigma_i), \alpha'_i(x, \varepsilon_i; \sigma_i), x, \varepsilon'_i; \sigma_i) \geq 0 \geq \Delta\Lambda_i(\alpha_i(x, \varepsilon'_i; \sigma_i), \alpha'_i(x, \varepsilon_i; \sigma_i), x, \varepsilon_i; \sigma_i)$. However, this contradicts the strict increasing difference condition in the expected returns (Lemma 1). □

Notice that finiteness of A_i does not play any role in proving Lemmas 1 and 2 beyond ensuring A_i exists and BR_i is nonempty. An implication of Lemma 1 is that every selection from $\text{BR}_i(\cdot; \sigma_i)$ is nondecreasing in ε_i for all i , x , σ_i (by the monotone selection theorem of Milgrom and Shannon (1994, Theorem 4)). Together with Lemma 2, they ensure that, for any given beliefs, each player's best response is a monotone pure strategy almost surely. The existence of an equilibrium then follows immediately from results developed in Athey (2001).

PROPOSITION 1. *Assume **M1**, **M2**, **M3**, **S1**, **S2**, and **S3**. Then a pure strategy Markov perfect equilibrium exists where each player's equilibrium strategy $\alpha_i(x, \varepsilon_i)$ is nondecreasing in ε_i for all i , x .*

PROOF. Under **S2** and **S3**, the regularity assumption A1 in Athey is satisfied with Λ_i as player's i objective function. Lemmas 1 and 2 imply that each player's best response to any Markov beliefs is a monotone pure strategy almost surely. Therefore, Λ_i satisfies the single crossing condition for games of incomplete information in Definition 3 of Athey. The proof then follows from Theorem 1 in Athey. □

Although we consider a dynamic game, by restricting the equilibrium concept to players using stationary Markov beliefs under the conditional independence and private values framework, the arguments used for static games in Athey are directly applicable.⁶ Athey also showed that finiteness of A_i can be replaced by compactness when the payoff

⁶The objective function (see the first display on p. 865) of the decision problem studied in Athey appears in a slightly different form than ours, where, instead of using a distribution of actions, she uses the strategy functions of opposing players as beliefs. However, the two approaches are analogous since any conditional distribution, σ_i , of \mathbf{a}_i given x_i uniquely determines monotone strategies $\alpha_i(\mathbf{s}_i) = (\alpha_i(s_{it}), \alpha_{-i}(\mathbf{s}_{-it}))$ for all x up to null sets on ε_i .

function is continuous in the players' actions. To apply her result in a dynamic setting, we also need to impose some continuity condition on the transition law of the states. Let $\underline{a}_i = \inf A_i$ and $\bar{a}_i = \sup A_i$, and let $G(x_{t+1}|x_t, \mathbf{a}_t)$ be the transition law of x_{t+1} conditioning on \mathbf{a}_t and x_t .

ASSUMPTION S4. For all i ,

- (i) $A_i = [\underline{a}_i, \bar{a}_i]$
- (ii) $u_i(a_i, \mathbf{a}_{-i}, x, \varepsilon_i)$ is continuous in $(a_i, \mathbf{a}_{-i}, \varepsilon_i)$ for all x
- (iii) $G(x'|x, a_i, \mathbf{a}_{-i})$ is continuous in (a_i, \mathbf{a}_{-i}) for all x, x' .

Assumptions M1, M2, M3, S2, and S4 ensure that the regularity condition in Athey (A1) is satisfied, and A_i exists and is continuous in a_i ; hence $\text{BR}_i(s_i; \sigma_i)$ is nonempty for all s_i, σ_i , since A_i is compact (Weierstrass theorem). Each player's best response for any given beliefs is also a monotone pure strategy almost surely (by replacing S3 with S4 in Lemmas 1 and 2). Then we have the following proposition.

PROPOSITION 2. Assume M1, M2, M3, S1, S2, and S4. Then a pure strategy Markov perfect equilibrium exists where each player's equilibrium strategy $\alpha_i(x, \varepsilon_i)$ is nondecreasing in ε_i for all i, x .

PROOF. Under S2 and S4, assumption A1 in Athey is satisfied with A_i as player's i objective function. It is easy to see that conditions (i)–(iii) in Theorem 2 of Athey are satisfied by our assumptions; in particular, for any finite $A'_1 \times \dots \times A'_I \subset A_1 \times \dots \times A_I$, a monotone pure (Markov) strategy exists by Proposition 1. The proof then follows from Theorem 2 in Athey (2001). \square

For modeling purposes, note that strict increasing differences do not imply that $\alpha_i(x, \varepsilon_i)$ is strictly increasing in ε_i . A sufficient condition for strict monotonicity is given by Edlin and Shannon (1998), which in our case requires that (i) $A_i(a_i, x, \varepsilon_i; \sigma_i)$ is continuously differentiable in a_i, ε_i and (ii) the best response satisfies the first order condition. Thus an intermediate case exists between purely continuous and discrete action games. For instance, when there are corner solutions, then the distribution of the action variable has both continuous and discrete components. Proposition 2 (and Theorem 2 in Athey) accommodates mass points as long as the payoff function remains continuous on the action space. However, the continuity requirement does exclude some interesting games. For example, although continuity in payoffs over opponents' mass points may be reasonable in Cournot oligopoly games, it rules out Bertrand-type pricing problems. A recent empirical study whose payoff structure satisfies the continuity requirement of Assumption S4 is the dynamic milk quota trading case in Hong and Shum (2010), where, an economic agent can have positive (negative) trade demand (supply), which is modeled continuously, or she can produce using an existing quota (mass point at zero). For further discussions of other games with discontinuities and the existence of their equilibria, see Athey (Section 4).

In this subsection, we have shown that games under increasing differences have a pure strategy equilibrium under weak primitive modeling conditions. Furthermore, Lemmas 1 and 2 show that players' decision problems also have unique solutions. The consequences of the lemmas are particularly important for inference, since analogous conditions ensure that the parameterized pseudo-decision problem gives a unique prediction of an optimal behavior almost surely. However, without further restrictions, games under increasing differences may also have multiple equilibria.⁷ In this paper, we only consider the estimation problem for games that either have a unique equilibrium or have observed data that have been generated from a single equilibrium.

2.4 Other dynamic models

Note that a single agent Markov decision problem is a special case of a game when $I = 1$, where the player's beliefs simplify to the Markov distribution of her own action given the states. Indeed, a class of popular games that is included in our general framework is built on the discrete decision problem studied in Rust (1987). These discrete games have been extensively studied in this literature (see the surveys of Akerberg et al. (2005) and Aguirregabiria and Mira (2010)) and they impose the following assumptions to model games with unordered discrete actions.

ASSUMPTION D (Discrete Choice Game). *For all i ,*

- (i) $A_i = \{0, \dots, K_i\}$.
- (ii) $\mathcal{E}_i = \mathbb{R}^{K_i+1}$ so that $\varepsilon_{it} = (\varepsilon_{it}(0), \dots, \varepsilon_{it}(K_i))$.
- (iii) *The distribution of ε_{it} is absolutely continuous with respect to the Lebesgue measure whose density is bounded on \mathcal{E}_i .*
- (iv) $u_i(a_i, \mathbf{a}_{-i}, x, \varepsilon_i) = \pi_i(a_i, \mathbf{a}_{-i}, x) + \sum_{k=0}^{K_i} \varepsilon_i(k) \mathbf{1}[a_i = k]$ for all \mathbf{a}_{-i}, x .

Under M1, M2, M3, and D, it is easy to see that the event where $\Lambda_i(a_i, s_{it}; \sigma) = \Lambda_i(a'_i, s_{it}; \sigma)$ has probability 0, so each player's best response for any given beliefs is a pure strategy almost surely; for further details, see AM and PSD, who characterized the equilibrium by choice probabilities analogous to our equation (4). Specifically, note that a vector of choice probabilities, $(P_i(0|x), \dots, P_i(K_i|x))$, is just a linear transformation of a vector of conditional distributions,

$$\begin{pmatrix} P_i(0|x) \\ \vdots \\ P_i(K_i|x) \end{pmatrix} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ -1 & \ddots & 0 & \vdots \\ \vdots & -1 & \ddots & 0 \\ 0 & \cdots & -1 & 1 \end{pmatrix} \begin{pmatrix} F_i(0|x) \\ \vdots \\ F_i(K_i|x) \end{pmatrix}, \quad (5)$$

⁷Recently, Mason and Valentinyi (2007) proposed some sufficient conditions for a unique equilibrium under increasing differences; specifically, by employing a stronger version of increasing differences and imposing a Lipschitz condition on the incremental return with respect to other players' actions.

where the transformation matrix has 1's on the main diagonal, -1 's on the subdiagonal, and 0's everywhere else.

The general model discussed in this section can also be adapted to accommodate games where players have more than one decision variable. This feature is useful for many oligopoly games, for instance, where the economic agents endogenously choose whether to participate in the market before deciding on the price or investment decisions. One can model such decision problems where players make sequential choices by combining the primitives from the games with a single action variable discussed previously; for a detailed discussion, see [Arcidiacono and Miller \(2008\)](#), [BBL](#), and [Srisuma \(2010\)](#).

3. ESTIMATION METHODOLOGY

We now parameterize $\{u_i\}_{i=1}^I$ by a finite-dimensional parameter $\theta \in \Theta \subset \mathbb{R}^p$ and update the notation for the payoff functions with $\{u_{i,\theta}\}_{i=1}^I$. We take $\{\beta_i\}_{i=1}^I$ as known. We do not impose any particular distribution on G as this is nonparametrically identified under weak regularity conditions. To keep the notation as simple as possible, we assume that the observed data are collected from games played over two periods across N markets. Specifically, we omit the time subscript and let $\{(\mathbf{a}_n, x_n, x'_n, \varepsilon_n)\}_{n=1}^N$ denote a random sample generated from a particular equilibrium when $\theta = \theta_0$, where x'_n is the only variable observed from the second period. We state this as an assumption that we maintain for the remainder of the paper.

ASSUMPTION E. *The data are generated by a Markov perfect equilibrium $(\boldsymbol{\alpha}, \boldsymbol{\sigma}) = (\alpha_1, \dots, \alpha_I, \sigma_1, \dots, \sigma_I)$ for some $\theta = \theta_0 \in \Theta$.*

The econometrician only observes $\{(\mathbf{a}_n, x_n, x'_n)\}_{n=1}^N$. The goal is to estimate θ_0 . Assumption E implies that $a_{in} = \alpha_i(x_n, \varepsilon_{in})$ for all i, n . We simply denote the conditional distribution of the equilibrium actions for each player by F_i and let $\mathbf{F} = (F_1, \dots, F_I)$, so that $\sigma_i = \prod_{l=1}^I F_l$ is the same for all i . For any $\theta \in \Theta$, we can then define the pseudo-decision problems where players use $\boldsymbol{\sigma}$ to construct the policy values. When each pseudo-decision problem has a unique solution, then there is a map, analogous to the previous section, that takes θ into $F_{i,\theta}$, the implied best response distribution of actions given σ_i . By construction, the equilibrium condition requires that $F_{i,\theta_0} = F_i$ for all i , which is the condition that motivates our minimum distance estimator. Therefore, our estimation strategy requires the construction of the distribution of the best response mapping analogous to that found in Section 2.2. Section 3.1 gives the outline of our minimum distance estimator.

We provide details regarding practical implementation in Section 3.2. The section ends with a brief discussion. In what follows, since we only consider the policy value functions and associated pseudo-decision problems generated from $\boldsymbol{\sigma}$, henceforth we suppress the dependence on beliefs.

