Identification and Estimation of Dynamic Discrete Games

Victor Aguirregabiria* Boston University

Pedro Mira* CEMFI, Madrid

May 2002

Abstract

This paper studies three econometric issues which appear in the estimation of dynamic discrete games: (1) the identification of strategic interactions among players; (2) the existence of multiple equilibria, that makes maximum likelihood estimation unfeasible even in relatively simple models; and (3) the exponential degree of complexity in the solution and estimation. First, we present conditions for nonparametric identification of players' payoff functions which apply regardless multiplicity of equilibria. Second, we propose a simple pseudo-maximum likelihood estimator (PMLE) and prove its consistency and asymptotic normality. We present also a version of this estimator that exploits randomization techniques to approximate value functions. This estimator can be computed in polynomial time, and it is asymptotically equivalent to the PMLE if $\sqrt{T}/R \to 0$ as $T \to \infty$, where T is the sample size and R is the number of simulations.

Keywords: Dynamic discrete games; Multiple equilibria; Identification; Simulation based estimation; Randomization.

JEL Classification: C13, C35, C63, C73.

Contact address: Victor Aguirregabiria. Boston University. Department of Economics. 270 Bay State Road. Boston, MA 02215. E-mail: vaguirre@bu.edu

*We want to thank comments received from Lanier Benkard, Michael Keane, Michael Manove, Andrea Moro, Antonio Merlo, and Steve Stern, as well as from seminar participants at the University of Virginia, and at the conference on *Numerically Intensive Economic Policy Analysis* at Queen's University.

1 Introduction

Empirical discrete games have proved to be useful econometric tools that provide parsimonious descriptions of strategic and social interactions among economic agents. The range of recent applications includes, among others, models of firms' entry (Bresnahan and Reiss, 1990, Berry, 1992, Reiss, 1996, Toivanen and Waterson, 2000), firms' spatial competition (Davis, 2000, and Seim, 2000), family labor supply (Kooreman, 1994), and models with social interactions (see Brock and Durlauf, 2001). However, two main problems have limited the scope of applications to relatively simple static games with a small number of players and choice alternatives: the problem of multiple equilibria, and the "curse of dimensionality" associated with the solution of these models. This paper studies these problems in a class of dynamic discrete games and develops techniques for the estimation of players' preferences. The paper analyzes also the identification of strategic interactions (or social interactions) in these games, i.e., the so called *reflection problem* (Manski, 1993). We show that certain types of dynamic games provide exclusion restrictions that can be exploited to identify nonparametrically strategic interactions. The rest of this introductory section describes previous work on these issues and the contribution of this paper.

Multiple equilibria. Multiple equilibria is a prevalent feature in most empirical games, discrete or not, where best response functions are non-linear in other players' actions. Models with multiple equilibria do not have a unique reduced form and this indeterminacy may pose practical and theoretical problems in the estimation of structural parameters. In particular, extremum estimators based on criterion functions which depend on reduced form probabilities, like maximum likelihood, can be unfeasible. Nevertheless, in games where unobservables appear additively in best response functions there is a relatively standard estimation procedure: orthogonality between unobservables and observable exogenous variables provides moment conditions that can be used to obtain consistent estimates of structural parameters. In fact, continuous differentiability and strict monotonicity of unobservables in best response functions is a sufficient condition to apply this approach.[1] However, these conditions

[1] Under continuous differentiability and strict monotonicity of unobservables in best response

do not hold in discrete games.

Econometricians have used two main approaches to deal with this problem in discrete games. A first approach has been to impose restrictions which guarantee equilibrium uniqueness. In particular, if strategic interactions have a recursive structure (i.e., player 1's payoff does not depend on other players' actions; player 2's payoff depends on player 1's action but not on other players' actions; etc) the equilibrium is unique and the model has a well-defined reduced form (see Heckman, 1978, Blundell and Smith, 1994, and Kooreman, 1994). However, this assumption imposes strong restrictions on players' strategic interactions which are not plausible in many applications. Furthermore, as shown by Tamer (2000), these restrictions may not be necessary for the identification of the model.[2] A second approach consists on exploiting only those predictions of the game which are invariant across the multiple equilibria. For instance, Bresnahan and Reiss (1990) and Berry (1992) consider static models of firms' entry where the number of entrants (but not their identity) is constant over the multiple equilibria. Their estimator maximizes a likelihood for the number of entrants, regardless their identity. However, this solution is problem-specific since an invariant function of the outcome may not always be available. It also implies a loss of efficiency since not all of the information contained in the sample is used. More importantly, this loss of information can make impossible to identify some parameters of interest.

The estimation approach that we proposed in this paper is based on a representation of equilibria in the space of players' choice probabilities. That is, equilibrium choice probabilities solve an equilibrium mapping that we call *best response probability mapping*. In a first stage, we obtain nonparametric (or semiparametric) estimates of equilibrium choice probabilities. This first stage identifies the equilibrium that players in our data are actually playing. In a second stage, our estimator of structural parameters maximizes a pseudo likelihood function where players' probabilities are

functions it is possible to invert these functions to obtain an expression where unobservables appear additively. This is the estimation approach proposed by Berry (1994) and implemented by Berry, Levinshon and Pakes (1995) in the context of Bertrand games in differentiated product markets.

[2]Notice also that in the context of dynamic games this recursive structure in one-period payoff functions does not guarantee the uniqueness of Markov Perfect equilibrium. The reason is that intertemporal payoffs do not have this recursive structure even when one-period payoffs do.

best responses to the equilibrium probabilities estimated in the first stage. We prove consistency and asymptotic normality of this pseudo maximum likelihood (PML) estimator. The procedure can be applied repeatedly in order to get a K-stage estimator with better statistical properties than the two-stage estimator. Two assumptions play an important role in our econometric approach: (1) players have incomplete information about other players' state variables, which allows us to represent Markov Perfect equilibria in terms of players' choice probabilities; and (2) players, or nature, do not randomize among multiple equilibria, which implies that population choice probabilities are equilibrium probabilities and not a mixture of equilibrium probabilities.

Curse of dimensionality. Computational costs in the solution and estimation of these models have also limited the range of empirical applications to static models with a relatively small number of players and choice alternatives. The cost of computing an equilibrium in these models increases exponentially with the number of players. Therefore, though single-agent models with discrete decision and state variables can be solved in polynomial time, the solution of discrete games requires exponential time (i.e., there is a "curse of dimensionality"). Recent work by Pakes and McGuire (1994 and 2002) provides an efficient stochastic algorithm to compute Markov-Perfect equilibria in dynamic discrete games. However, the typical *nested fixed-point algorithms* used to estimate single agent dynamic models and static games (see Berry, 1992, Rust, 1994, or Seim, 2000) require the repeated solution of the model for each trial value of the vector of parameters to be estimated. Therefore, the cost of estimating these models using these nested algorithms is several orders of magnitude larger than solving the model just once.

We show that the cost of implementing our PML estimator is of the same order of magnitude as one Newton iteration for solving the dynamic game. Therefore, this cost is several orders of magnitude smaller than the cost of a nested fixed-point algorithm. However, just one Newton iteration can be computationally very expensive for some dynamic games, and its degree of complexity is also exponential in the number of players. For that reason, we also propose a version of the PML estimator that exploits randomization techniques to approximate value functions (see Rust, 1997, and Pakes and McGuire, 2002). This estimator can be calculated in polynomial time, and it is asymptotically equivalent to the PML estimator if $\sqrt{T}/R \to 0$ as $T \to \infty$, where T is the sample size and R is the number of simulations. In this sense, this method breaks the "curse of dimensionality" in the estimation of this class of models.

Identification and reflection problem. We show that, given the discount factor and the probability distribution of unobservables, expected payoffs (i.e., players' payoff functions integrated over the probability distribution of other players' equilibrium strategies) are nonparametrically identified. This identification holds with or without multiple equilibria as long as players (or nature) do not switch over time among different equilibria. Given expected payoffs, the identification of payoff functions (i.e., not integrated over other players' equilibrium strategies) is subjected to the so called *reflection problem* (Manski, 1993), and therefore exclusion restrictions are needed. We discuss exclusion restrictions that appear in the context of dynamic games.

The rest of the paper is organized as follows. Section 2 describes the class of models considered in this paper, presents the basic assumptions, and defines a Markov-perfect equilibrium in this class of models. Section 3 discusses the identification of primitives. Section 4 presents our pseudo maximum likelihood estimator (PMLE). In section 5, we analyze the degree complexity in the implementation of the PMLE and propose a randomized version of this estimator. We conclude and summarize in section 6. Proofs of different results are provided in the Appendix.

2 A dynamic discrete game

This section presents a dynamic discrete game with incomplete information. In order to make some of the discussions more intuitive we consider a model where a finite number of companies decide the number of outlets to operate in a local market and the spatial location of these outlets.[3]

2.1 Framework and basic assumptions

The market is a city divided into a finite number of non-overlapping locations or cells. Locations are characterized by demand conditions which can change over time

[3] Davis (2000) and Seim (2000) consider similar models in a static context.

(e.g., number of households, distribution of household income, age, etc). Let d_t be the vector with the characteristics of all locations at period t. There are N firms operating in the market, which we index by $i \in I = \{1, 2, ...N\}$. At every discrete period t firms decide simultaneously how many stores to operate and their locations. Profits are bounded from above such that the maximum number of stores is finite. Therefore, a firm's set of choice alternatives, A, is discrete and finite. We represent the decision of firm i at period t by the variable $a_{it} \in A = \{1, ..., J\}$.