3.1 Minimum distance approach

To formally define $F_{i,\theta}$, we need to construct the pseudo-decision problem. As in Section 2.2, we begin by incorporating Assumptions M1–M3 to simplify the nature of future expected returns under σ . The (policy) value function, here written recursively, for any θ is

$$V_{i,\theta}(s_{in}) = E[u_{i,\theta}(\mathbf{a}_{in}, s_{in})|s_{in}] + \beta_i E[V_{i,\theta}(s'_{in})|s_{in}].$$

Under M1 and M3, by the law of iterated expectation, the ex ante value, $E[V_{i,\theta}(s_{in})|x_n]$, can be written as the solution to the matrix equation

$$m_{i,\theta} = r_{i,\theta} + \mathcal{L}_i m_{i,\theta}, \quad (6)$$

where $m_{i,\theta}$ and $r_{i,\theta}$ are J -dimensional vectors whose j th entries are $m_{i,\theta}(x^j) = E[V_{i,\theta}(s_{in})|x_n = x^j]$ and $r_{i,\theta}(x^j) = E[u_{i,\theta}(\mathbf{a}_n, s_{in})|x_n = x^j]$, respectively, and \mathcal{L}_i is a J by J matrix whose (j, k) th entry is $\beta_i \times \Pr[x'_n = x^k | x_n = x^j]$. Since \mathcal{L}_i is the product between β_i and a stochastic matrix, $I - \mathcal{L}_i$ is invertible, ensuring the existence and uniqueness of $m_{i,\theta}$ for all (i, θ) .⁸ Under M1, by the law of iterated expectation, the choice-specific expected future return, $E[V_{i,\theta}(s'_{in})|x_n, a_{in}]$, is a linear transform of the ex ante value,

$$g_{i,\theta} = \mathcal{H}_i m_{i,\theta}, \quad (7)$$

where, for all (a_i, x) , $g_{i,\theta}(a_i, x) = E[V_{i,\theta}(s'_{in})|x_n = x, a_{in} = a_i]$, and $\mathcal{H}_i \phi(a_i, x) = \sum_{x' \in X} \phi(x') G_i(x'|x, a_i)$ for any $\phi: X \rightarrow \mathbb{R}$, where G_i is the transition law of x'_n conditioning on (x_n, a_{in}) . Then, under M1 and M2, the parameterized objective function for the pseudo-decision problem is given by

$$\Lambda_{i,\theta}(a_i, x, \varepsilon_i) = E[u_{i,\theta}(a_i, \mathbf{a}_{-in}, x_n, \varepsilon_i)|x_n = x] + \beta_i g_{i,\theta}(a_i, x). \quad (8)$$

For $u_{i,\theta}$ that satisfies the modeling assumptions analogous to those in Sections 2.3 and 2.4, $\Lambda_{i,\theta}(\cdot, x_n, \varepsilon_{in})$ has a unique maximizer on A_i almost surely. We denote its corresponding best response function by $\alpha_{i,\theta}$, so that

$$\alpha_{i,\theta}(x, \varepsilon_i) = \arg \max_{a_i \in A_i} \Lambda_{i,\theta}(a_i, x, \varepsilon_i). \quad (9)$$

Then the pseudo-model implied distribution function can be written as an outcome of the map (cf. equation (3))

$$F_{i,\theta} = \Psi_{i,\theta} \left(\prod_{l=1}^I F_l \right), \quad (10)$$

where $F_{i,\theta}(a_i|x) = \int \mathbf{1}[\alpha_{i,\theta}(x, \varepsilon_i) \leq a_i] dQ_i(\varepsilon_i)$ for all (a_i, x) .⁹

⁸This is a special case of footnote 2. The existence of $(I - \mathcal{L}_i)^{-1}$ can also be seen to follow directly from the dominant diagonal theorem since the sum of the (nonnegative) elements in each row of \mathcal{L}_i is $\beta_i < 1$ (Taussky (1949)).

By construction, the equilibrium condition implies that $F_{i,\theta} = F_i$ when $\theta = \theta_0$. We consider the limiting objective function that measures an L^2 distance between $F_{i,\theta}(\cdot|x)$ and $F_i(\cdot|x)$ over the support of A_i for all i and x :

$$M(\theta) = \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (F_{i,\theta}(a_i|x) - F_i(a_i|x))^2 \mu_{i,x}(da_i)$$

for some measures $\{\mu_{i,x}\}_{i=1, x=X}^{\mathcal{I}, X}$. The issues of identification and the choice of measures are discussed in Section 4. For now, we suppose $M(\theta)$ has a unique minimum at zero when $\theta = \theta_0$.

3.2 Implementation

In practice, $\Psi_{i,\theta}$ and \mathbf{F} are unknown, so we replace them by their empirical counterparts. Our estimator minimizes the sample analog of $M(\theta)$. The estimation procedure therefore proceeds in two stages. The first stage estimates the pseudo-model implied distributions. The second stage chooses θ to minimize their distance with the distribution of actions from the data. For the convenience of the reader, in Table 1 we tabulate various elements and their possible estimators from equations (8) and (10) that are used to define $F_{i,\theta}$.

The elements from the linear equations can be found in (6) and (7). We also let $E_n[\psi(w_n)|x_n = x]$ denote an empirical version of $E[\psi(w_n)|x_n = x]$ for any function ψ

TABLE 1. List of variables with definitions and some possible estimators for any i, a_i, x, x' .

Variable	Definition	Possible Estimator
<i>From the data</i>		
$p_X(x)$	$\Pr[x_n = x]$	$\widehat{p}_X(x) = \frac{1}{N} \sum_{n=1}^N \mathbf{1}[x_n = x]$
$p_{X',X}(x', x)$	$\Pr[x'_n = x', x_n = x]$	$\widehat{p}_{X',X}(x', x) = \frac{1}{N} \sum_{n=1}^N \mathbf{1}[x'_n = x', x_n = x]$
$G_i(x' x, a_i)$	$\Pr[x'_n = x' x_n = x, a_{in} = a_i]$	\widehat{G}_i depends on a_{in}
$F_i(a_i x)$	$\Pr[a_{in} \leq a_i x_n = x]$	$\widehat{F}_i(a_i x) = \frac{1}{N} \sum_{n=1}^N \mathbf{1}[a_{in} \leq a_i, x_n = x] / \widehat{p}_X(x)$
<i>Linear equations</i>		
$r_{i,\theta}(x)$	$E[u_{i,\theta}(a_{in}, \mathbf{a}_{-in}, x_n, \varepsilon_{in}) x_n = x]$	$\widehat{r}_{i,\theta}$ depends on a_{in}
$\mathcal{L}_i\phi(x)$	$\beta_i E[\phi(x'_n) x_n = x]$	see equation (12) below
$\mathcal{H}_i\phi(a_i, x)$	$E[\phi(x'_n) x_n = x, a_{in} = a_i]$	$\widehat{\mathcal{H}}_i$ depends on a_{in}
$m_{i,\theta}(x)$	$E[V_{i,\theta}(s_{in}) x_n = x]$	$\widehat{m}_{i,\theta} = (I - \widehat{\mathcal{L}}_i)^{-1} \widehat{r}_{i,\theta}$
$g_{i,\theta}(a_i, x)$	$E[V_{i,\theta}(s'_{in}) x_n = x, a_{in} = a_i]$	$\widehat{g}_{i,\theta} = \widehat{\mathcal{H}}_i (I - \widehat{\mathcal{L}}_i)^{-1} \widehat{r}_{i,\theta}$

⁹For the discrete action games considered in Section 2.4 (under Assumption D), there is no need to solve the pseudo-decision problem at all since the choice probabilities (hence distribution functions) have a one-to-one relationship with the normalized expected returns (Hotz and Miller (1993)). In particular, when the vectors of the unobserved states are also independent and identically distributed (i.i.d.) extreme values, then $F_{i,\theta}(a_i|x) - F_{i,\theta}(a_i - 1|x) = \int \mathbf{1}[\alpha_{i,\theta}(x, \varepsilon_i) = a_i] dQ_i(\varepsilon_i)$ has a closed form in the expected returns (for instance, see AM).

of w_n , which can be any vectors from the sample. In particular, since x_n is a discrete random variable, a possible candidate of $E_n[\psi(w_n)|x_n = x]$ is simply $\frac{1}{N} \sum_{n=1}^N \psi(w_n) \mathbf{1}[x_n = x] / \widehat{p}_X(x)$.

First stage distribution of best responses A feasible estimator for $F_{i,\theta}$ can be obtained by estimating $\Lambda_{i,\theta}$ and simulating ε_{in} as follows.

Step 1 Estimate the elements of $\Lambda_{i,\theta}$. From (8), let

$$\widehat{\Lambda}_{i,\theta}(a_i, x, \varepsilon_i) = E_n[u_{i,\theta}(a_i, \mathbf{a}_{-in}, x_n, \varepsilon_i)|x_n = x] + \beta_i \widehat{g}_{i,\theta}(a_i, x) \quad \text{for all } (a_i, x, \varepsilon_i).$$

Using equations (6) and (7), $g_{i,\theta}$ satisfies

$$g_{i,\theta} = \mathcal{H}_i(I - \mathcal{L}_i)^{-1} r_{i,\theta}. \tag{11}$$

Therefore, $\widehat{g}_{i,\theta}$ can be obtained from $(\widehat{r}_{i,\theta}, \widehat{\mathcal{L}}_i, \widehat{\mathcal{H}}_i)$, estimators for $(r_{i,\theta}, \mathcal{L}_i, \mathcal{H}_i)$, which we now consider.

Estimation of $r_{i,\theta}$. The estimation of $r_{i,\theta}$ is complicated by the fact that we do not observe $\{\varepsilon_{in}\}_{n=1}^N$. Estimable games in this literature impose modeling assumptions that allow $r_{i,\theta}$ to be nonparametrically identified for all θ . For examples, unordered discrete action games (under Assumption D) make use of the well known Hotz and Miller (1993) inversion theorem to identify and estimate $r_{i,\theta}$, and for games with monotone actions, the identification and estimation of $r_{i,\theta}$ rely on the quantile invariance between a_{in} and ε_{in} . To illustrate, we consider the purely continuous and discrete action cases under monotonicity.

EXAMPLE 1. Suppose $\alpha_i(x, \varepsilon_i)$ is strictly increasing in ε_i almost everywhere on \mathcal{E}_i for all i, x . Then the inverse of α_i exists and we denote it by ρ_i , which is defined by the relation $\rho_i(\alpha_i(x, \varepsilon_i), x) = \varepsilon_i$ for all i, x_i, ε_i . It follows that $F_i(a_i|x) = Q_i(\rho_i(a_i, x))$. Thus $\varepsilon_{in} = Q_i^{-1}(F_i(a_{in}|x_n))$ and we can generate the private value $\widehat{\varepsilon}_{in}$ by $Q_i^{-1}(\widehat{F}_i(a_{in}|x_n))$. Then one candidate for $\widehat{r}_{i,\theta}(x)$ is $E_n[u_{i,\theta}(\mathbf{a}_n, x_n, \widehat{\varepsilon}_{in})|x_n = x]$.

EXAMPLE 2. Suppose $\alpha_i(x, \varepsilon_i)$ is weakly increasing in ε_i almost everywhere on \mathcal{E}_i for all i, x . Let $\{a_i^k\}_{k=1}^{K_i}$ be an increasing sequence of possible actions for some $K_i < \infty$. Although the inverse of α_i does not exist, by monotonicity, we have $\mathcal{E}_i = \bigcup_{k=1}^{K_i} C_k(x)$, where $C_k(x) = [Q_i^{-1}(F_i(a_i^{k-1}|x)), Q_i^{-1}(F_i(a_i^k|x))]$ for $k > 1$. Therefore, the cutoff values where the optimal action jumps to higher actions are identified. In particular,

$$r_{i,\theta}(x) = \sum_{k=1}^{K_i} \Pr[a_{in} = a_i^k | x_n = x] \int_{C_k(x)} E[u_{i,\theta}(a_i^k, \mathbf{a}_{-in}, x_n, \varepsilon_i) | x_n = x] dQ_i(\varepsilon_i).$$

Then, for instance, we can estimate $r_{i,\theta}(x)$ by replacing $\Pr[a_{in} = a_i^k | x_n = x]$ with $\frac{1}{N} \sum_{n=1}^N \mathbf{1}[a_{in} = a_i^k, x_n = x] / \widehat{p}_X(x)$ and estimate $\int_{C_k(x)} E[u_{i,\theta}(a_i^k, \mathbf{a}_{-in}, x_n, \varepsilon_i) | x_n = x] dQ_i(\varepsilon_i)$ by replacing $\int_{\widehat{C}_k(x)} E_n[u_{i,\theta}(a_i^k, \mathbf{a}_{-in}, x_n, \varepsilon_i) | x_n = x] d\widehat{Q}_i(\varepsilon_i)$ with $\widehat{C}_k(x) = [Q_i^{-1}(\widehat{F}_i(a_i^{k-1}|x)), Q_i^{-1}(\widehat{F}_i(a_i^k|x))]$.

The mixed continuous case can also be straightforwardly dealt with by using a combination of the two techniques above, since we can write

$$r_{i,\theta}(x) = \Pr[a_{in} \in A_i^C | x_n = x] E[u_{i,\theta}(a_{in}, \mathbf{a}_{-in}, x_n, \varepsilon_{in}) | x_n = x, a_{in} \in A_i^C] \\ + \Pr[a_{in} \in A_i^D | x_n = x] E[u_{i,\theta}(a_{in}, \mathbf{a}_{-in}, x_n, \varepsilon_{in}) | x_n = x, a_{in} \in A_i^D],$$

where A_i^D denotes the support of A_i that a_{in} has positive mass points and A_i^C is the complement set of A_i^D with respect to A_i .

Estimation of \mathcal{L}_i . The variable \mathcal{L}_i can be represented by a J by J matrix of conditional probabilities. A simple estimator for \mathcal{L}_i is the frequency estimator whose (j, k) th element satisfies

$$\widehat{\mathcal{L}}_i(j, k) = \begin{cases} \beta_i \widehat{p}_{X', X}(x^k, x^j) / \widehat{p}_X(x^j), & \text{if } \widehat{p}_X(x^j) > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

An appealing feature of the frequency estimator is that $(I - \widehat{\mathcal{L}}_i)^{-1}$ necessarily exists as discussed previously.

Estimation of \mathcal{H}_i . The variable \mathcal{H}_i is a conditional expectation operator defined by G_i , the transition law of x'_n conditioning on a_{in} and x_n . The nature of the nonparametric estimator of G_i depends on whether a_{in} is continuous, discrete, or mixed. For an estimator \widehat{G}_i of G_i , $\widehat{\mathcal{H}}_i$ is defined as $\widehat{\mathcal{H}}_i \phi(a_i, x) = \sum_{x' \in X} \phi(x') \widehat{G}_i(x' | x, a_i)$ for any a_i, x and any function $\phi : X \rightarrow \mathbb{R}$.

EXAMPLE 1 (Continued). There are many nonparametric estimators that can be used to estimate a conditional expectation. One candidate is a Nadaraya–Watson type estimator, where $\widehat{G}_i(x' | x, a_i) = \frac{1}{N} \sum_{n=1}^N \mathbf{1}[x'_n = x', x_n = x] K_h(a_{in} - a_i) / \frac{1}{N} \sum_{n=1}^N \mathbf{1}[x_n = x] K_h(a_{in} - a_i)$ and $K_h(\cdot) = \frac{1}{h} K(\frac{\cdot}{h})$ denotes a user-chosen kernel and h is the bandwidth.

EXAMPLE 2 (Continued). Since all variables are discrete, we can simply use the frequency estimator $\widehat{G}_i(x' | x, a_i) = \sum_{n=1}^N \mathbf{1}[x'_n = x', x_n = x, a_{in} = a_i] / \sum_{n=1}^N \mathbf{1}[x_n = x, a_{in} = a_i]$ whenever $\sum_{n=1}^N \mathbf{1}[x_n = x, a_{in} = a_i] > 0$ and define $\widehat{G}_i(x' | x, a_i)$ to be zero otherwise.

For the mixed continuous case, a candidate for $\widehat{G}_i(x' | x, a_i)$ can be constructed in the same way as one of the two examples above, depending on whether a_i lies in the support of A_i that has positive mass.

Estimation of $\widehat{g}_{i,\theta}$. This is simply the sample analog of equation (11), that is, $\widehat{g}_{i,\theta} = \widehat{\mathcal{H}}_i(I - \widehat{\mathcal{L}}_i)^{-1} \widehat{r}_{i,\theta}$, which can be obtained following equations (6) and (7). First, for any $\widehat{r}_{i,\theta}, \widehat{m}_{i,\theta}$ can be estimated by a matrix multiplication: $\widehat{m}_{i,\theta} = (I - \widehat{\mathcal{L}}_i)^{-1} \widehat{r}_{i,\theta}$. Then, for any a_i, x , $\widehat{g}_{i,\theta}(a_i, x) = \sum_{x' \in X} \widehat{m}_{i,\theta}(x') \widehat{G}_i(x' | x, a_i)$. Note that $\widehat{\mathcal{L}}_i$ and $\widehat{\mathcal{H}}_i$ do not depend on θ .

Step 2 Estimate $F_{i,\theta}$. Having obtained the pseudo-objective function $\widehat{\Lambda}_{i,\theta}$, the implied best response and distributions are

$$\widehat{\alpha}_{i,\theta}(x, \varepsilon_i) = \arg \max_{a_i \in A_i} \{ \widehat{\Lambda}_{i,\theta}(a_i, x, \varepsilon_i) \} \quad \text{and}$$

$$\widehat{F}_{i,\theta}(a_i | x) = \int \mathbf{1}[\widehat{\alpha}_{i,\theta}(x, \varepsilon_i) \leq a_i] dQ_i(\varepsilon_i),$$

respectively. As shown in Section 2, the issue of the existence and uniqueness of solutions to $\widehat{\Lambda}_{i,\theta}(a_i, x, \varepsilon_i)$ depends crucially on the modeling of $u_{i,\theta}$. It is easy to see that we also have existence and uniqueness in the finite sample when conditions in Sections 2.3 and 2.4 hold for $u_i = u_{i,\theta}$ for all θ with the examples given above.

Note that $\widehat{F}_{i,\theta}(a_i|x)$ is a random distribution function of $\widehat{\alpha}_{i,\theta}(s_{in})$, conditioning on the event that $x_n = x$. In particular, $\widehat{F}_{i,\theta}$ is generally different from \widehat{F}_i , even when $\theta = \theta_0$ since the randomness of the former comes from the construction of the pseudo-model, while the latter is driven purely by the data. Although we know the distribution of ε_{in} , $\widehat{F}_{i,\theta}$ generally does not have a closed form and is generally infeasible; special cases do exist for unordered discrete action games; see AM and PSD. We denote a feasible estimator for $F_{i,\theta}$ by $\widetilde{F}_{i,\theta}$, which can be obtained by simulation. For instance, in our numerical studies, we use

$$\widetilde{F}_{i,\theta}(a_i|x) = \frac{1}{R} \sum_{r=1}^R \mathbf{1}[\widehat{\alpha}_{i,\theta}(x, \varepsilon_i^r) \leq a_i], \quad (13)$$

where $\{\varepsilon_i^r\}_{r=1}^R$ denotes a random sample drawn from the known distribution of ε_{in} .

Second stage optimization Given the estimators $(\widetilde{F}_{i,\theta}, \widehat{F}_i)$ for $(F_{i,\theta}, F_i)$, a class of L^2 -distance functions can be constructed from (potentially random) measures $\{\mu_{i,x}\}_{i \in \mathcal{I}, x \in X}$ defined on the support of A_i :

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

When A_i is finite, it is natural to choose each $\mu_{i,x}$ to be a count measure, where \widehat{M}_N can then be written as $\sum_{i \in \mathcal{I}} \sum_{x \in X} \sum_{a_i \in A_i} (\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_i(a_i|x))^2 \mu_{i,x}(\{a_i\})$. Our minimum distance estimator minimizes $\widehat{M}_N(\theta)$. The statistical properties of the estimator depend on the choice of $\{\mu_{i,x}\}_{i \in \mathcal{I}, x \in X}$: we discuss the selection of these measures in Section 4.

A REMARK ON SEMIPARAMETRIC ESTIMATION. Our methodology naturally generalizes to the case when x_n is a continuous random variable (or vector), where equation (6) becomes a linear integral equation of type II that has a well posed solution (see Srisuma and Linton (2012)). In this case, regardless of whether a_{in} is continuous or discrete, the estimation problem is semiparametric since \mathcal{L}_i becomes an operator on an infinite-dimensional space. Under Assumption M3, if a_{in} has a continuous component, then estimating \mathcal{H}_i also leads to a semiparametric problem. However, the dimensionality of an infinite-dimensional parameter is always 1, since each player forms an expectation based only on her action in the pseudo-decision problem. This is in contrast to the forward simulation method of BBL, where estimating value functions requires future states to be sequentially drawn from the estimator of G (not G_i) that is a conditional distribution conditioning on the actions of *all* players. In our case, the nonparametric dimensionality problem is determined by the total number of continuous variables present in a_{in} and x_n .

3.3 A discussion

Having gone through our two-step procedure in detail, we can now put its practical advantages in relation to its full solution counterpart into perspective. In particular, an analogous estimator can be defined by a two stage procedure similar to the one described above, where Step 1, in the first stage, now requires the equilibrium beliefs to be computed for each θ . Even if we have unlimited computational resources, multiple equilibria give rise to multiple beliefs, leading to more than one model implied distributions of actions. Without the indeterminacy issue, solving for the equilibrium numerically is also nontrivial: it typically involves fixed point iterations of some nonlinear functional equation, for example, see Pakes and McGuire (1994). The additional numerical cost required to solve for the equilibrium of dynamic games repeatedly is generally considered infeasible.

We use the insight from Hotz and Miller (1993) and its extension to dynamic games (AM and PSD), where we only consider the beliefs observed from the data that leads to the pseudo-model. As described in the previous section, there are no multiplicity issues associated with the pseudo-decision problem for the two main classes of games where players' actions are modeled to be monotone in the unobserved states or to be unordered discrete. Given the beliefs, the implied value functions and objective functions for the pseudo-decision problem are also easy to compute for each θ . Particularly, in Step 1 of the first stage, all the elements we require to estimate the continuation value function, $g_{i,\theta}$, either have explicit functional forms or are nonparametrically identified, hence they are easy to program (for instance, see Table 1).

We also comment on the prospect of solving equation (6), which we can think of as inverting the estimate of the matrix $I - \mathcal{L}_i$. Although not all estimators of \mathcal{L}_i lead to a nonsingular estimator of $I - \mathcal{L}_i$, a simple frequency estimator does. Importantly, since we estimate \mathcal{L}_i nonparametrically, suppose $I - \widehat{\mathcal{L}}_i$ is invertible; this inversion only has to be done once. In addition, similar to Hotz et al. (1994) and BBL, we can also take advantage of the linear structure of the (policy) value equation. Specifically, when the parameterization of θ in $u_{i,\theta}$ is linear, so that $u_{i,\theta} = \theta^\top u_{i,0}$ for some p -dimensional vector $u_{i,0}$, then $r_{i,\theta}$ can be written as $\theta^\top r_{i,0}$, where $r_{i,0}$ is a p -dimensional vector such that the $r_{i,0}(x) = E[u_{i,0}(\mathbf{a}_n, s_{in}) | x_n = x]$ for all x . In matrix notation, $r_{i,\theta} = \mathcal{R}_i \theta$, where \mathcal{R}_i is a $J \times p$ matrix whose j th row is $r_{i,0}^\top(x^j)$. Then $m_{i,\theta}$ equals $(I - \mathcal{L}_i)^{-1} \mathcal{R}_i \theta$ and for the choice-specific expected future return, $g_{i,\theta}$ in equation (11) simplifies to $\mathcal{H}_i (I - \mathcal{L}_i)^{-1} \mathcal{R}_i \theta$, where $\mathcal{H}_i (I - \mathcal{L}_i)^{-1} \mathcal{R}_i$ does not depend on θ .