At the beginning of period t a firm is characterized by two vectors of state variables, x_{it} and ε_{it}, which affect its profitability. Variables in x_{it} are common knowledge for all firms in the market, but the vector ε_{it} is private information of firm i.[4] Let $x_t \equiv (d_t, x_{1t}, x_{2t}, ..., x_{Nt})$ and $\varepsilon_t \equiv (\varepsilon_{1t}, \varepsilon_{2t}, ..., \varepsilon_{Nt})$ be the vectors of common knowledge and private information variables, respectively. A firm's current profits depends on x_t, on its own private information ε_{it}, and on the vector of firms' current decisions, $a_t \equiv (a_{1t}, a_{2t}, ..., a_{Nt})$. Let $\tilde{\Pi}_i(a_t, x_t, \varepsilon_{it})$ be firm i's current profit function. We assume that $\{x_t, \varepsilon_t\}$ follows a controlled Markov process with transition probability $p(x_{t+1}, \varepsilon_{t+1} | a_t, x_t, \varepsilon_t)$. This transition probability is common knowledge.

A firm decides its network of stores to maximize expected discounted intertemporal profits:

$$E\left\{\sum_{s=t}^{\infty} \beta^{s-t} \; \tilde{\Pi}_i(a_s, x_s, \varepsilon_{is}) \mid x_t, \varepsilon_{it}\right\} \tag{1}$$

where $\beta \in (0, 1)$ is the discount factor. The primitives of the model are the profits functions $\{\tilde{\Pi}_i(.) : i = 1, 2, ...N\}$, the transition probability $p(.|.)$, and the discount factor β. We consider the following assumptions on these primitives.

ASSUMPTION 1: Private information appears additively in the profit function. That is,

$$\tilde{\Pi}_i(a_t, x_t, \varepsilon_{it}) = \Pi_i(a_t, x_t) + \varepsilon_{it}(a_{it})$$

where $\Pi_i(.)$ is a real valued function, and $\varepsilon_{it} \equiv \{\varepsilon_{it}(1), \varepsilon_{it}(2), ..., \varepsilon_{it}(J)\} \in R^J$ is a vector of real valued random variables.

[4] For instance, some variables which could enter in x_{it} are the firm's network of outlets at previous period or the firm's previous market share. Managerial ability at different outlets could be a component of ε_{it}.

ASSUMPTION 2: The transition probability $p(x_{t+1}, \varepsilon_{t+1}|a_t, x_t, \varepsilon_t)$ factors as:

$$p(x_{t+1}, \varepsilon_{t+1}|a_t, x_t, \varepsilon_t) = p_\varepsilon(\varepsilon_{t+1})\ f(x_{t+1} \mid a_t, x_t)$$

That is: (1) given firms' decisions at period t, private information variables do not affect the transition of common knowledge variables; (2) future realizations of private information variables are independent of current decisions and common knowledge variables; and (3) private information variables are independently and identically distributed over time. We also assume that private information is independently distributed across firms: $p_\varepsilon(\varepsilon_t) = \prod_{i=1}^N g_i(\varepsilon_{it})$, where, for any firm i, $g_i(.)$ is a density function which is absolutely continuous with respect to the Lebesgue measure.

ASSUMPTION 3: Common knowledge variables have a discrete and finite support: $x_t \in X \equiv \{x^1, x^2, ..., x^{|X|}\}$ where $|X|$ is a finite number.

2.2 Strategies and Bellman equations

The game has a Markov structure, and we assume that firms play Markov strategies. That is, if $\{x_t, \varepsilon_{it}\} = \{x_s, \varepsilon_{is}\}$ then firm i's decisions at periods t and s are the same. Therefore, we can omit the time subindex and use x' and ε' to denote next period state variables. Let $\alpha = \{\alpha_i(x, \varepsilon_i)\}$ be a set of strategy functions or decision rules, one for each firm, with $\alpha_i : X \times R^J \rightarrow A$. Associated with a set of strategy functions α we can define a set of *conditional choice probabilities* $P^\alpha = \{P_i^\alpha(a_i|x)\}$ such that,

$$P_i^\alpha(a_i|x) \equiv \Pr\left[\alpha_i(x, \varepsilon_i) = a_i\ |x\right] = \int I\left\{\alpha_i(x, \varepsilon_i) = a_i\right\}\ g_i(\varepsilon_i)\ d\varepsilon_i \qquad (2)$$

where $I\{.\}$ is the indicator function. The probabilities $\{P_i^\alpha(a_i|x) : a_i \in A\}$ represent the expected behavior of firm i from the point of view of the rest of the firms when firm i follows its strategy in α.

Let $\pi_i(a_i, x; P_{-i}^\alpha)$ be firm i's expected profit if it chooses alternative a_i and the other firms behave according to their respective strategies in α.[5] Since private information

[5] In the terminology of Harsanyi (1995), the profits functions $\Pi_i(a_1, a_2, ..., a_N, x)$ are *conditional payoff functions*, and the expected profits $\pi_i(a_i, x; P_{-i}^\alpha)$ are *semi-conditional payoff functions*.

variables are independent across firms, it is clear that:[6]

$$\pi_i(a_i, x; P^\alpha_{-i}) = \sum_{a_{-i} \in A^{N-1}} \left(\prod_{j \neq i} P^\alpha_j(a_j|x) \right) \Pi_i(a_i, a_{-i}, x) \tag{3}$$

Let $\tilde{V}_i(x, \varepsilon_i; P^\alpha_{-i})$ be the value of firm i if this firm behaves optimally now and in the future given that the other firms follow their strategies in α. By Bellman's principle of optimality, we can write:

$$\tilde{V}_i(x, \varepsilon_i; P^\alpha_{-i}) = \max_{a_i \in A} \left\{ \pi_i(a_i, x; P^\alpha_{-i}) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} \left[\int \tilde{V}_i(x', \varepsilon'_i; P^\alpha_{-i}) \ g(\varepsilon'_i) \ d\varepsilon'_i \right] \bar{f}_i(x'|x, a_i; P^\alpha_{-i}) \right\} \tag{4}$$

where $\bar{f}_i(x'|x, a_i; P^\alpha_{-i})$ is the transition probability of x conditional on firm i choosing a_i and the other firms behaving according to α:

$$\bar{f}_i(x'|x, a_i; P^\alpha_{-i}) = \sum_{a_{-i} \in A^{N-1}} \left(\prod_{j \neq i} P^\alpha_j(a_j|x) \right) f(x'|x, a_i, a_{-i}) \tag{5}$$

It is convenient to define value functions integrated over private information variables. Let $V_i(x; P^\alpha_{-i})$ be the integrated value function $\int \tilde{V}_i(x, \varepsilon_i; P^\alpha_{-i}) \ g(d\varepsilon_i)$. Based on this definition and equation (4), we can obtain the *integrated Bellman equation*:

$$V_i(x; P^\alpha_{-i}) = \int \max_{a_i \in A} \left\{ \pi_i(a_i, x; P^\alpha_{-i}) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} V_i(x'; P^\alpha_{-i}) \ \bar{f}_i(x'|x, a_i; P^\alpha_{-i}) \right\} g_i(d\varepsilon_i) \tag{6}$$

The right hand side of equation (6) is a contraction mapping in the space of value functions (see Rust et al, 2002, and Aguirregabiria and Mira, 2002). Therefore, for each firm, there is a unique function $V_i(x; P^\alpha_{-i})$ that solves this functional equation for given α.

2.3 Markov-perfect equilibria

Let $\phi_i(x, \varepsilon_i; P^\alpha_{-i})$ be the best response of firm i if other players follow their strategies in α. That is,

$$\phi_i(x, \varepsilon_i; P^\alpha_{-i}) = \arg\max_{a_i \in A} \left\{ v_i(a_i, x; P^\alpha_{-i}) + \varepsilon_i(a_i) \right\} \tag{7}$$

[6] In a static game, firm i's *best response* to other firms' strategies in α is $\arg\max_{a_i \in A}$ $\{\pi_i(a_i, x; P^\alpha_{-i}) + \varepsilon_i(a_i)\}$, and we can use these best response functions to define (Bayesian) Nash equilibria.

where $\{v_i(a_i, x; P_{-i}^{\alpha})\}$ are the conditional choice value functions:

$$v_i(a_i, x; P_{-i}^{\alpha}) \equiv \pi_i(a_i, x; P_{-i}^{\alpha}) + \beta \sum_{x' \in X} V_i(x'; P_{-i}^{\alpha}) \ \bar{f}_i(x'|x, a_i; P_{-i}^{\alpha}) \tag{8}$$

So far α is arbitrary and does not necessarily describe the equilibrium behavior of other firms. The following definition characterizes equilibrium strategies of all firms as best responses to one another.

DEFINITION: A stationary Markov-perfect Equilibrium (MPE) in this game is a set of strategy functions α^* such that for any firm i and for any $(x, \varepsilon_i) \in X \times R^J$,

$$\alpha_i^*(x, \varepsilon_i) \quad = \quad \phi_i(x, \varepsilon_i; P_{-i}^{\alpha*}) \tag{9}$$

Following Milgrom and Weber (1985) we can also represent a MPE in probability space.[7] Let α^* be a set of MPE strategies, and let P^* be the set of conditional choice probabilities associated with these strategies. Then, $P^* = \Phi(P^*)$, where for any vector of probabilities P, $\Phi(P) = \{\Phi_i(a_i|x; P_{-i})\}$, and:

$$\Phi_i(a_i|x; P_{-i}) = \Lambda_i(a_i|\{v_i(j, x; P_{-i}) : j \in A\}) \tag{10}$$

where $\Lambda_i(a_i|\{c_j : j \in A\})$ is the operator $\int I\left(a_i = \arg\max_{j \in A}\left[c_j + \varepsilon_i(j)\right]\right) g_i(d\varepsilon_i)$. We call the functions Φ_i *best response probability functions.* Given our assumptions on the distribution of private information, best response probability functions are continuous in P, and therefore an equilibrium in probabilities exits. In general, the equilibrium is not unique.