In practice, the researcher has the freedom to choose any estimators for $r_{i,\theta}$, \mathcal{L}_i , and \mathcal{H}_i . Therefore, it is also straightforward to carry out our methodology in a fully parametric framework by parameterizing \mathcal{L}_i and \mathcal{H}_i . In particular, under the Markovian framework, \mathcal{L}_i and \mathcal{H}_i can be estimated independently of the dynamic parameters; they can then be used to transform the estimator of $r_{i,\theta}$ as discussed in Step 1 and then all of the above subsequent steps remain valid.

4. INFERENCE

Before we proceed to the asymptotic theorems, it is important to first consider whether the minimum distance approach suggested in the previous section provides a sensible

method to uncover θ_0 from the data. In particular, similar to other two-step estimators in the literature, the extent of what we can learn about θ_0 is restricted to the pseudo-best response functions $\{\alpha_{i,\theta}\}_{\theta \in \Theta}$ defined in (9). Therefore, it is appropriate to speak of identification in terms of the pseudo-model generated by the data.

DEFINITION 2. The set $\Theta_0 = \{\alpha_{i,\theta}(x, \varepsilon_{in}) = \alpha_i(x, \varepsilon_{in}) \text{ a.s. for all } (i, x)\}$ is called the *identified set*.

DEFINITION 3. The variable θ_0 is said to be *identified* if Θ_0 is a singleton set.

In Section 4.1, we show that, for the class of games discussed previously, $\{F_{i,\theta}\}_{\theta \in \Theta}$ contains the same identifying information on the identified set in the sense that the following two conditions are equivalent:

$$\alpha_{i,\theta}(x, \varepsilon_{in}) = \alpha_i(x, \varepsilon_{in}) \quad \text{a.s. for all } (i, x) \text{ iff } \theta \in \Theta_0, \quad (14)$$

$$F_{i,\theta}(a_{in}|x) = F_i(a_{in}|x) \quad \text{a.s. for all } (i, x) \text{ iff } \theta \in \Theta_0. \quad (15)$$

Section 4.2 then takes the identified set to be a singleton and provides conditions for our minimum distance estimator to be consistent and asymptotically normal.

4.1 Equivalence of identification conditions

We consider the parameterized versions of games discussed in Section 2.3. Specifically, let Assumptions S1', S3', and S4' be identical to Assumptions S1, S3, and S4 everywhere except that u_i is replaced by $u_{i,\theta}$ and all conditions imposed on the former are assumed to hold for the latter for all θ . In what follows, we denote the probability measure of ε_{in} by \mathcal{Q}_i . We begin with games that have finite actions.

PROPOSITION 3. Assume M1, M2, M3, S1', S2, and S3'. Then conditions (14) and (15) are equivalent.

PROOF. Suppose for each i , $A_i = \{a_i^1, \dots, a_i^{K_i}\}$. Then condition (15) only has to be checked on A_i .

Suppose (14) holds. The implication is immediate for $\theta \in \Theta_0$. Let $D_{i,x,\theta} = \{\alpha_{i,\theta}(x, \varepsilon_{in}) \neq \alpha_i(x, \varepsilon_{in})\}$. When $\theta \notin \Theta_0$, there exists some i, x , such that $\mathcal{Q}_i(D_{i,x,\theta}) > 0$. Let $D_{i,x,\theta}(k)$ denote $D_{i,x,\theta} \cap \{\alpha_i(x, \varepsilon_{in}) = a_i^k\}$, and let $k^* = \min\{k : \mathcal{Q}_i(D_{i,x,\theta}(k)) > 0\}$. By Assumption S2 and the monotonicity of $\alpha_{i,\theta}(x, \cdot)$ and $\alpha_i(x, \cdot)$, we have $F_{i,\theta}(a_i|x) = F_i(a_i|x)$ for all $a_i < a_i^{k^*}$ and $\mathcal{Q}_i(\{\alpha_{i,\theta}(x, \varepsilon_{in}) = a_i^{k^*}\}) \neq \mathcal{Q}_i(\{\alpha_i(x, \varepsilon_{in}) = a_i^{k^*}\})$. Therefore, $F_{i,\theta}(a_i^{k^*}|x) \neq F_i(a_i^{k^*}|x)$.

Suppose (15) holds. If $\theta \in \Theta_0$, then $\mathcal{Q}_i(\{\alpha_{i,\theta}(x, \varepsilon_{in}) = a_i^k\}) = \mathcal{Q}_i(\{\alpha_i(x, \varepsilon_{in}) = a_i^k\})$ for all k ; hence it follows from Assumption S2 and the monotonicity of $\alpha_{i,\theta}(x, \cdot)$ and $\alpha_i(x, \cdot)$ that $\mathcal{Q}_i(D_{i,x,\theta}) = 0$ for all i, x . If $\theta \notin \Theta_0$, let $k^* = \min\{k : F_{i,\theta}(a_i^k|x) - F_i(a_i^{k-1}|x) \neq F_i(a_i^k|x) - F_i(a_i^{k-1}|x)\}$, where we define $F_{i,\theta}(a_i^0|x) = F_i(a_i^0|x) = 0$. By Assumption S2 and the monotonicity of $\alpha_{i,\theta}(x, \cdot)$ and $\alpha_i(x, \cdot)$, it follows that $\{\alpha_{i,\theta}(x, \varepsilon_{in}) \leq a_i\}$ and $\{\alpha_i(x, \varepsilon_{in}) \leq a_i\}$ may differ only on a \mathcal{Q}_i null set for $a_i < a_i^{k^*}$. Therefore, $\mathcal{Q}_i(\{\alpha_{i,\theta}(x, \varepsilon_{in}) = a_i^{k^*}\} \Delta \{\alpha_i(x, \varepsilon_{in}) = a_i^{k^*}\}) > 0$.¹⁰ \square

¹⁰For any sets A, B , $A \Delta B = (A \cup B) \setminus (A \cap B)$ denotes the symmetric difference between A and B .

An equivalence result is also available when the distribution of a_{in} is continuous, that is, the best response is strictly monotone in ε_i .

PROPOSITION 4. *Assume M1, M2, M3, S1', S2, and S4', and for all i, x, θ , that $\alpha_{i,\theta}(x, \varepsilon_i)$ is strictly increasing in ε_i . Then conditions (14) and (15) are equivalent.*

PROOF. The inverse of $\alpha_{i,\theta}(x, \cdot)$ exists and is unique for all i, x, θ by strict monotonicity. We denote the inverse by $\rho_{i,\theta}(\cdot, x)$, so that $\rho_{i,\theta}(\alpha_{i,\theta}(x, \varepsilon_i), x) = \varepsilon_i$ for all $i, \theta, x, \varepsilon_i$. Then, for any a_i, x ,

$$\begin{aligned} F_{i,\theta}(a_i|x) &= \Pr[\alpha_{i,\theta}(x, \varepsilon_{in}) \leq a_i | x_n = x] \\ &= \Pr[\varepsilon_{in} \leq \rho_{i,\theta}(a_i, x) | x_n = x] \\ &= Q_i(\rho_{i,\theta}(a_i, x)). \end{aligned}$$

Since Q_i is a bijection map, as it is strictly increasing (Assumption S2), the one-to-one correspondence between $\alpha_{i,\theta}$ and $\rho_{i,\theta}$ for all θ completes the equivalence claim. \square

We have an analogous result when the distribution of a_{in} has finite mass points as well as a continuous contribution. For notational simplicity, we consider games where each action variable has a single mass point at the lower boundary of the support.

PROPOSITION 5. *Assume M1, M2, M3, S1', S2, and S4', and for all i, x, θ , that there exists $\underline{\varepsilon}_{i,x,\theta} \in \mathcal{E}_i$ such that $\alpha_{i,\theta}(x, \varepsilon_i) = \underline{a}_i$ for all $\varepsilon_i \leq \underline{\varepsilon}_{i,x,\theta}$ and $\alpha_{i,\theta}(x, \varepsilon_i)$ is strictly increasing in ε_i for $\varepsilon_i > \underline{\varepsilon}_{i,x,\theta}$. Furthermore, $\underline{\varepsilon}_{i,x,\theta} = \underline{\varepsilon}_{i,x} > \underline{\varepsilon}_i$ whenever $\theta \in \Theta_0$. Then conditions (14) and (15) are equivalent.*

PROOF. We only consider θ such that $\underline{\varepsilon}_{i,x,\theta} > \underline{\varepsilon}_i$. As seen previously, we repeatedly make use of Assumption S2 and the monotonicity of $\alpha_{i,\theta}(x, \cdot)$ and $\alpha_i(x, \cdot)$.

Suppose (14) holds. The implication is immediate for $\theta \in \Theta_0$. If $\theta \notin \Theta_0$, then for some i, x , either (i) $\underline{\varepsilon}_{i,x,\theta}^0 \neq \underline{\varepsilon}_{i,x}$ so that $\alpha_{i,\theta}(x, \varepsilon_i)$ and $\alpha_i(x, \varepsilon_i)$ do not agree when $\varepsilon_i \in (\min\{\underline{\varepsilon}_{i,x,\theta}^0, \underline{\varepsilon}_{i,x}\}, \max\{\underline{\varepsilon}_{i,x,\theta}^0, \underline{\varepsilon}_{i,x}\})$, in which case $F_{i,\theta}(\underline{a}_i|x) \neq F_i(\underline{a}_i|x)$, or (ii) $\underline{\varepsilon}_{i,x,\theta}^0 = \underline{\varepsilon}_{i,x}$ so that strict monotonicity implies $\alpha_{i,\theta}(x, \cdot)$ and $\alpha_i(x, \cdot)$ must have different inverses, hence different implied distribution functions.

Suppose (15) holds. The implication is now obvious for $\theta \in \Theta_0$. If $\theta \notin \Theta_0$, then either (i) $F_{i,\theta}(\underline{a}_i|x) \neq F_i(\underline{a}_i|x)$, in which case $Q_i(\{\alpha_{i,\theta}(x, \varepsilon_{in}) = \underline{a}_i\}) \neq Q_i(\{\alpha_i(x, \varepsilon_{in}) = \underline{a}_i\})$, or (ii) the one-to-one correspondence between the best responses and their implied distribution functions implies that $\{\alpha_{i,\theta}(x, \varepsilon_{in}) \neq \alpha_i(x, \varepsilon_{in})\}$ has a positive measure. \square

When θ_0 is identified, equivalence between conditions (14) and (15) means that a minimum distance criterion function can be constructed so that it has a unique mini-

mum only at θ_0 . For instance, it is sufficient that for all i, x , any $E \subset A_i$ that has positive probability measure with respect to the distribution of a_{in} also has a positive measure on $\mu_{i,x}$. The equivalence of information content on the identified set between the pseudo-best response function and the implied distribution is not restricted to games with monotone strategies. Conditions (14) and (15) are also equivalent for the discrete choice games studied in AM and PSD. Since (15) can be stated in terms of the choice probabilities (see (5)), the equivalence condition follows from the one-to-one relationship between the choice probabilities and the optimal decision rule using Hotz and Miller's (1993) well known inversion result (see also Lemma 1 of Pesendorfer and Schmidt-Dengler (2003)).