Equilibrium probabilities solve the *coupled* fixed-point problem defined by (6) and (10). Given a set of probabilities P we obtain value functions V_i as solutions of the N fixed point problems in (6); and given these value functions we obtain best response probabilities using the right hand-side of equation (10).

2.4 An alternative equilibrium mapping

Based on results in Aguirregabiria and Mira (2002), we now define an equilibrium mapping that avoids the solution of the N fixed point problems in (6). This alternative equilibrium mapping is an important building block in our estimation method.

[7] Milgrom and Weber consider both discrete-choice and continuous-choice games. In their terminology $\{P_i^{\alpha}\}$ are called *distributional strategies*, and P^* is an *equilibrium in distributional strategies.*

Let $P^* \equiv \{P_i(a_i|x)\}$ be an equilibrium in choice probabilities, and let V_1^*, V_2^*, ..., V_N^* be the firms' value functions associated with this equilibrium. Since equilibrium probabilities are best responses, we can rewrite the Bellman equation (6) as,

$$V_i^*(x) = \sum_{a_i \in A} P_i^*(a_i|x) \left[\pi_i(a_i, x; P_{-i}^*) + e_i(a_i; P_i^*(x))\right] + \beta \sum_{x' \in X} V_i^*(x') \; \bar{f}(x'|x; P^*) \quad (11)$$

where $\bar{f}(x'|x; P^*)$ is the transition probability of x induced by P^*.[8] The term $e_i(a_i; P_i^*(x))$ is the expectation of $\varepsilon_i(a)$ conditional on x and on alternative a_i being the optimal response for player i. By Hotz-Miller proposition (Hotz and Miller, 1993), these conditional expectations are only functions of a_i and $P_i^*(x)$.

Taking equilibrium probabilities as given, expression (11) describes the vector of values V_i^* as the solution of a system of linear equations. In vector form:

$$\left[I_{|X|} - \beta \; F(P^*)\right] V_i^* = \sum_{a_i \in A} P_i^*(a_i) * \left[\pi_i(a_i; P_{-i}^*) + e_i(a_i; P_i^*)\right] \quad (12)$$

where $I_{|X|}$ is the identity matrix; $F(P^*)$ is a matrix with transition probabilities $\tilde{f}(x'|x; P^*)$; and $P_i^*(a_i)$, $\pi_i(a_i; P_{-i}^*)$ and $e_i(a_i; P_i^*)$ are vectors of dimension $|X|$ with the obvious definitions. Let $\Gamma_i(P^*) \equiv \{\Gamma_i(x; P^*) : x \in X\}$ be the solution to this system of linear equations, such that $V_i^*(x) = \Gamma_i(x; P^*)$. For arbitrary probabilities P, not necessarily in equilibrium, $\Gamma_i(.)$ can be interpreted as a *valuation operator:* that is, $\Gamma_i(x; P)$ is the *expected value of firm i if all firms (including firm i) behave today and in the future according to their choice probabilities in P*. Therefore, we can represent the equilibrium condition as $P^* = \Psi(P^*)$, where $\Psi(P) \equiv \{\Psi_i(a_i|x; P)\}$ and:

$$\Psi_i(a_i|x; P) = \Lambda_i(a_i|\{\hat{v}_i(j, x; P)\}) \quad (13)$$

where now the conditional choice values $\hat{v}_i$ are based on the operator $\Gamma_i(x; P)$ instead of on the values $V_i(x; P_{-i})$, i.e., $\hat{v}_i(j, x; P) \equiv \pi_i(j, x; P_{-i}) + \beta \sum_{x' \in X} \Gamma_i(x; P) \; \bar{f}_i(x'|x, j; P_{-i})$.

The only difference between equilibrium mappings Φ_i and Ψ_i is that in the second mapping we use the valuation operators $\Gamma_i(x; P)$ instead of the values $V_i(x; P_{-i})$. To evaluate Φ_i one has to solve N dynamic programming problems, but to obtain Γ_i one

[8] That is, $\bar{f}(x'|x; P^*) = \sum_{a \in A^N} \left(\prod_{i=1}^{N} P_i^*(a_i|x)\right) f(x'|x, a) = \sum_{a_i \in A} P_i^*(a_i|x) \bar{f}_i(x'|x, a_i)$.

has to solve N systems of linear equations. Of course, this does not necessarily mean that solving for an equilibrium using mapping Ψ is cheaper that using mapping Φ, because the number of iterations needed with each of these operators can be different. However, in the context of the estimation of the model, we will see that using mapping Ψ instead of Φ can provide significant computational gains.

2.5 An example: A conditional logit dynamic game

We illustrate the previous definitions and results in the context of a dynamic game version of McFadden's conditional logit (McFadden, 1984). Payoff functions have the following form:

$$\tilde{\Pi}_i(a, x, \varepsilon_i) = y_1(a_i, x_i)\ \theta_1 + y_2(a_i, x_{-i})\ \theta_2 + y_3(a_i, a_{-i})\ \theta_3 + \varepsilon_i(a_i) \tag{14}$$

where θ_1, θ_2, and θ_3 are column vectors of structural parameters, and $y_1(.)$, $y_2(.)$, and $y_3(.)$ are row vectors of functions of decision and state variables. Following Manski's terminology: (1) the term $y_1(a_i, x_i)\theta_1$ captures *correlated effects*; (2) the component $y_2(a_i, x_{-i})\theta_2$ represents *contextual effects*; and (3) the term $y_3(a_i, a_{-i})\theta_3$ captures strategic interactions or *endogenous effects*. Private information variables $\{\varepsilon_i(a) : a \in A\}$ are independently and identically distributed over time, over individuals and over choice alternatives with Extreme value type 1 distribution.

Expected payoffs induced by a vector of probabilities P are:

$$\pi_i(a_i, x; P_{-i}) = y_1(a_i, x_i)\ \theta_1 + y_2(a_i, x_{-i})\ \theta_2 + s_i\left[a_i; P_{-i}(x)\right]\theta_3 \tag{15}$$

where $s_i\left[a_i; P_{-i}(x)\right] = \sum_{a_{-i}} \left(\prod_{j \neq i} P_j(a_j|x)\right) y_3(a_i, a_{-i})$. Given the extreme value distribution of ε, the (integrated) Bellman equation is:

$$V_i(x; P_{-i}) = \ln\left[\sum_{j=1}^{J} \exp\left\{\pi_i(j, x; P_{-i}) + \beta \sum_{x'} V_i(x'; P_{-i})\ \bar{f}_i(x'|x, j; P_{-i})\right\}\right] \tag{16}$$

And best response probability functions are:

$$\Phi_i(a_i|x; P_{-i}) = \frac{\exp\left\{\pi_i(a_i, x; P_{-i}) + \beta \sum_{x'} V_i(x'; P_{-i})\ \bar{f}_i(x'|x, a_i; P_{-i})\right\}}{\sum_{j=1}^{J} \exp\left\{\pi_i(j, x; P_{-i}) + \beta \sum_{x'} V_i(x'; P_{-i})\ \bar{f}_i(x'|x, j; P_{-i})\right\}} \tag{17}$$

We show in the Appendix that in this model the equilibrium mapping Ψ_i has the following form:

$$\Psi_i(a|x;P) = \frac{\exp\{\ \tilde{y}_{1i}(a,x;P)\ \theta_1 + \tilde{y}_{2i}(a,x;P)\ \theta_2 + \tilde{s}_i(a,x;P)\ \theta_3 + \tilde{e}_i(a,x;P)\ \}}{\sum_{j=1}^{J} \exp\{\ \tilde{y}_{1i}(j,x;P)\ \theta_1 + \tilde{y}_{2i}(j,x;P)\ \theta_2 + \tilde{s}_i(j,x;P)\ \theta_3 + \tilde{e}_i(j,x;P)\}} \tag{18}$$

where the values $\tilde{y}_{1i}$, $\tilde{y}_{2i}$, $\tilde{s}_i$ and $\tilde{e}_i$ can be obtained from choice and transition probabilities, without knowing the structural parameters $\{\theta_1, \theta_2, \theta_3\}$. In Section 4, we exploit this property to obtain a simple estimator of structural parameters.

3 Econometric identification of primitives

3.1 Data and data generating process

A researcher observes firms' actions and common knowledge state variables at T independent markets: $\{a_m, x_m, x'_m : m = 1, 2, ..., T\}$ where m is the market subindex, and T is the number of markets, which is large.[9] For the moment, we consider that the econometrician observes all common knowledge state variables, but we relax this assumption later in this section. The researcher is interested in the estimation of the primitives of the model $\theta \equiv \{\Pi_i, g_i, f, \beta : i \in I\}$. Under Assumption 2 the transition probability f can be identified from transition data: $f(x'|x,a) = \Pr(x'_m = x'|x_m = x, a_m = a)$. Therefore, in this section we treat f as known and study the identification of the rest of the primitives from the conditional distribution $\Pr(a_m|x_m)$.