4.2 Asymptotic theorems

We state the regularity conditions for our theorems in terms of the distribution functions and their estimators. These conditions are more informative than the usual high level conditions as they allow us to highlight key features of the minimum distance estimator. They are also flexible enough to cover all the games considered in this paper, and to admit any estimators for $F_{i,\theta}$ and F_i as long as the conditions below are satisfied. Indeed, our Theorems 1 and 2 are also applicable to any estimation problem based on minimizing the distance of conditional distribution functions outside the context of dynamic games. The proofs of Theorems 1 and 2 can be found in Appendix B.

Specific to our application, for some estimators $(\tilde{F}_{i,\theta}, \hat{F}_i)$ of $(F_{i,\theta}, F_i)$, recall that the objective function is

$$\hat{M}_N(\theta) = \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (\tilde{F}_{i,\theta}(a_i|x) - \hat{F}_i(a_i|x))^2 \mu_{i,x}(da_i),$$

where $\tilde{F}_{i,\theta}$ is a feasible estimator for $F_{i,\theta}$. However, $\tilde{F}_{i,\theta}$ may generally not be smooth in θ due to simulation (see (13)). We denote a smooth version of \hat{M}_N by M_N , where $\tilde{F}_{i,\theta}$ is replaced by $\hat{F}_{i,\theta}$, an infeasible estimator of $F_{i,\theta}$, and denote its limiting function by M , so that

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

$$M(\theta) = \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} (F_{i,\theta}(a_i|x) - F_i(a_i|x))^2 \mu_{i,x}(da_i).$$

The minimum distance estimator is defined to be any sequence $\hat{\theta}$ that satisfies

$$\hat{M}_N(\hat{\theta}) \leq \inf_{\theta \in \Theta} \hat{M}_N(\theta) + o_p(N^{-1}).$$

ASSUMPTION A1.

- (i) The set Θ is a compact subset of \mathbb{R}^p .
- (ii) For all i, a_i, x , $F_{i,\theta}(a_i|x)$ and $F_i(a_i|x)$ exist, and $F_{i,\theta}(a_i|x) = F_i(a_i|x)$ if and only if $\theta = \theta_0$.

(iii) For all i, a_i, x , $F_{i,\theta}(a_i|x)$ is continuous on Θ .

(iv) For all i, x , $\mu_{i,x}$ is a nonrandom finite measure on A_i that dominates the distribution of a_{in} .

(v) For all i, x , $\sup_{(\theta,a_i) \in \Theta \times A_i} |\tilde{F}_{i,\theta}(a_i|x) - \hat{F}_{i,\theta}(a_i|x)| = o_p(1)$.

(vi) For all i, x , $\sup_{(\theta,a_i) \in \Theta \times A_i} |\hat{F}_{i,\theta}(a_i|x) - F_{i,\theta}(a_i|x)| = o_p(1)$.

(vii) For all i, x , $\sup_{a_i \in A_i} |\hat{F}_i(a_i|x) - F_i(a_i|x)| = o_p(1)$.

Assumption A1(ii) is the point-identification assumption of the pseudo-model. Assumption A1(iv) ensures that the measures used to define the objective function do not lose any identifying information on θ_0 . In application, A_i is generally compact, hence finiteness of the measures is a mild assumption. Note that the integral representation of \tilde{M}_N , M_N , and M encompasses games with discrete, continuous, or mixed discrete–continuous actions. When A_i is finite, $\mu_{i,x}$ is a count measure and it is sufficient to choose measures that put positive weights on each point of A_i . For a purely continuous action game, the domination condition is satisfied by choosing any measure dominated by the Lebesgue measure, for instance, the uniform measure. For an intermediate case with a_{in} that has a mixture of discrete and continuous distributions, $\mu_{i,x}$ is simply a combination of the count and continuous measures. We can also allow the measures to be random. Specifically, we can also use any random measure $\hat{\mu}_{i,x}$ as long as it converges (weakly) to $\mu_{i,x}$ that satisfy the finiteness and dominant conditions; one such candidate is the empirical measure, which puts equal mass on each observed data point (a_{in}, x_n) and puts zero measure outside of it.¹¹ Assumption A1(i)–(iv) imply that $M(\theta)$ has a well separated minimum over a compact set at θ_0 . The remaining conditions require our estimators for the distribution functions to be uniformly consistent, which can generally be verified using empirical process theory (see van der Vaart and Wellner (1996)). Note that A1(v) is not relevant if $\tilde{F}_{i,\theta}$ is feasible. An important special case is when $\tilde{F}_{i,\theta}$ is the naive Monte Carlo integration estimator. Suppose $\tilde{F}_{i,\theta}$ is defined as in (13). Then

$$\begin{aligned} \tilde{F}_{i,\theta}(a_i|x) - \hat{F}_{i,\theta}(a_i|x) &= \frac{1}{R} \sum_{r=1}^R \mathbf{1}[\hat{\alpha}_{i,\theta}(x, \varepsilon_i^r) \leq a_i] \\ &\quad - \int \mathbf{1}[\hat{\alpha}_{i,\theta}(x, \varepsilon_i) \leq a_i] dQ_i(\varepsilon_i), \end{aligned} \tag{16}$$

so that A1(v) is expected to hold as $R \rightarrow \infty$ by an application of Glivenko–Cantelli theorem. Assumption A1(vi) requires a standard equicontinuity condition and uniform consistent estimation of the parameters in the first stage. Assumption A1(vii) follows from the classical uniform law of large numbers.

THEOREM 1 (Consistency). *Under Assumption A1, $\hat{\theta} \xrightarrow{p} \theta_0$.*

¹¹The proofs in Appendix B can be lengthened, leading to the same asymptotic results for random measures $\{\hat{\mu}_{i,x}\}_{i \in \mathcal{I}, x \in X}$, where $\hat{\mu}_{i,x}$ converges weakly to $\mu_{i,x}$ for all (i, x) , using repeated applications of the continuous mapping theorem (see Ranga Rao (1962)).

To show asymptotic normality we require additional assumptions. In what follows, let \rightsquigarrow denote weak convergence and let $l^\infty(A_i)$ denote the space of all bounded functions on A_i .

ASSUMPTION A2.

- (i) *The true parameter θ_0 lies in the interior of Θ .*
- (ii) *For all i , a_i , x , $F_{i,\theta}(a_i|x)$ and $\widehat{F}_{i,\theta}(a_i|x)$ are twice continuously differentiable in θ in a neighborhood of θ_0 , and $\int_{A_i} \frac{\partial}{\partial \theta_l} F_{i,\theta}(a_i|x) \mu_{i,x}(da_i)$ and $\int_{A_i} \frac{\partial^2}{\partial \theta_l \partial \theta_{l'}} F_{i,\theta}(a_i|x) \mu_{i,x}(da_i)$ exist for all l, l' for θ in a neighborhood of θ_0 .*
- (iii) *The matrix $\sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial}{\partial \theta} F_{i,\theta}(a_i|x) \frac{\partial}{\partial \theta^\top} F_{i,\theta}(a_i|x) \mu_{i,x}(da_i)$ is positive definite at $\theta = \theta_0$.*
- (iv) *For all i, l, x , $\sup_{a_i \in A_i} |\frac{\partial}{\partial \theta_l} \widehat{F}_{i,\theta}(a_i|x) - \frac{\partial}{\partial \theta_l} F_{i,\theta_0}(a_i|x)| = o_p(1)$ as $\|\theta - \theta_0\| \rightarrow 0$.*
- (v) *For all i, l, l', x , $\sup_{a_i \in A_i} |\frac{\partial^2}{\partial \theta_l \partial \theta_{l'}} \widehat{F}_{i,\theta}(a_i|x) - \frac{\partial^2}{\partial \theta_l \partial \theta_{l'}} F_{i,\theta_0}(a_i|x)| = o_p(1)$ as $\|\theta - \theta_0\| \rightarrow 0$.*
- (vi) *For all i, x , $\sup_{a_i \in A_i} |\widetilde{F}_{i,\theta}(a_i|x) - \widehat{F}_{i,\theta}(a_i|x)| = o_p(1/\sqrt{N})$ uniformly in a neighborhood of θ_0 .*
- (vii) *For all i, x , $\sqrt{N}(\widehat{F}_i(\cdot|x) - F_i(\cdot|x)) \rightsquigarrow \mathbb{V}_{i,x}$, where $\mathbb{V}_{i,x}$ is a tight Gaussian process that belongs to $l^\infty(A_i)$.*
- (viii) *For all i, x , $\sqrt{N}(\widehat{F}_{i,\theta_0}(\cdot|x) - F_{i,\theta_0}(\cdot|x)) \rightsquigarrow \mathbb{W}_{i,x}$, where $\mathbb{W}_{i,x}$ is a tight Gaussian process that belongs to $l^\infty(A_i)$.*
- (ix) *For all i, x , $\sqrt{N}(\widehat{F}_{i,\theta_0}(\cdot|x) - \widehat{F}_i(\cdot|x)) \rightsquigarrow \mathbb{T}_{i,x}$, where $\mathbb{T}_{i,x}$ is a tight Gaussian process that belongs to $l^\infty(A_i)$.*

Condition A2(i)–(v) are standard regularity and smoothness assumptions. Since $F_{i,\theta}(a_i|x)$ is twice continuously differentiable in θ (near θ_0), sufficient conditions for A2(iv) and (v) are uniform consistency of the first and second derivatives of $\widehat{F}_{i,\theta}$ to $F_{i,\theta}$ respectively (cf. A1(vi)). Assumption A2(vi) imposes a rate for the simulation error. If $\widetilde{F}_{i,\theta}$ is defined by (13), then $\sqrt{N}(\widetilde{F}_{i,\theta} - \widehat{F}_{i,\theta})$ is an empirical process (see equation (16)) that is expected to satisfy the Donsker theorem. The remaining conditions assume that uniform central limit theorems hold on A_i . When A_i is finite, the uniform limit theorem reduces to the multivariate central limit theorem where the tightness condition is trivially satisfied; otherwise, these can be verified using empirical process theory (cf. A1(v)–(vii)). Specifically, A2(viii) captures the effects from using a first step estimator, which typically can be verified by showing the linearization of $\sqrt{N}(\widehat{F}_{i,\theta_0} - F_{i,\theta_0})$ satisfies the Donsker's theorem. When the limiting distributions in A2(vii) and (viii) are jointly Gaussian, which is expected to hold in most applications, A2(ix) immediately follows from the continuous mapping theorem.

THEOREM 2 (Asymptotic Normality). *Under Assumptions A1 and A2,*

$$\sqrt{N}(\widehat{\theta} - \theta_0) = \left(\frac{\partial^2}{\partial \theta \partial \theta^\top} M(\theta_0) \right)^{-1} \sqrt{N} \frac{\partial}{\partial \theta} M_N(\theta_0) + o_p(1)$$

and $\sqrt{N}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, \mathcal{W}^{-1}\mathcal{V}\mathcal{W}^{-1})$, where

$$\mathcal{V} = \lim_{N \rightarrow \infty} \text{var} \left(2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \left(\frac{\partial}{\partial \theta} F_{i, \theta_0}(a_i | x) \right. \right. \\ \left. \left. \times (\hat{F}_{i, \theta_0}(a_i | x) - \hat{F}_i(a_i | x)) \right) \mu_{i, x}(da_i) \right), \quad (17)$$

$$\mathcal{W} = 2 \sum_{i \in \mathcal{I}} \sum_{x \in X} \int_{A_i} \frac{\partial}{\partial \theta} F_{i, \theta_0}(a_i | x) \frac{\partial}{\partial \theta^\top} F_{i, \theta_0}(a_i | x) \mu_{i, x}(da_i). \quad (18)$$

The asymptotic distribution of our estimator shows no effect of using the feasible estimator $\tilde{F}_{i, \theta}$ instead of $\hat{F}_{i, \theta}$. In order to perform inference a feasible estimator for the asymptotic variance is required. Bootstrapping is a natural candidate to estimate the standard error in this setting.¹² In a closely related framework, [Kasahara and Shimotsu \(2008a\)](#) developed a bootstrap procedure for a parametric discrete decision model that can be applied to discrete action games (under Assumption D). Recently, [Cheng and Huang \(2010\)](#) provided some general conditions to validate the use of the bootstrap as an inferential tool for a general class of semiparametric M -estimators when the objective function is not smooth. We show in the next section that bootstrapping performs well with our minimum distance estimator.