Let $P^0 \equiv \{\Pr(a_m = a|x_m = x) : (a,x) \in A^N \times X\}$ be the *true* conditional distribution of a_m in the population. Let D be a mapping from the space of primitives Θ to the space of distributions Ω such that $D(\theta)$ is the set of Markov-perfect equilibria (in choice probabilities) associated with θ. In general, there will be values of the primitives for which the model has multiple equilibria. This implies that D is not a function but a correspondence. In this context, how the population P^0 has been generated? Is it the result of randomization over the different equilibria in some

[9] We consider asymptotics in the number of markets because this is the most common framework in empirical applications in industrial organization. However, all the results in the paper can be extended to the case of asymptotics in the number of time periods with a small number of players and one market.

particular set $D(\theta^0)$? Do players (or nature) select different elements of $D(\theta^0)$ at different markets or at different periods of time? We consider the following assumption about the *data generating process* or *equilibrium selection device.*

ASSUMPTION 4 (No sunspots): Given a value for the primitives of the model, $\theta = \{\Pi_i, g_i, f, \beta : i \in I\}$*, players (or nature) select only one equilibrium in* $D(\theta)$ *and they do not switch to other equilibria as long as* θ *does not change.*

This assumption implies that there is at least one value θ^0 such that $P^0 \in D(\theta^0)$. In contrast, when players switch or randomize among different equilibria, P^0 is a mixture of the equilibria in $D(\theta^0)$ and this mixture does not belong to $D(\theta^0)$. Assumption 4 plays an important role in the identification results and in the estimation method that we present in this paper.

3.2 Identification of payoff functions

We present sufficient conditions for the nonparametric identification of payoff functions $\{\Pi_i(a, x)\}$ taking the discount factor and the distribution of players' private information as given. We proceed in two parts. First, Proposition 1 establishes sufficient conditions for nonparametric identification of expected payoff functions $\{\pi_i(a_i, x; P^0_{-i})\}$. Then, we consider the identification of payoffs $\{\Pi_i(a, x)\}$ taking expected profits as given.

PROPOSITION 1: Consider the following conditions: (1) Assumptions 1 to 4; (2) $P^0 \equiv \{\Pr(a_m = a | x_m = x) : (a, x) \in A^N \times X\}$ *is nonparametrically identified; (3)* $\{g_i : i \in I\}$ *and* β *are known; and (4) We impose the normalization* $\Pi_i(a_i = J, a_{-i}, x) = 0$ *for any* (i, a_{-i}, x)*. Under these conditions, expected payoffs* $\{\pi_i(a_i, x; P^0_{-i})\}$ *are nonparametrically identified.*

Notice that this Proposition holds with or without multiple equilibria. That is, under Assumption 4 the conditions for the identification of expected payoffs (and for that matter, of payoffs) do not depend on whether the model has multiple equilibria or not.

Condition (2) holds if the econometrician observes $\{a, x\}$ over its full support $A^N \times X$. However, P^0 can be nonparametrically identified under weaker conditions.

In particular, P^0 can be identified when the econometrician does not observe some common knowledge state variables. Let $x_m = \{z_m, \xi_m\}$, where z_m is observable to the researcher and ξ_m is unobservable.

ASSUMPTION 5: ξ_m is independent of z_m and is independently and identically distributed over markets with probability distribution $\lambda(\xi)$. Furthermore, ξ_m has discrete and finite support $\Upsilon = \{\xi^1, \xi^2, ..., \xi^L\}$, where $\lambda(\xi) > 0$ for any $\xi \in \Upsilon$.

It is clear that $\tilde{P}^0 \equiv \{\Pr(a_m = a | z_m = z) : (a, x) \in A^N \times Z\}$ is identified from the data. But Proposition 1 requires the identification of P^0. To understand why P^0 can be identified in this model, notice that the independence of private information variables imposes restrictions on the probabilities P^0: i.e., conditional on common knowledge variables, players' actions should be independent. Therefore, once we condition on observable state variables z_m, the spatial correlation between the actions of any two players should be explained by unobservable market characteristics. More formally, assumptions 2 and 5 imply the following relationship between $\tilde{P}^0$ and P^0 and λ:

$$\tilde{P}^0(a|z) = \sum_{\xi \in \Upsilon} \lambda(\xi) \left[\prod_{i=1}^{N} P_i^0(a_i|x, \xi) \right] \tag{19}$$

These restrictions can be used to identify P^0 and λ from $\tilde{P}^0$. We establish sufficient conditions for identification in the following Lemma.

LEMMA 1: Consider the following conditions: (1) Assumptions 1, 2, 3 and 5; (2) $\tilde{P}^0 \equiv \{\Pr(a_m = a | z_m = z) : (a, z) \in A^N \times Z\}$ is nonparametrically identified; (3) Order condition, $L \leq int\left(J^N / [N(J-1)+1]\right)$; and (4) Rank condition, for any $(a_i, z) \in A \times Z$ and any ξ and ξ' in Υ, $P_i^0(a_i|z, \xi) \neq P_i^0(a_i|z, \xi')$. Under these conditions λ and P^0 are nonparametrically identified.

The conditions in Lemma 1 do not impose any restriction on the support of z.[10] Conditions (3) and (4) can be relaxed if we incorporate assumptions on Z and on the variability of $\tilde{P}^0$ over Z. Notice, that the order condition for identification implies that when the number of players increases we can allow for more points in the support of unobservable market characteristics. For instance, in a binary choice game we need

[10] In fact, this Lemma applies also when there are not observable variables z.

at least three players to identify P^0 and λ when $L = 2$; we need four players when $L = 3$; and we can identify a support Υ with nine points when the number of players is six.

We now consider the identification of payoff functions Π_i given choice probabilities and expected payoffs. Remember that:

$$\pi_i^0(a_i, z, \xi) \equiv \pi_i(a_i, z, \xi; P_{-i}^0) = \sum_{a_{-i}} P_{-i}^0(a_{-i}|z, \xi)\ \Pi_i(a_i, a_{-i}, z, \xi) \tag{20}$$

For each value (a_i, z, ξ) we have only one equation and J^{N-1} unknown payoffs Π_i. It is clear that without further restrictions we cannot identify Π_i from P^0 and π_i^0. First, we need an *order condition* to identify Π_i: the number of payoff values in $\{\Pi_i\}$ should not be larger than the number of expected payoffs $\{\pi_i^0\}$. However, the problem of identifying Π_i is more complicated than just imposing this order condition. It is related to the so called *reflection problem* which appears in models with social interactions (see Manski, 1993 and 1995).

To illustrate this point, we consider the specification of payoffs in the example in section 2.5. For simplicity, we incorporate the unobservable variable ξ additively.

$$\Pi_i(a, z, \xi) = y_1(a_i, z_i)\ \theta_1 + y_2(a_i, z_{-i})\ \theta_2 + y_3(a_i, a_{-i})\ \theta_3 + y_4(a_i)\ \xi \tag{21}$$

Since $\pi_i^0(a_i, z, \xi) = E(\Pi_i(a, z, \xi)|a_i, z, \xi)$, we can write $\pi_i^0(a_i, z, \xi) = \Pi_i(a, z, \xi) + u_i$, where u_i is orthogonal to z and ξ. Therefore, we have the following regression equation:

$$\pi_i^0(z, \xi) = y_1(z_i)\ \theta_1 + y_2(z_{-i})\ \theta_2 + y_3(a_{-i})\ \theta_3 + \xi + u_i \tag{22}$$

where we have omitted a_i as an argument for notational simplicity.. We know $\pi_i^0(z, \xi)$, $y_1(z_i)$, $y_2(z_{-i})$, $y_3(a_{-i})$ and the probability distribution of ξ, and we want to estimate $\theta = \{\theta_1, \theta_2, \theta_3\}$. In this regression $y_3(a_{-i})$ is endogenous because players' actions depend on the unobservables ξ and u_i. Notice that a_{-i} is endogenous in this regression even if there is not unobservable common knowledge state variables, i.e., $\xi = 0$ with probability one. In principle, we do not have exclusion restrictions to instrument $y_3(a_{-i})$ because z is an argument in the regression. Therefore, without further restrictions, the identification of θ_i relies on functional form assumptions on $y_1(.)$ and

$y_2(.)$.[11]

Suppose that we can partitioned z_i in two sub-vectors, z_i^A and z_i^B, such that: (1) P*layer i's* payoff depends on $z^A = \{z_1^A, z_2^A, ..., z_N^A\}$ and on its own z_i^B but not on z_{-i}^B, i.e., $y_2(z_{-i}) = y_2(z_{-i}^A)$; and (2) the matrix $E(z_{-i}^B\ y_3(a_{-i}))$ has full-column rank. Under these conditions, θ is identified for any specification of $y_1(.)$, $y_2(.)$, and $y_3(.)$. This exclusion restriction applies both to static and dynamic games. However, there are some dynamic games where this type of restriction seems more plausible. For instance, consider a dynamic game where firms face costs of changing their decision from period $t-1$ to period t. That is the case in models with adjustment costs or switching costs. Suppose that a firm's payoff function depends on the current actions of all firms (i.e., contemporaneous strategic interactions) and on its own previous action (i.e., adjustment costs), but not on other firms previous actions (i.e., not lagged strategic interactions). In this context, if $E(a_{-i,t-1} y_3(a_{-it}))$ is full-column rank, and strategic interactions are identified.

4 Estimation

Consider that the primitives of the model are known up to a vector of structural parameters θ, and that primitives are continuously differentiable in θ. Assumption 4 implies that there is at least one value θ^0 such that P^0 is an equilibrium associated with θ^0. Furthermore, under the conditions in Proposition 1 and the exclusion restrictions described in previous section, θ^0 is unique. The researcher wants to estimate θ^0 given a sample $\{a_m, z_m : m = 1, 2, ..., T\}$. First, we describe a maximum likelihood estimator of θ^0 and discuss its theoretical and practical limitations.

4.1 Maximum likelihood estimation

In principle, we could follow a maximum likelihood approach to estimate θ^0. Let $\Theta \subset R^K$ be the space of the parameter vector θ, where Θ is compact. For $\theta \in \Theta$, let $D(\theta)$ be the set of equilibria, in probability space, associated with θ: i.e., $D(\theta) =$

[11] Notice that a functional form assumption on $y_3(a_{-i})$ (e.g., $y_3(a_{-i}) = \sum_{j \neq i} a_j$) does not help to solve this identification problem.

$\{P^{*1}(\theta), P^{*2}(\theta), ...\}$.

DEFINITION (Equilibrium type): Let θ and θ' be two elements of Θ. And let $P(\theta) \in D(\theta)$ and $P'(\theta') \in D(\theta')$ be two equilibria associated with θ and θ', respectively. We say that $P(\theta)$ and $P'(\theta')$ belong to the same *equilibrium type* if $\lim_{\theta' \to \theta} P'(\theta') = P(\theta)$.

For simplicity, suppose that there is a finite number Q of equilibrium types, which we index by q. In general, not all the equilibrium types exist for every $\theta \in \Theta$. let $\Theta_q \subseteq \Theta$ be the subset of parameter vectors for which equilibrium q exists.

For any type q, define the following *equilibrium type-specific* log-likelihood functions:

$$l^q(\theta) = \begin{cases} \frac{1}{T} \sum_{m=1}^{T} \sum_{i=1}^{N} \ln P_i^{*q}(a_{im}|x_m;\theta) & if \quad \theta \in \Theta_q \\ -\infty & if \quad \theta \notin \Theta_q \end{cases} . \tag{23}$$

Under assumption 4, the population probabilities P^0 belong to some equilibrium type. If we knew this equilibrium type, say q_0, we would maximize $l^{q_0}(\theta)$ with respect to θ and obtain the MLE of θ^0. Assumptions 1 to 4 and the differentiability of primitives with respect to θ guarantee that $l^{q_0}(\theta)$ is a continuously differentiable function. Therefore, if Θ_{q_0} is a compact set, this estimator is consistent, asymptotically normal and efficient.

Although we do not know the equilibrium type of P^0, we can obtain the MLE of θ^0 using the following procedure. First, for any type q, let $\hat{\theta}^q$ be the maximum likelihood estimator of θ^0 conditional on P^0 being type q: that is, $\hat{\theta}^q = \arg\max_{\theta \in \Theta_q} l^q(\theta)$. Then, the maximum likelihood estimator is defined as:

$$\hat{\theta}_{MLE} = \hat{\theta}^q \Leftrightarrow l^q(\hat{\theta}^q) = \max_{q' \in \{1,2,..,Q\}} l^{q'}(\hat{\theta}^{q'}) \tag{24}$$

It is straightforward to prove that, under Assumptions 1 to 4 and the additional assumption that all the sets Θ_q are compact, this estimator is consistent, asymptotically normal and efficient. That is, under these conditions we choose $\hat{\theta}^{q_0}$ asymptotically with probability one, and $\hat{\theta}^{q_0}$ is CAN and AE.

However, this estimator has important theoretical and practical limitations. First, in models with multiple equilibria some sets Θ_q (including Θ_{q_0}) may not be compact and this can result in the inconsistency of this estimator. Second, for the implementation of this method we should know all the equilibrium types that the model has on

Θ. This is computationally impractical even for relatively simple models. And third, to obtain an *equilibrium type-specific* estimator, say $\hat{\theta}^q$, we need an algorithm that guarantees that for different values of θ we always select equilibrium type q. This can be a very difficult task for some types of equilibria (see McKelvey and McLennan, 1996).

Given these important computational problems, the previous ML estimator has not been used in empirical applications. Instead, the approach in some applications has been to maximize a "likelihood function" where, for every trial value of θ, the researcher lets the algorithm that searches for an equilibrium "decide" the probabilities that should enter in the likelihood. In general, the sample criterion obtained in this way is not a well defined likelihood function and it does not have a clear statistical interpretation. The statistical properties of the estimator that maximizes this criterion function are unknown. In fact, given that when θ changes the algorithm may "jump" between equilibrium types, this criterion function can be very discontinuous and typical theorems for extremum estimators do not apply.

4.2 Pseudo maximum likelihood estimator

Consider the following pseudo likelihood function:

$$\tilde{l}_T(P^0, \theta) = \frac{1}{T} \sum\nolimits_{m=1}^{T} \sum\nolimits_{i=1}^{N} \ln \Psi_i(a_{im}|x_m; P^0, \theta) \tag{25}$$

where Ψ is the equilibrium mapping defined in section 2. For arbitrary θ this pseudo likelihood is not equal to the likelihood associated with the equilibrium type of P^0, i.e., $\tilde{l}_T(P^0, \theta) \neq l_T^{q_0}(\theta)$. However, the true and the pseudo likelihoods are equal at $\theta = \theta^0$: i.e., by the equilibrium condition and Assumption 4, $\Psi(P^0, \theta^0) = P^{*q_0}(\theta^0)$. Based on this property, we define the following pseudo-maximum likelihood estimator. Let $\hat{P}^0$ be an initial nonparametric estimator of P^0. The pseudo maximum likelihood estimator of θ^0 is the value of θ that maximizes the pseudo likelihood $\tilde{l}_T(\theta, \hat{P}^0)$:

$$\hat{\theta}_{PMLE} = \arg\max_{\theta \in \Theta} \ \tilde{l}_T(\hat{P}^0, \theta) \tag{26}$$

The following Proposition shows that this estimator is consistent and asymptotically normal under mild regularity conditions.

PROPOSITION 2: Consider that Θ *is a compact set and that* θ^0 *uniquely maximizes* $\tilde{l}_\infty(P^0,\theta) \equiv E\left(\sum_i \ln \Psi_i(a_{im}|x_m;P^0,\theta)\right)$ *in* Θ*. Let* $\hat{P}_T^0$ *be a root-T consistent and asymptotically normal estimator of* P^0 *such that* $\left[\sqrt{T}(\hat{P}_T^0 - P^0)' \;;\; \sqrt{T}\ \partial\tilde{l}_T(\theta^0,P^0)/\partial\theta'\right] \rightarrow_d N(0,\Omega^0)$*, where* $\Omega^0 \equiv \begin{bmatrix} \Omega^0_{PP} & \Omega^0_{P\theta} \\ \Omega^0_{\theta P} & \Omega^0_{\theta\theta} \end{bmatrix}$*,* Ω^0_{PP} *is covariance matrix in the asymptotic distribution of* $\hat{P}_T^0$*, and* $\Omega^0_{\theta\theta} \equiv \partial^2\tilde{l}_\infty(P^0,\theta^0)/\partial\theta\partial\theta'$*. Under these conditions and Assumptions 1 to 4 the estimator* $\hat{\theta}_{PMLE}$ *converges a.s. to* θ^0*, it is asymptotically normal, and the variance of its asymptotic distribution is*

$$V^0 = (\Omega^0_{\theta\theta})^{-1} + (\Omega^0_{\theta\theta})^{-1}\left[\Delta^0_{\theta P}\ \Omega^0_{PP}\ \Delta^{0\prime}_{\theta P} + \Omega^0_{\theta P}\ \Delta^{0\prime}_{\theta P} + \Delta^0_{\theta P}\ \Omega^0_{P\theta}\right](\Omega^0_{\theta\theta})^{-1}$$

where $\Delta^0_{\theta P} \equiv \partial^2\tilde{l}_\infty(P^0,\theta^0)/\partial\theta\partial P'$*.*

Root-T consistency and asymptotic normality of the first stage nonparametric estimator (together with regularity conditions) are sufficient to guarantee root-T consistency and asymptotic normality of the pseudo maximum likelihood estimator. A sufficient condition for asymptotic efficiency is that the Jacobian matrix $\partial^2\Psi_i(a_{im}|x_m;P^0,\theta^0)/\partial\theta\partial P'$ is zero. To see this, notice that this condition implies that: (1) $\Delta^0_{\theta P} = 0$, and therefore $V^0 = (\Omega^0_{\theta\theta})^{-1}$; and (2) by the implicit function theorem, $\partial\Psi_i(a_{im}|x_m;P^0,\theta^0)/\partial\theta = \partial P_i^{*q_0}(a_{im}|x_m;\theta^0)/\partial\theta$, and therefore score and pseudo-score are equal at θ^0 and $V^0 = (\Omega^0_{\theta\theta})^{-1} = (\partial^2 l^{q_0}_\infty(\theta^0)/\partial\theta\partial\theta')^{-1}$, which is the covariance matrix of the MLE. In other words, under the zero Jacobian matrix condition Fisher's information matrix is block diagonal and the asymptotic variance of the first stage estimator does not affect the asymptotic variance of the second stage estimator. In Aguirregabiria and Mira (2002) we proved that this condition is satisfied for the class of discrete choice dynamic programming models considered in Rust (1987, 1994), among others. However, in general this condition does not hold for the mapping Ψ_i in the class of dynamic games that we consider in this paper.

There are several reasons why this estimator is of interest. First, it solves the problem of indeterminacy associated with multiple equilibria. Second, repeated solutions of the dynamic game are avoided and this can result in significant computational gains. Furthermore, for some models the pseudo likelihood function $\tilde{l}_T(P^0,\theta)$ is globally concave in θ, which in general it is not the case for the likelihood function. That is the case in the conditional logit game presented in Section 2.5. This can also reduce significantly the computational cost in the estimation.