A REMARK ON SEMIPARAMETRIC ESTIMATION. Theorems 1 and 2 are applicable to both parametric and semiparametric problems. In the context of dynamic games, the first stage estimators (finite and/or infinite dimensional) are defined implicitly in our objective function \hat{M}_N through $\tilde{F}_{i, \theta}$. The uniform consistency and functional central limit theory requirements in A1 and A2 are standard for a minimum distance estimator. These uniformity conditions can be verified using modern empirical process theory under weak conditions. In particular, for the simulation estimator defined in (13), [Andrews \(1994, “type IV class”\)](#) and [Chen, Linton, and van Keilegom \(2003, Theorem 3.2\)](#) provided conditions for the Donsker theorem to hold in a parametric and semiparametric setting, respectively.¹³

Possible extensions In this paper we have focused on a consistent estimation method for a large class of dynamic games. However, there are two important aspects of our estimators we have not discussed. These are the issues of efficiency and finite sample bias.

¹²Recently, [Ackerberg, Chen, and Hahn \(2012\)](#) proposed a way to simplify semiparametric inference when unknown functions are estimated by the method of sieves. They considered, as a specific example, a class of discrete action games, where they focused on estimating finite conditional moment models and also required the objective function to be smooth. Therefore, despite the fact that our theorems admit sieves estimators, their results are generally not applicable to our estimator or to other notable estimators in this literature (e.g., the iterative estimator of [Aguirregabiria and Mira \(2007\)](#) and the inequality estimator of BBL).

¹³[Srisuma \(2010\)](#) gave a set of primitive conditions where Assumptions A1 and A2 are satisfied for a single agent problem that coincides with the purely continuous action game in Section 2.3 when $I = 1$.

Our minimum distance estimator is not efficient. For example, when A_i is finite, we can create large vectors of the conditional distribution of actions across all players, action choices, and observable states, and then our objective function is a special case of the asymptotic least squares estimators analogous to the setup in PSD with a diagonal weighting matrix. In principle, we can provide a more efficient estimator by considering a more general metric to match the distribution functions and constructing efficient weights (that rely on a consistent preliminary estimator). However, the efficient weights generally require estimates of $\frac{\partial}{\partial \theta} F_{i,\theta}(a_i|x)$ for all a_i, x , which rely on further numerical approximations when the feasible estimator of $F_{i,\theta}$ is not smooth (for recent results on statistical properties of estimators with numerical derivatives, see [Hong, Mahajan, and Nekipelov \(2010\)](#)). The issue of efficient estimation for this class of games is a challenging and interesting problem in both theory and practice, especially in a semiparametric model.

Another important concern for two-step estimators is the bias in small sample. In a single agent discrete choice setting, [Aguirregabiria and Mira \(2002\)](#) proposed iteration methods that appear to improve the finite sample performance of their estimators. [Kasahara and Shimotsu \(2008a\)](#) gave a theoretical explanation of Aguirregabiria and Mira's findings; their idea is that a fixed point constraint of the pseudo-model implied choice probabilities provides an iteration operator that can be used to reduce the bias in the first stage estimation. Although such an iteration procedure may not converge, especially in a game setting ([Pesendorfer and Schmidt-Dengler \(2010\)](#)), [Kasahara and Shimotsu \(2012\)](#) recently provided an alternative iteration method that leads to a consistent estimator even when the fixed point constraint is not a contraction (hence it need not ensure global convergence). The frameworks that the aforementioned papers consider are games under Assumption D in Section 2.4. Since equation (10) also represents a fixed point constraint, it will be interesting to study whether analogous iterative schemes can be developed for other class of games such as those considered in this paper.

5. NUMERICAL EXAMPLES

We apply our methodology described in Section 4 to estimate two simulated dynamic models with continuous actions. We construct our minimum distance estimators based on the estimators proposed in Table 1 and Example 1; first, in a semiparametric dynamic price setting problem for a single agent firm and, second, in a parametric framework, where we use our estimator and BBL's to estimate a repeated Cournot duopoly game. Since it is generally difficult to solve a dynamic optimization problem, the models below are kept simple so as to generate the data. It is easy to check that both examples below satisfy conditions M1, M2, M3, S1', S2, and S4', so that monotone pure strategy equilibria exist and players only employ monotone best response strategies.

DESIGN 1 (Markov Decision Problem). At every period, each firm faces the demand function

$$D_\theta(a, x, \varepsilon) = \bar{D} - \theta_1 a + \theta_2(x + \varepsilon),$$

where a denotes the price, x is the demand shifter (e.g., some observable measure of the consumer's satisfaction), and ε is the firm's private demand shock. The term \bar{D} can be interpreted as a constant market size and (θ_1, θ_2) denote the parameters that represent the market elasticities that lie in $\mathbb{R}^+ \times \mathbb{R}^+$. The firm's profit function is

$$u_\theta(a, x, \varepsilon) = D_\theta(a, x, \varepsilon)(a - c),$$

where c denotes a constant marginal cost. The price-setting decision today affects the demand for the next period. Specifically, x_n takes a value of either 1 or -1 , and its transitional distribution is summarized by $\Pr[x'_n = -1 | x_n, a_n = a] = \frac{a - \underline{a}}{\bar{a} - \underline{a}}$, where \underline{a} and \bar{a} denote the minimum and maximum possible prices, respectively. The evolution of private shocks is completely random and transitory, and ε_n is distributed uniformly on $[-1, 1]$. The firm chooses price a_n to maximize its discounted expected profit, where future payoff is discounted by $\beta = 0.9$. The values of (\bar{D}, c) are assigned to be $(3, 1)$ and the data are generated using the optimal decision when $\theta = (1, 0.5)$. We generate 500 replications of the controlled Markov processes with sample size $N \in \{20, 100, 200\}$, where each decision series spans five time periods. This leads to three sets of experiments with the total sample size, NT , of 100, 500, and 1000.

We have two estimators, denoted by $\hat{\theta}^{\text{UM}}$ and $\hat{\theta}^{\text{EM}}$, that minimize the objective functions constructed using the uniform and empirical measures, respectively. For the non-parametric estimator of the transition law, $G(x'|x, a)$, we use a truncated fourth order kernel based on the density of a standard normal random variable (see Rao (1983)). For each replication, we experiment with three different bandwidths $\{h_\varsigma = 1.06s(NT)^{-\varsigma} : \varsigma = \frac{1}{6}, \frac{1}{7}, \frac{1}{8}\}$; the order of the bandwidth is chosen to be consistent with a derivative of the one-dimensional kernel estimator for a density or regression derivative (for example, see Hansen (2008)).¹⁴ We simulate the pseudo-distribution function using $N \log(N)$ random draws. The number of bootstrap draws is 99.

We report the bias, median of the bias, standard deviation, coverage probability of 95% confidence interval based on a standard normal approximation, and the bootstrapped standard errors and coverage probabilities from the bootstrapped distributions. Tables 2 and 3 give the results for θ_1 and θ_2 , respectively, where the bootstrapped values are given in italics.

We make the following general observations for our estimators across all bandwidths and measures: (i) the median of the bias is similar to the mean; (ii) the estimators are consistent, as N increases the bias and the standard deviation converges to zero; (iii) the performance of the bootstrapped standard errors steadily approaches the true value with increasing sample size and appears to be consistent; (iv) the coverage probabilities improve with sample size, although the results for θ_1 are closer to the nominal value than are those for θ_2 , the bootstrapped confidence intervals appear to perform reasonably well, and even favorably in some cases, relative to the normal approximations with

¹⁴When $A_{i,\theta}$ is smooth, by the implicit function theorem, $\alpha_{i,\theta}$ is a smooth functional of $E[\frac{\partial}{\partial a} u_{i,\theta}(\cdot, \mathbf{a}_{-in}, x_n, \cdot) | x_n = \cdot]$ and $\frac{\partial}{\partial a} g_{i,\theta}$ from $\frac{\partial}{\partial a_i} A_{i,\theta}(\alpha_{i,\theta}(s_i), s_i) = 0$. Since $E_n[\frac{\partial}{\partial a} u_{i,\theta}(\cdot, \mathbf{a}_{-in}, x_n, \cdot) | x_n = \cdot]$ converges at the parametric rate, the rate of convergence of $\hat{\alpha}_{i,\theta}$ is determined by $\frac{\partial}{\partial a} \hat{g}_{i,\theta}$.

TABLE 2. Monte Carlo results (Markov decision process).

NT	s	$\hat{\theta}_1^{UM}$				$\hat{\theta}_1^{EM}$			
		Bias	Mbias	Std	95%	Bias	Mbias	Std	95%
100	1/6	-0.0101	0.0118	0.1716	0.9560	0.0023	0.0205	0.1722	0.9640
		—	—	0.2594	0.9960	—	—	0.2370	0.9960
	1/7	0.0014	0.0248	0.1543	0.9600	0.0128	0.0305	0.1515	0.9640
500	1/8	—	—	0.2370	0.9900	—	—	0.2326	0.9840
	1/8	0.0067	0.0296	0.1569	0.9680	0.0233	0.0409	0.1406	0.9580
		—	—	0.2227	0.9740	—	—	0.2129	0.9740
1000	1/6	0.0009	0.0024	0.0695	0.9360	0.0025	0.0040	0.0693	0.9360
		—	—	0.0784	0.9760	—	—	0.0784	0.9800
	1/7	0.0044	0.0086	0.0620	0.9480	0.0065	0.0091	0.0617	0.9440
1000	1/8	—	—	0.0706	0.9720	—	—	0.0708	0.9720
	1/8	0.0105	0.0155	0.0574	0.9380	0.0121	0.0164	0.0583	0.9340
		—	—	0.0797	0.9660	—	—	0.0649	0.9700
1000	1/6	-0.0023	0.0003	0.0511	0.9380	-0.0025	-0.0007	0.0507	0.9400
		—	—	0.0552	0.9540	—	—	0.0552	0.9540
	1/7	0.0021	0.0045	0.0475	0.9500	0.0028	0.0044	0.0474	0.9540
1000	1/8	—	—	0.0500	0.9600	—	—	0.0500	0.9640
	1/8	0.0073	0.0086	0.0457	0.9460	0.0075	0.0081	0.0450	0.9460
		—	—	0.0463	0.9500	—	—	0.0462	0.9460

Note: The bandwidth used in the nonparametric estimation is $h_s = 1.06s(NT)^{-s}$, where s is the standard deviation of $\{a_{nt}\}_{n=1, t=1}^{N, T}$.

infeasible variance at larger sample sizes. Therefore, the bootstrap appears to offer one reasonable mode to perform inference for our estimator.