Example: Consider the conditional logit model in section 2.5. In this model, the probabilities $\Psi_i(a|x;\hat{P}^0,\theta)$ that enter in the pseudo-likelihood function have the following logistic form:

$$\Psi_i(a_i|x_m;\hat{P}^0,\theta) = \frac{\exp\{\ \tilde{y}^0_{1im}(a_i)\ \theta_1 + \tilde{y}^0_{2im}(a_i)\ \theta_2 + \tilde{s}^0_{im}(a_i)\ \theta_3 + \tilde{e}^0_{im}(a_i)\ \}}{\sum_{j=1}^{J}\exp\{\ \tilde{y}^0_{1im}(j)\ \theta_1 + \tilde{y}^0_{2im}(j)\ \theta_2 + \tilde{s}^0_{im}(j)\ \theta_3 + \tilde{e}^0_{im}(j)\}}$$

where, given $\hat{P}^0$, the values $\tilde{y}^0_{1im}$, $\tilde{y}^0_{2im}$, $\tilde{s}^0_{im}$, and $\tilde{e}^0_{im}$ can be obtained by solving a system of linear equations with dimension $|X|(J-1)$ (see the Appendix). Therefore, the pseudo likelihood function is just the likelihood of a standard conditional logit model. It is simple to incorporate unobservable common knowledge state variables in this model. For instance, consider that θ_1, θ_2 and θ_3 are vectors of random coefficients which depend on ξ, such that $\theta_{1m} = \bar{\theta}_1 + \sigma_1\xi_m$, where $\bar{\theta}_1$ and σ_1 are vectors of parameters. We can estimate this random coefficients logit model by simulated maximum likelihood as in McFadden and Train (2000).

A potentially important limitation of this PMLE is that initial nonparametric estimates of P^0 can be very imprecise, and this will imply large asymptotic variances and significant finite sample biases in the estimates of structural parameters. We propose two ways of dealing with this problem. First, it is possible to exploit the parametric assumption about the distribution of $\varepsilon' s$ to obtain initial semiparametric estimates of P^0. These estimates will be more precise than nonparametric estimates. Second, we can apply our procedure recursively to obtain a sequence of consistent estimators that may have better asymptotic and finite sample properties than the initial estimator. Let $\hat{\theta}^1_{PMLE}$ be our initial estimator. Associated with this estimator we can obtain consistent estimates of choice probabilities that exploit the structure of the model. Let $\hat{P}^1$ be this vector of estimated probabilities, where $\hat{P}^1_i(a_i|x) = \Psi_i(a_i|x;\hat{P}^0,\hat{\theta}^1_{PMLE})$. Then, we can obtain a new pseudo-maximum likelihood estimator as $\hat{\theta}^2_{PMLE} = \arg\max_{\theta\in\Theta}\tilde{l}_T(\theta,\hat{P}^1)$. In general, we can obtain a K-stage PML estimator as:

$$\hat{\theta}^K_{PMLE} = \arg\max_{\theta\in\Theta}\ \tilde{l}_T(\hat{P}^{K-1},\theta) \qquad (27)$$

where, for $K \geq 1$, $\hat{P}_i^K(a_i|x) = \Psi_i(a_i|x; \hat{P}^{K-1}, \hat{\theta}_{PMLE}^K)$.

5 Complexity and a randomized PML estimator

5.1 Degree of complexity of the PML estimator

The dimension or size of a dynamic discrete game can be described in terms of three integers: the number of players, N, the number of actions per player J, and the number of possible states, $|X|$. Notice that $x = \{x_1, x_2, ..., x_N\}$ where x_i is a player-specific state variable. Therefore, the dimension of the state space depends on the number of players, and we can write $|X| = S^N$, where S is the dimension of the space where x_i lives. The size of the estimation problem depends also on the number of observations, T.

First, we measure the computational cost of implementing the PML estimator in a problem of size $\{N, J, S, T\}$. Following the literature of complexity theory, the *worst case complexity* of a problem of size n, denoted by $comp(n)$, is the minimal number of simple algebraic operations (i.e., sums or multiplications of two numbers) necessary to solved the hardest possible problem with that size. This measure of complexity is generally represented by asymptotic lower or upper bounds. For instance, $comp(n) = O(h(n))$ means that the function $h(n)$ is an asymptotic upper bound for the complexity of a problem with size n, that is, $\lim_{n \to \infty} |comp(n)/h(n)| < \infty$.

Table 1 presents the degree of complexity of different computational tasks which should be performed in the implementation of the PML estimator. For all these tasks, complexity is exponential in the number of players. Therefore, there is a "curse of dimensionality" in the estimation of discrete games with a relatively large number of players. Notice that this problem exists even when the game is static because for static models we should also compute expected payoffs π_i, and this computation is subject to a curse of dimensionality.

5.2 Randomized PML estimator

We present here a version of the PML estimator that exploits Monte Carlo techniques to approximate expected profits and the valuation operators $\Gamma_i(P)$. This estimator exploits the randomization method proposed by Rust (1997) in the context of solving single agent dynamic programming models.[12] We show that this estimator can be calculated in polynomial time, and that it is asymptotically equivalent to the PMLE when the number of Monte Carlo simulations and the sample size go to infinity. In this sense, this estimator breaks the curse of dimensionality in the estimation of this class of dynamic discrete games.

First, we follow Rust (1997) to define a randomized version of our equilibrium mapping Ψ. Let $\tilde{u} = \{u_1, u_2, ...\}$ be an infinite sequence of independent random draws from a $U(0,1)$. This sequence is fixed and it does not depend on any of the primitives of the model or on the sample. Let $\tilde{u}_R$ denote the first R elements of the sequence $\tilde{u}$. Given transition probabilities $f = \{f(x'|x,a)\}$, choice probabilities $P = \{P_i(a_i|x)\}$ and $\tilde{u}_R$, we generate R independent random draws from the unconditional (ergodic) distribution of the state variables x.[13] Let $\tilde{X}_R = \{\tilde{x}_r : r = 1, 2, ..., R\}$ be this set of random draws. For every $\tilde{x} \in \tilde{X}_R$ we obtain R independent random draws of players' decisions using the conditional choice probabilities $\{P_i(.|\tilde{x})\}$ and $\tilde{u}_R$.[14] Let $\{\tilde{a}_r(\tilde{x}) : r = 1, 2, ..., R\}$ be this set of simulations, where $\tilde{a}_r(\tilde{x}) = \{\tilde{a}_{1r}(\tilde{x}), \tilde{a}_{2r}(\tilde{x}), ..., \tilde{a}_{Nr}(\tilde{x})\}$. The sequence $\tilde{u}_R$ is an argument in all the simulators that we define below. However, for notational simplicity, we omit $\tilde{u}_R$ as argument and instead we use a subindex R: i.e., for any function $h(.)$, $h_R(.) \equiv h(.; \tilde{u}_R)$.

Given these random draws, we define the following simulators of expected profits and transition probabilities. For every $(a_i, \tilde{x}) \in A \times X_R$:

$$\tilde{\pi}_{i,R}(a_i, \tilde{x}; P, \theta) = \frac{1}{R} \sum_{r=1}^{R} \Pi_i(a_i, \tilde{a}_{-ir}(\tilde{x}), \tilde{x}; \theta) \tag{28}$$

[12] Rust shows that his randomization tehnique solves the curse of dimensionality in the solution of discrete choice dynamic programming problems with continuous state variables. Here the problem is slightly different because our state space is discrete, though there is a curse of dimensionality associated with the number of players. Also, our main concern here is estimation.

[13] In principle, there is a curse of dimensionality in the computation of probabilities in f and P, and in the calculation of the unconditional distribution of x given f and P. However, we show below that for the implementation of our estimator we need estimates of these probabilities at only R values of x, where R does not increase exponentially with the size of the problem.

[14] By the conditional independence of players' actions, the computation of these draws is not subject to a curse of dimensionality.

And for every $(a_i, \tilde{x}, \tilde{x}') \in A \times X_R \times X_R$:

$$\hat{f}_{i,R}(\tilde{x}'|\tilde{x}, a_i; P) = \frac{1}{R}\sum_{r=1}^{R} f(\tilde{x}'|\tilde{x}, a_i, \tilde{a}_{-ir}(\tilde{x})) \tag{29}$$

Since $\sum_{\tilde{x}' \in X_R} \hat{f}_{i,R}(\tilde{x}'|\tilde{x}, a_i; P)$ does not necessarily sum to one, we consider the following simple normalization (Rust, 1997):

$$\tilde{f}_{i,R}(\tilde{x}'|\tilde{x}, a_i; P) = \frac{\hat{f}_{i,R}(\tilde{x}'|\tilde{x}, a_i; P)}{\sum_{r=1}^{R} \hat{f}_{i,R}(\tilde{x}_r|\tilde{x}, a_i; P)} \tag{30}$$

Now, we use the previous simulators to define randomized versions of the mappings $\Gamma_i(P;\theta)$ and $\Psi_i(P,\theta)$. The randomized operator (or simulator) $\tilde{\Gamma}_{i,R}(P,\theta)$ is the $R \times 1$ vector of values $\tilde{V}_i$ that solves the system of linear equations:

$$\left[I_R - \beta \ \tilde{F}_R(P)\right] \tilde{V}_i = \sum_{a_i \in A} P_i(a_i) * \left[\tilde{\pi}_{i,R}(a_i; P, \theta) + e_i(a_i; P_i)\right] \tag{31}$$

where $\tilde{F}_R(P)$ is the $R \times R$ matrix of transition probabilities that we obtain using the simulators $\tilde{f}_{i,R}(\tilde{x}'|\tilde{x}, a_i; P)$. Finally, the simulator of the equilibrium mapping Ψ_i is:

$$\tilde{\Psi}_{i,R}(a_i|x; P, \theta) = \Lambda_i(a_i \mid \{\tilde{v}_{i,R}(j, x; P, \theta)\}) \tag{32}$$

where $\{\tilde{v}_{i,R}(j, x; P, \theta)\}$ are simulators of conditional choice value functions,

$$\tilde{v}_{i,R}(j, x; P, \theta) \equiv \tilde{\pi}_{i,R}(j, x; P, \theta) + \beta \sum_{r=1}^{R} \tilde{\Gamma}_{i,R}(\tilde{x}_r; P, \theta) \ \tilde{f}_{i,R}(\tilde{x}_r|x, a_i; P) \tag{33}$$

It is clear that, by a strong law of large numbers, all these simulators are unbiased as R goes to infinity. They are also continuously differentiable in θ and P. Furthermore, given that the random draws $\tilde{u}_R$ do not change with θ and P, these simulators are stochastically equicontinuous in θ. The degree of complexity of calculating simulated expected profits and transition probabilities is $O(NJR^2)$ and $O(NJR^3)$, respectively. And the computational cost of solving the system of equations that defines $\tilde{\Gamma}_{i,R}$ has order of magnitude $O(R^3)$.

Let $\hat{P}^0$ be a consistent estimator of P^0. The simulated PML estimator of θ^0 is the value of θ that maximizes in Θ the simulated pseudo likelihood function,

$$Q_{T,R}(\hat{P}^0, \theta) = \frac{1}{T}\sum_{m=1}^{T}\sum_{i=1}^{N} \ln \tilde{\Psi}_{i,R}(a_{im}|x_m; \hat{P}^0, \theta) \tag{34}$$

Notice that to construct this simulated likelihood we need to evaluate the mapping $\tilde{\Psi}_i$ at the sample points $\{x_m : m = 1, 2, ..., T\}$. Therefore, these sample values should belong to X_R, and R should be greater or equal than T.[15] The remaining $R-T$ random draws are generated using the transition probability $f(.|x, a)$ for values (x, a) taken randomly from the sample. Also, notice that to obtain all the previous simulators we do not have to estimate choice probabilities at every point in the space X. We only need to estimate these probabilities at the R values in X_R. Therefore, the degree of complexity of evaluating the simulated pseudo likelihood is $O(R^3)$.

The following proposition establishes the asymptotic properties of this estimator.

PROPOSITION 3: Under the conditions in Proposition 2 the simulated pseudo-maximum likelihood estimator (S-PMLE) is asymptotically equivalent to the PMLE if $\sqrt{T}/R$ goes to zero as T goes to infinity.

6 Conclusions

In this paper we have presented several results and techniques to deal with some common problems which appear in the estimation of dynamic discrete games.

First, we showed that, if players do not randomize among multiple equilibria, the econometric identification of payoff functions does not depend on whether the model has multiple or unique equilibria. Although it is well known that there is not a necessary relationship between multiple equilibria and under-identification, non-uniqueness makes a model more susceptible of being under-identified (Jovanovic, 1989). Therefore, finding conditions for identification is particularly relevant in models with multiple equilibria.

Second, we illustrated the analogy between the problem of identifying strategic interactions in our model and Manski's reflection problem. We presented exclusion restrictions that identify strategic interaction. These exclusion restrictions are based on two economic assumptions: (1) players face adjustment costs when changing their decisions; and (2) strategic interactions occur contemporaneously and not with a lag.

[15] Since we have a random sample, we can always consider sample values as random draws from the unconditional distribution of x.

We think that these are plausible assumptions in some applications.

Third, we propose a pseudo maximum likelihood estimator (PMLE) and obtained its asymptotic properties. This method has several attractive features: it solves the problem of indeterminacy associated with multiple equilibria, and it is computationally cheaper than solving the model just once. In Aguirregabiria and Mira (2002) we showed that in the context of single agent dynamic discrete models this PMLE is asymptotically efficient. However, we show here that this property does not hold in the case of static or dynamic discrete games. For that reason, the sequential K-stage version of this estimator might be particularly useful for this type of models.

The cost of implementing the PMLE increases exponentially with the number of players. This can make this method unfeasible even when the number of players is relatively small, like 20 or 30 players. For that reason, we propose a simulated version of this estimator that exploits randomization techniques proposed by Rust (1997). We show that this simulated PMLE is asymptotically equivalent to the PMLE as long as the number of Monte Carlo simulations grow faster than the sample size. Preliminary experiments with this estimator (not reported here) show encouraging results.

APPENDIX

[A.1] MAPPING Ψ_i IN A CONDITIONAL LOGIT GAME

Given the Extreme value distribution of ε_i, the mapping Ψ_i has the following form:

$$\Psi_i(a_i|x;P) = \frac{\exp\left\{\pi_i(a_i,x;P_{-i}) + \beta\sum_{x'}\Gamma_i(x';P_{-i})\ \bar{f}_i(x'|x,a_i;P_{-i})\right\}}{\sum_{j=1}^{J}\exp\left\{\pi_i(j,x;P_{-i}) + \beta\sum_{x'}\Gamma_i(x';P_{-i})\ \bar{f}_i(x'|x,j;P_{-i})\right\}} \quad \text{(A.1.1)}$$

In general, the mapping $\Gamma_i(P)$ has the following vector form:

$$\Gamma_i(P) = [I - \beta\ F(P)]^{-1}\sum_{a_i\in A} P_i(a_i) * [\pi_i(a_i;P_{-i}) + e_i(a_i;P_i)] \quad \text{(A.1.2)}$$

where $F(P)$ is the $|X|\times|X|$ matrix of transitions $\tilde{f}(x'|x;P)$. Given our specification of payoff function in the conditional logit model, we have that the $|X|\times 1$ vectors of expected payoffs $\pi_i(a_i;P_{-i})$ can be written as:

$$\pi_i(a_i;P_{-i}) = Y_1(a_i)\theta_1 + Y_2(a_i)\theta_2 + S(a_i,P_{-i})\theta_3$$

where $Y_j(a_i)$ is a $|X|\times K_j$ matrix with rows $y_j(a,x)$; and $S_i(P)$ is a $|X|\times K_3$ matrix with rows $s_i(a,P_{-i}(x))$. Also, given the extreme value distribution of ε:

$$e_i(a_i;P_i) = \gamma - \ln(P_i(a_i)),$$

where γ is Euler's constant. Therefore, we can write,

$$\Gamma_i(P) = W_i^1(P)\ \theta_1 + W_i^2(P)\ \theta_2 + W_i^S(P)\ \theta_3 + W_i^e(P) \quad \text{(A.1.3)}$$

where:

$$\begin{aligned} W_i^j(P) &= [I-\beta\ F(P)]^{-1}\textstyle\sum_{a_i\in A} P_i(a_i) * Y_j(a_i) \\ W_i^S(P) &= [I-\beta\ F(P)]^{-1}\textstyle\sum_{a_i\in A} P_i(a_i) * S(a_i,P_{-i}) \\ W_i^e(P) &= [I-\beta\ F(P)]^{-1}\textstyle\sum_{a_i\in A} P_i(a_i) * e_i(a_i;P_i) \end{aligned}$$

The matrices W have the following interpretation. The $m-th$ row of $W_i^1(P)$ is the expected and discounted value of current and future realizations of y_1 for player i given that the current state is x^m and that agents behave, now and in the future, according to choice probabilities in P. The other matrices have the same interpretation for the other components of the payoff function.

Solving the previous expression of $\Gamma_i(P)$ in the mapping Ψ_i we get:

$$\Psi_i(a|x;P) = \frac{\exp\{\ \tilde{y}_{1i}(a,x;P)\ \theta_1 + \tilde{y}_{2i}(a,x;P)\ \theta_2 + \tilde{s}_i(a,x;P)\ \theta_3 + \tilde{e}_i(a,x;P)\ \}}{\sum_{j=1}^{J}\exp\{\ \tilde{y}_{1i}(j,x;P)\ \theta_1 + \tilde{y}_{2i}(j,x;P)\ \theta_2 + \tilde{s}_i(j,x;P)\ \theta_3 + \tilde{e}_i(j,x;P)\}} \tag{A.1.4}$$

where:

$$\begin{aligned}
\tilde{y}_{ji}(a,x;P) &= y_j(a,x_i) + \beta \textstyle\sum_{x'} \bar{f}_i(x'|x,a;P_{-i})\ W_i^j(x';P) \\
\tilde{s}_i(a;P) &= s_i(a,x;P_{-i}(x)) + \beta \textstyle\sum_{x'} \bar{f}_i(x'|x,a;P_{-i})\ W_i^S(x';P) \\
\tilde{e}_i(a,x;P) &= \beta \textstyle\sum_{x'} \bar{f}_i(x'|x,a;P_{-i})\ W_i^e(x';P)
\end{aligned}$$

[A.2] PROOF OF PROPOSITION 1

Let P^0 be the true population choice probabilities. The model imposes two sets of restrictions on P^0: (1) choice probabilities are best responses to one another: $P_i^0(a_i|x) = \Phi_i(a_i|x;P_{-i}^0)$; and (2) the Bellman equation implies a relationship between value functions and choice probabilities: $V_i(x;P_{-i}^0) = \Gamma_i(x;P^0)$. Notice that the mapping $\Psi_i(a_i|x;P)$ is just the mapping $\Phi_i(a_i|x;P_{-i})$ where we have solved the condition $V_i(x;P_{-i}) = \Gamma_i(x;P)$. Therefore, the two sets of conditions can be summarized in just one: $P_i^0(a_i|x) = \Psi_i(a_i|x;P^0)$. We prove here that this set of equations identifies expected profits $\pi_i(a_i,x;P_{-i}^0)$.