DESIGN 2 (Cournot Game). We use a variant of a repeated Cournot duopoly competition studied in PSD. We specify a linear inverse demand function

$$D_{\theta}(\mathbf{a}, x) = x(\bar{D} - \theta_1(a_1 + a_2)),$$

where a_i denotes the quantity supplied by player i , x is the demand shifter that rotates the slope of the demand curve, and \bar{D} represents the market size similar to Example 1. The parameter space for (θ_1, θ_2) is $\mathbb{R}^+ \times \mathbb{R}^+$. Each firm has a private stochastic marginal cost, driven by ε_i , so that the profit function for each period is

$$u_{i,\theta}(a_i, a_j, x, \varepsilon_i) = a_i(D(\mathbf{a}, x) - \theta_2\varepsilon_i) \quad \text{for } i, j = 1, 2 \text{ and } i \neq j.$$

The distribution of ε_{in} is normal with mean 0 and variance 1, and is distributed independently across players, time, and other variables. The observable state is the stochastic demand coefficient x_n that has 0.5 probability of taking values 2 or 4, independently of previous actions and states. Thus an equilibrium exists; in particular, the symmetric strategy profile where each player maximizes her expected static profit (a noncooperative Nash equilibrium) in every period is an equilibrium. We add a dynamic dimension

TABLE 3. Monte Carlo results (Markov decision process).

NT	s	$\widehat{\theta}_2^{\text{JM}}$				$\widehat{\theta}_2^{\text{EM}}$			
		Bias	Mbias	Std	95%	Bias	Mbias	Std	95%
100	1/6	0.1140	0.0546	0.2633	0.9440	0.1079	0.0493	0.2460	0.9380
		—	—	<i>0.3245</i>	<i>0.9940</i>	—	—	<i>0.3118</i>	<i>0.9940</i>
	1/7	0.0986	0.0580	0.2213	0.9320	0.0985	0.0575	0.2257	0.9420
		—	—	<i>0.3126</i>	<i>0.9920</i>	—	—	<i>0.3040</i>	<i>0.9940</i>
	1/8	0.1030	0.0583	0.2267	0.9380	0.0987	0.0524	0.2110	0.9360
		—	—	<i>0.2969</i>	<i>0.9920</i>	—	—	<i>0.2885</i>	<i>0.9940</i>
500	1/6	0.0381	0.0370	0.0878	0.9220	0.0386	0.0371	0.0889	0.9160
		—	—	<i>0.1091</i>	<i>0.9740</i>	—	—	<i>0.1086</i>	<i>0.9740</i>
	1/7	0.0383	0.0307	0.0860	0.9200	0.0373	0.0317	0.0867	0.9180
		—	—	<i>0.1078</i>	<i>0.9580</i>	—	—	<i>0.1026</i>	<i>0.9700</i>
	1/8	0.0374	0.0308	0.0839	0.9060	0.0367	0.0312	0.0854	0.9140
		—	—	<i>0.3005</i>	<i>0.9460</i>	—	—	<i>0.0964</i>	<i>0.9460</i>
1000	1/6	0.0338	0.0319	0.0699	0.9260	0.0332	0.0310	0.0691	0.9240
		—	—	<i>0.0753</i>	<i>0.9560</i>	—	—	<i>0.0753</i>	<i>0.9560</i>
	1/7	0.0317	0.0275	0.0668	0.9320	0.0308	0.0255	0.0670	0.9260
		—	—	<i>0.0704</i>	<i>0.9420</i>	—	—	<i>0.0706</i>	<i>0.9440</i>
	1/8	0.0316	0.0256	0.0648	0.9240	0.0310	0.0232	0.0650	0.9220
		—	—	<i>0.0662</i>	<i>0.9300</i>	—	—	<i>0.0665</i>	<i>0.9300</i>

Note: The bandwidth used in the nonparametric estimation is $h_s = 1.06s(NT)^{-s}$, where s is the standard deviation of $\{a_{nt}\}_{n=1, t=1}^{N, T}$.

to our estimation problem by misspecifying the model (see below). Our data are generated from the symmetric equilibrium in the static duopoly game, where \bar{D} is normalized to 1, we use $\theta_0 = (0.2, 0.2)$, and the discounting factor is 0.9. For each simulation, we generate $N \in \{100, 500, 1000\}$ independent draws from the equilibrium. The experiment is repeated 500 times for each N .

For our estimators, as previously, we use two estimators constructed from the objective functions with uniform measures and empirical measures. We allow for a particular misspecification such that our agent maximizes the objective function (cf. (8)) in the pseudo-optimization stage,

$$\tilde{\Lambda}_{i, \theta}(a_i, s_i) = E[u_{i, \theta}(a_i, \mathbf{a}_{-in}, x_n, \varepsilon_i) | x_n = x] + \beta_i \tilde{g}_i(a_i, x),$$

where \tilde{g}_i is a linear function of a_i that has a random slope, varying randomly with each player and state. The slope of \tilde{g}_i converges to zero at the parametric rate, and it is determined by a random draw from a normal distribution with mean zero and variance $\frac{1}{N}$. We simulate the pseudo-distribution function using $N \log(N)$ random draws.

We also consider two versions of BBL estimators: one is based on choosing an alternative strategy by an additive perturbation and the other by multiplicative perturbation. For additive perturbations, each inequality is represented by an alternative strategy $\tilde{\alpha}_1(\cdot; \eta_1)$ for some $\eta_1 \in \mathbb{R}$ such that $\tilde{\alpha}_1(s_i; \eta_1) = \alpha_{\theta_0}(s_i) + \eta_1$ for all $s_i \in S_i$, where α_{θ_0} is the (symmetric) optimal strategy estimable from the data. We draw η_1 from a normal distribution with mean 0 and variance 0.5. For multiplicative perturbation, each

inequality is represented by an alternative strategy $\tilde{\alpha}_2(\cdot; \eta_2)$ for some $\eta_2 \in \mathbb{R}$ such that $\tilde{\alpha}(s_i; \eta_2) = \eta_2 \alpha_{\theta_0}(s_i)$ for all $s_i \in S_i$. We draw η_2 from a normal distribution with mean 1 and variance 0.5. The BBL type objective functions are constructed based on using $N_I \in \{300, 600\}$ randomly drawn inequalities and the number of simulations used to compute the expected returns is 2000.¹⁵ BBL estimators correctly ignore the dynamics and estimate the repeated static game.

We show in the second example of Appendix A.1 that the parameters in the Cournot game are identified. However, with BBL's approach, we also show that the class of additive perturbations preserves the identifying information of θ_{01} but not θ_{02} , in the sense that the expected returns from employing the optimal strategies that generate the data (with $\theta = \theta_0$) are always at least as large as the returns from additively perturbed strategies for all $\theta' = (\theta_{01}, \theta'_2)$ with any value of θ'_2 . On the other hand, the inequalities based on multiplicative perturbations can preserve the identifying information of both θ_{01} and θ_{02} .

We report the bias, median of the bias, standard deviation, interquartile range scaled by 1.349 (which approximately equals the standard deviation for a normal variable), coverage probability of 95% confidence interval based on a standard normal approximation, and mean square error. Tables 4 and 5 give the results for our estimators, with and without the misspecified the dynamics, and BBL's estimators, constructed using additive and multiplicative perturbations, of θ_1 and θ_2 , respectively.

For θ_1 , as expected, all estimators appear to be consistent, and from looking at the coverage probabilities and comparing the standard deviation with the scaled interquartile range, are well approximated by a normal distribution. For θ_2 , as before, our estimators appear to be consistent and asymptotically normal. BBL's estimators of θ_2 show several interesting characteristics. The first general observation is that estimators obtained by using multiplicative perturbations perform better, as expected, at least for larger sample sizes; they also appear to be consistent, but seem to be less well approximated by a normal distribution compared to our estimators. For the estimators based on additive perturbations, the bias appears to increase with sample size, which can be explained by looking at the mathematical details of our examples in Appendix A, since the loss of identification only materializes in the limit. However, its standard deviation is decreasing with sample size, although it does so at an increasingly slower rate compared to the multiplicative perturbations. It is also unclear from our small scale studies what role the number of inequalities has on the statistical properties of BBL's estimators; for instance, we see that more inequalities lead to an improvement in the mean squared error for additive perturbations but not for multiplicative perturbations.

6. CONCLUSION

The discrete Markov decision process studied in Rust (1987) provided a useful framework to model and estimate dynamic games of incomplete information. In this paper,

¹⁵The number of inequalities and simulations we use represents the upper bound values that BBL used in their simulation studies, which conform to their asymptotic theorems. Specifically, see Assumption S2(iii) on p. 1348, N_I is allowed to grow to infinity at any rate, while the number of simulations is required to go to infinity at a faster rate than \sqrt{N} .

TABLE 4. Monte Carlo results (Cournot game).

N	$\hat{\theta}_1$	Bias	Mbias	Std	Iqr	95%	Mse
100	UM	0.0000	-0.0001	0.0014	0.0013	0.9460	0.0000
	UM-M	0.0001	-0.0001	0.0026	0.0027	0.9460	0.0000
	EM	-0.0004	-0.0004	0.0014	0.0013	0.9400	0.0000
	EM-M	-0.0003	-0.0005	0.0026	0.0028	0.9580	0.0000
	AP-L	-0.0000	-0.0000	0.0016	0.0017	0.9560	0.0000
	AP-H	-0.0001	-0.0001	0.0017	0.0015	0.9540	0.0000
	MP-L	0.0002	0.0000	0.0020	0.0019	0.9440	0.0000
	MP-H	0.0001	-0.0000	0.0021	0.0020	0.9400	0.0000
500	UM	0.0000	0.0000	0.0007	0.0007	0.9580	0.0000
	UM-M	0.0000	-0.0000	0.0012	0.0012	0.9580	0.0000
	EM	-0.0001	-0.0000	0.0007	0.0007	0.9540	0.0000
	EM-M	-0.0000	-0.0001	0.0012	0.0012	0.9560	0.0000
	AP-L	0.0000	0.0000	0.0008	0.0008	0.9380	0.0000
	AP-H	-0.0000	-0.0000	0.0007	0.0008	0.9580	0.0000
	MP-L	0.0000	0.0000	0.0010	0.0009	0.9580	0.0000
	MP-H	-0.0000	0.0000	0.0010	0.0009	0.9460	0.0000
1000	UM	-0.0000	0.0000	0.0004	0.0004	0.9460	0.0000
	UM-M	-0.0000	0.0000	0.0008	0.0008	0.9420	0.0000
	EM	-0.0001	-0.0000	0.0004	0.0004	0.9480	0.0000
	EM-M	-0.0001	-0.0000	0.0008	0.0008	0.9480	0.0000
	AP-L	-0.0000	0.0000	0.0006	0.0005	0.9380	0.0000
	AP-H	0.0001	0.0000	0.0005	0.0005	0.9440	0.0000
	MP-L	0.0000	0.0000	0.0007	0.0005	0.9720	0.0000
	MP-H	0.0000	0.0000	0.0007	0.0007	0.9460	0.0000

Note: UM and EM are our minimum distance estimators for the static games obtained from using uniform and empirical measures, respectively; UM-M and EM-M are their misspecified counterparts. AP-L and AP-H are BBL's estimators obtained from using additive perturbations with 300 and 600 inequalities, respectively. MP-L and MP-H are BBL's estimators obtained from using multiplicative perturbations with 300 and 600 inequalities, respectively.

we propose a two-step methodology, in a similar spirit to [Hotz and Miller \(1993\)](#), using the pseudo-model to estimate popular Markovian games studied in the literature. The pseudo-model is particularly useful in the estimation of games since it can avoid the practical and statistical complications when the actual model has multiple equilibria, as well as generally reducing the computational burden relative to the full solution approach. We give precise conditions that extend the scope of the pseudo-model—traditionally used to model games where players' actions are discrete and unordered (e.g., AM and PSD)—to games where players' actions are monotone in their private values that can be discrete, continuous, or mixed. We also show that pure strategy Markov equilibria exist for these estimable monotone choice games. Our estimator is defined to minimize the distance between the distribution of actions implied by the data and the pseudo-model that is motivated by a characterization of the equilibrium. Since the distribution functions are defined on the familiar Euclidean space, given an identified (pseudo-) model, we suggest simple metrics for constructing objective functions that can be used for consistent estimation. In contrast, BBL's method requires selection of

TABLE 5. Monte Carlo results (Cournot game).

N	$\hat{\theta}_2$	Bias	Mbias	Std	Iqr	95%	Mse
100	UM	-0.0009	-0.0011	0.0119	0.0131	0.9580	0.0001
	UM-M	0.0054	0.0053	0.0142	0.0132	0.9500	0.0002
	EM	0.0033	0.0029	0.0140	0.0139	0.9360	0.0002
	EM-M	0.0205	0.0164	0.0255	0.0206	0.8960	0.0011
	AP-L	-0.0174	-0.0421	0.2613	0.2153	0.9300	0.0686
	AP-H	0.0008	-0.0062	0.1520	0.1390	0.9480	0.0231
	MP-L	0.0268	0.0047	0.2623	0.2009	0.9400	0.0695
	MP-H	0.0217	0.0064	0.2753	0.2623	0.9500	0.0762
500	UM	-0.0002	-0.0003	0.0052	0.0050	0.9580	0.0000
	UM-M	0.0005	0.0006	0.0055	0.0053	0.9500	0.0000
	EM	0.0006	0.0002	0.0059	0.0055	0.9520	0.0000
	EM-M	0.0036	0.0036	0.0070	0.0069	0.9320	0.0001
	AP-L	-0.0241	-0.0828	0.2012	0.1380	0.9380	0.0411
	AP-H	-0.0150	-0.0248	0.1388	0.1097	0.9340	0.0195
	MP-L	-0.0010	0.0039	0.0945	0.0117	0.9260	0.0089
	MP-H	0.0046	0.0043	0.1191	0.0841	0.9400	0.0142
1000	UM	0.0001	-0.0000	0.0037	0.0038	0.9460	0.0000
	UM-M	0.0004	0.0006	0.0039	0.0039	0.9600	0.0000
	EM	0.0006	0.0004	0.0042	0.0046	0.9560	0.0000
	EM-M	0.0019	0.0022	0.0047	0.0045	0.9380	0.0000
	AP-L	-0.0288	-0.0943	0.1833	0.1024	0.9400	0.0344
	AP-H	-0.0168	-0.0295	0.1284	0.1141	0.9360	0.0168
	MP-L	0.0021	0.0000	0.0643	0.0046	0.9280	0.0041
	MP-H	-0.0054	0.0005	0.0820	0.0455	0.9080	0.0068

Note: UM and EM are our minimum distance estimators for the static games obtained from using uniform and empirical measures, respectively; UM-M and EM-M are their misspecified counterparts. AP-L and AP-H are BBL's estimators obtained from using additive perturbations with 300 and 600 inequalities, respectively. MP-L and MP-H are BBL's estimators obtained from using multiplicative perturbations with 300 and 600 inequalities, respectively.

alternative strategies, where a suitable choice of objective functions may be less obvious, especially when actions are continuously distributed. We illustrate the importance of choosing objective functions for consistent estimation in finite samples with a Monte Carlo study and provide the theoretical explanations in Appendix A.

There are several directions for future research. We focus on consistent estimation and have not provided an efficient estimator in this paper. Our methodology also appears to be amenable to adoption of an iterative scheme along the lines of Aguirregabiria and Mira (2002, 2007) and Kasahara and Shimotsu (2012) that may reduce the small sample bias of the first step estimator. Last, although we do not contribute to the development of ways to deal with unobserved heterogeneity and the related issues regarding multiple equilibria, we believe the recent progress made in the studies of dynamic discrete choice models, for example, the nonparametric finite mixture results of Kasahara and Shimotsu (2008b) or methods that take advantage of finite dependence structure in Arcidiacono and Miller (2008), can be adapted and extended to estimate the dynamic games considered in this paper.

REFERENCES

- Ackerberg, D., L. Benkard, S. Berry, and A. Pakes (2005), “Econometric tools for analyzing market outcome.” In *Handbook of Econometrics*, Vol. 6 (J. Heckman and E. Leamer, eds.), North-Holland, Amsterdam. [549, 560]
- Ackerberg, D., X. Chen, and J. Hahn (2012), “A Practical asymptotic variance estimator for two-step semiparametric estimators.” *Review of Economics and Statistics*, 94, 481–498. [573]
- Aguirregabiria, V. and P. Mira (2002), “Swapping nested fixed point algorithm: A class of estimators for discrete Markov decision models.” *Econometrica*, 70, 1519–1543. [574, 580]
- Aguirregabiria, V. and P. Mira (2007), “Sequential estimation of dynamic discrete games.” *Econometrica*, 75, 1–53. [550, 573, 580]
- Aguirregabiria, V. and P. Mira (2010), “Dynamic discrete choice structural models: A survey.” *Journal of Econometrics*, 156, 38–67. [549, 560]
- Andrews, D. W. K. (1994), “Empirical processes methods in econometrics.” In *Handbook of Econometrics*, Vol. 4 (R. F. Engle and D. McFadden, eds.), North-Holland, Amsterdam. [573]
- Arcidiacono, P. and R. A. Miller (2008), “CCP estimation of dynamic discrete choice models with unobserved heterogeneity.” Working paper. [561, 580]
- Athey, S. (2001), “Single crossing properties and the existence of pure strategy equilibria in games of incomplete information.” *Econometrica*, 69, 861–889. [551, 557, 558, 559]
- Bajari, P., C. L. Benkard, and J. Levin (2007), “Estimating dynamic models of imperfect competition.” *Econometrica*, 75, 1331–1370. [550]
- Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2009), “Semiparametric estimation of a dynamic game of incomplete information.” Working paper, University of Minnesota. [550, 554]
- Chen, X., O. B. Linton, and I. van Keilegom (2003), “Estimation of semiparametric models when the criterion function is not smooth.” *Econometrica*, 71, 1591–1608. [573]
- Cheng, G. and J. Z. Huang (2010), “Bootstrap consistency for general semiparametric M-estimation.” *The Annals of Statistics*, 38, 2884–2915. [573]
- Domínguez, M. A. and I. N. Lobato (2004), “Consistent estimation of models defined by conditional moment restrictions.” *Econometrica*, 72, 1601–1615. [551]
- Doraszelski, U. and M. Satterthwaite (2010), “Computable Markov-perfect industry dynamics.” *RAND Journal of Economics*, 41, 215–243. [551]
- Edlin, A. and C. Shannon (1998), “Strict monotonicity in comparative statics.” *Journal of Economic Theory*, 81, 201–219. [559]

Gowrisankaran, G., C. Lucarelli, P. Schmidt-Dengler, and R. Town (2010), "Government policy and the dynamics of market structure: Evidence from Critical Access Hospitals." Working paper. [550]

Hansen, B. (2008), "Uniform convergence rates for kernel estimation with dependent data." *Econometric Theory*, 24, 726–748. [575]

Hong, H., A. Mahajan, and D. Nekipelov (2010), "Extremum estimation and numerical derivatives." Working paper, Stanford University. [574]

Hong, H. and M. Shum (2010), "Pairwise-difference estimation of a dynamic optimization model." *Review of Economics Studies*, 77, 273–304. [559]

Hotz, V. and R. A. Miller (1993), "Conditional choice probabilities and the estimation of dynamic models." *Review of Economic Studies*, 60, 497–531. [549, 550, 563, 564, 567, 570, 579]

Hotz, V., R. A. Miller, S. Sanders, and J. Smith (1994), "A simulation estimator for dynamic models of discrete choice." *Review of Economic Studies*, 61, 265–289. [567]

Kasahara, H. and K. Shimotsu (2008a), "Pseudo-likelihood estimation and bootstrap inference for structural discrete Markov decision models." *Journal of Econometrics*, 146, 92–106. [573, 574]

Kasahara, H. and K. Shimotsu (2008b), "Nonparametric identification of finite mixture models of dynamic discrete choices." *Econometrica*, 77, 135–175. [580]

Kasahara, H. and K. Shimotsu (2012), "Sequential estimation of structural models with a fixed point constraint." *Econometrica*, 80, 2303–2319. [574, 580]

Khan, S. and E. Tamer (2009), "Inference on endogenously censored regression models using conditional moment inequalities." *Journal of Econometrics*, 152, 104–119. [551]

Kreyszig, E. (1989), *Introduction to Functional Analysis With Applications*. Wiley, New York. [555]

Maskin, E. and J. Tirole (2001), "Markov perfect equilibrium I: Observable actions." *Journal of Economic Theory*, 100, 191–219. [554]

Mason, R. and A. Valentinyi (2007), "The existence and uniqueness of monotone pure strategy equilibrium in Bayesian games." Discussion Paper Series in Economics and Econometrics, Southampton University. [560]

Milgrom, P. and C. Shannon (1994), "Monotone comparative statics." *Econometrica*, 62, 157–180. [558]

Pakes, A. and P. McGuire (1994), "Computing Markov perfect Nash equilibrium: Numerical implications of a dynamic differentiated product model." *RAND Journal of Economics*, 25, 555–589. [567]

Pesendorfer, M. and P. Schmidt-Dengler (2003), "Identification and estimation of dynamic games." Working paper, NBER. [570]

Pesendorfer, M. and P. Schmidt-Dengler (2008), "Asymptotic least squares estimator for dynamic games." *Reviews of Economics Studies*, 75, 901–928. [550, 556]

Pesendorfer, M. and P. Schmidt-Dengler (2010), "Sequential estimation of dynamic discrete games: A comment." *Econometrica*, 78, 833–842. [574]

Prakasa Rao, B. L. S. (1983), *Nonparametric Functional Estimation*. Probability & Mathematical Statistics Monograph. Academic Press, New York. [575]

Ranga Rao, R. (1962), "Relations between weak and uniform convergence of measures with applications." *Annals of Mathematical Statistics*, 33, 659–680. [571]

Rust, J. (1987), "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." *Econometrica*, 55, 999–1033. [554, 560, 578]

Rust, J. (1994), "Estimation of dynamic structural models: problems and prospects part I: Discrete decision processes." In *Proceedings of the 6th World Congress of the Econometric Society*, Cambridge University Press, Cambridge. [549]

Ryan, S. (2012), "The costs of environmental regulation in a concentrated industry." *Econometrica*, 80, 1019–1061. [550]

Santos, C. D. (2010), "Sunk costs of R&D, trade and productivity: The moulds industry case." Working paper, University of Alicante. [550]

Schrumpf, P. (2011), "Identification and estimation of dynamic games with continuous states and controls." Working paper, University of British Columbia. [550]

Srisuma, S. (2010), "Minimum distance estimation for a class of Markov decision processes." Working paper, LSE. [561, 573]

Srisuma, S. and O. B. Linton (2012), "Semiparametric estimation of Markov decision processes with continuous state space." *Journal of Econometrics*, 166, 320–341. [552, 554, 566]

Tamer, E. (2003), "Incomplete simultaneous discrete response model with multiple equilibria." *Review of Economic Studies*, 70, 147–165. [550]

Tausky, O. (1949), "A recurring theorem on determinants." *American Mathematical Monthly*, 56, 672–676. [562]

Topkis, D. M. (1998), *Supermodularity and Complementarity*. Princeton University Press, Princeton. [557]

van der Vaart, A. W. and J. A. Wellner (1996), *Weak Convergence and Empirical Processes*. Springer-Verlag, New York. [571]