For notational simplicity, we use the superindex 0 in those functions which depend on P^0 and we omit this vector as an explicit argument, e.g., $\pi_i^0(a_i,x) = \pi_i(a_i,x;P_{-i}^0)$. Define the conditional choice value functions:

$$v_i^0(a_i,x) \equiv \pi_i^0(a_i,x) + \beta \sum_{x'\in X} \Gamma_i^0(x')\ \bar{f}_i^0(x'|x,a_i) \tag{A.2.1}$$

Notice that by the equilibrium condition and by the definition of mapping Ψ_i we can write:

$$P_i^0(a_i|x) = \Psi_i^0(a_i|x) \equiv \int I\left\{a_i = \arg\max_{j\in A}\left[v_i^0(j,x) + \varepsilon_i(j)\right]\right\}\ g_i(d\varepsilon_i) \tag{A.2.2}$$

Under Assumption 1 and 2, the system of equations on (A.2.2) implies a one-to-one relationship between P_i^0 and the set of value differences $\{v_i^0(a,x) - v_i^0(J,x)\}$ (see Proposition 1 in Hotz and Miller, 1993). Let Q_i be this mapping, such that $v_i^0(a,x) - v_i^0(J,x) = Q_i(a,x;P_i^0)$. An important property of the mapping Q_i is that it only depends on the distribution of the unobservables ε_i. In particular, it does not depend on the other primitives of the model, i.e., discount factor, profits or transition probabilities of common knowledge variables. Therefore, we can treat $Q_i(a, x; P_i^0)$ as known values.

Taking into account the definition of conditional choice value function we can write:

$$\pi_i^0(a_i, x) + \beta \sum_{x' \in X} \left[\bar{f}_i(x'|x, a_i) - \bar{f}_i(x'|x, J) \right] \Gamma_i^0(x') = Q_i(a, x; P_i^0) \tag{A.2.3}$$

where we have used that $\pi_i^0(J, x) = 0$ by our normalization assumption. Notice that $\Gamma_i^0(x)$ depends on expected profits, and therefore it is not obvious that these profits are identified from the previous system of equations. In particular, there might be more than one set of expected profit functions that solve this system of equations.

Writing the system of equations (A.2.3) in vector form, and taking into account that $\Gamma_i^0 = [I - \beta F^0]^{-1} \left[\sum_{j=1}^{J} P_i^0(j) * (\pi_i^0(j) + e_i^0(j)) \right]$, we can write:

$$\pi_i^0(a_i) + \beta \left[F_i^0(a_i) - F_i^0(J) \right] \left[I - \beta F^0 \right]^{-1} \left[\sum_{j=1}^{J} P_i^0(j) * \left(\pi_i^0(j) + e_i^0(j) \right) \right] = Q_i(a_i; P_i^0), \tag{A.2.4}$$

where $\pi_i^0(a)$, Γ_i^0 and $Q_i(a_i; P_i^0)$ are $|X| \times 1$ vectors, and $F_i^0(a)$ are $|X| \times |X|$ matrices. Rearranging terms we have that:

$$B_i(a, 1)\ \pi_i^0(1) + B_i(a, 2)\ \pi_i^0(2) + ... + B_i(a, J-1)\ \pi_i^0(J-1) = C_i(a) \tag{A.2.5}$$

where $C_i(a)$ is a $|X| \times 1$ vector, and $\{B_i(a, j)\}$ are $|X| \times |X|$ matrices with the following form:

$$C_i(a) = Q_i(a_i; P_i^0) - \beta \left[F_i^0(a_i) - F_i^0(J) \right] \left[I - \beta\ F^0 \right]^{-1} \left[\sum_{j=1}^{J} P_i^0(j) * e_i^0(j) \right] \tag{A.2.6}$$

and,

$$B_i(a, j) = \begin{cases} I + \beta \left[F_i^0(a) - F_i^0(J) \right] \left[I_{|X|} - \beta F^0 \right]^{-1} diag\{P_i^0(a)\} & if \quad a = j \\ \beta \left[F_i^0(a) - F_i^0(J) \right] \left[I_{|X|} - \beta F^0 \right]^{-1} diag\{P_i^0(j)\} & if \quad a \neq j \end{cases} \tag{A.2.7}$$

Both the matrices $\{B_i(a, j)\}$ and the vectors $C_i(a)$ only depend on $\{f, \beta, g_i, P_i^0 : i \in I\}$, i.e., they do not depend on profit functions.

Since we have one of these expressions for each value $a \in A_{-J}$, we have a linear system with $|X|(J-1)$ equations and unknowns, where the unknowns are the one-period expected profits. We can write this system for firm i as: $B_i\ \pi_i^0 = C_i$, where $\pi_i^0 \equiv \{\pi_i^0(1)', \pi_i^0(2)', ..., \pi_i^0(J-1)'\}'$, and B_i is the $|X|(J-1) \times |X|(J-1)$ matrix:

$$B_i = I_{|X|(J-1)} + \beta \left(\tilde{1}_{J-1} \otimes \begin{pmatrix} [F_i^0(1) - F_i^0(J)]\,[I - \beta F^0]^{-1} \\ \vdots \\ [F_i^0(J-1) - F_i^0(J)]\,[I - \beta F^0]^{-1} \end{pmatrix} \right) diag\{P_i^0\} \tag{A.2.8}$$

where $\tilde{1}_{J-1}$ is a $1 \times (J-1)$ vector of ones. It is simple to show that this matrix is invertible. Therefore, there is a unique vector of expected profits π_i^0 which is consistent with P^0, f, g_i and β.

[A.3] PROOF OF LEMMA 1:

We consider here the *worst case scenario* in which there are not observable variables z. Let $\tilde{P}^0 = \{\Pr(a_m = a) : a \in A^N\}$ be the distribution of a_m, which is identified from the data. Let $\lambda \equiv \{\lambda(\xi) : \xi \in \Upsilon\}$ be the probability distribution of ξ, and let $P^0 \equiv \{P_i^0(a_i|\xi) : a_i \in A,\ \xi \in \Upsilon,\ i \in I\}$ represent player's choice probabilities. We want to identify λ and P^0 from $\tilde{P}^0$. Independence of private information variables implies that for any $a \in A^N$: $\tilde{P}^0(a) = \sum_{\xi \in \Upsilon} \lambda(\xi) \left[\prod_{i=1}^N P_i^0(a_i|\xi)\right]$. Or in vector form, $\tilde{P}^0 = H(\lambda, P^0)$. The order condition for identification is $\dim(\tilde{P}^0) \geq \dim(\lambda) + \dim(P^0)$, which implies: $J^N - 1 \geq (L-1) + N(J-1)L$, or $L \leq int\left(J^N / [N(J-1)+1]\right)$. The rank condition for identification requires the Jacobian matrix of $H(\lambda, P^0)$ to be full-column rank. Condition (6) in Lemma 1 and $\lambda(\xi) > 0$ for any $\xi \in \Upsilon$ guarantee that this rank condition holds.

[A.4] PROOF OF PROPOSITION 2

Consistency of PMLE: Notice that: (a) $\tilde{l}_\infty(P, \theta)$ is continuous on a compact set, so it is uniformly continuous; (b) $\tilde{l}_T(P, \theta)$ converges a.s. and uniformly in (P, θ) to $\tilde{l}_\infty(P, \theta)$; and (c) $\hat{P}_T^0$ converges a.s. to P^0. Under (a)-(c), $\tilde{l}_T(\hat{P}_T^0, \theta)$ converges a.s. and uniformly in θ to $\tilde{l}_\infty(P^0, \theta)$ (Lemma 24.1 in Gourieroux and Monfort, vol. II, page 392). Then, given that $\tilde{l}_\infty(P^0, \theta)$ has a unique maximum in Θ at θ^0, $\hat{\theta}_T \equiv \arg\max_{\theta \in \Theta} \tilde{l}_T(\hat{P}_T^0, \theta)$ converges a.s. to θ^0 (Property 24.2 in Gourieroux and Monfort, vol. II, page 387).

Asymptotic distribution of PMLE: First order conditions of optimality imply that with probability approaching one $\partial\tilde{l}_T(\hat{P}^0_T,\hat{\theta}_T)/\partial\theta = 0$. Since $\tilde{l}_T(P,\theta)$ is twice continuously differentiable, we can apply the stochastic mean value theorem to $\partial\tilde{l}_T(.,.)/\partial\theta$ between $(\hat{P}^0_T,\hat{\theta}_T)$ and (P^0,θ^0). Let $(\bar{P}_T,\bar{\theta}_T)$ be the vector of mean values which. By consistency of $(\hat{P}^0_T,\hat{\theta}_T)$, these mean values also converge a.s. to (P^0,θ^0). By the mean value theorem,

$$0 = \frac{\partial\tilde{l}_T(\hat{P}^0_T,\hat{\theta}_T)}{\partial\theta} = \frac{\partial\tilde{l}_T(P^0,\theta^0)}{\partial\theta} + \frac{\partial\tilde{l}_T(\bar{P}_T,\bar{\theta}_T)}{\partial\theta\partial P'}(\hat{P}^0_T - P^0) + \frac{\partial\tilde{l}_T(\bar{P}_T,\bar{\theta}_T)}{\partial\theta\partial\theta'}(\hat{\theta}_T - \theta^0) \quad \text{(A.4.1)}$$

Rearranging terms,

$$\sqrt{T}(\hat{\theta}_T - \theta^0) = -\left[\frac{\partial\tilde{l}_T(\bar{P}_T,\bar{\theta}_T)}{\partial\theta\partial\theta'}\right]^{-1}\left\{\frac{\partial\tilde{l}_T(\bar{P}_T,\bar{\theta}_T)}{\partial\theta\partial P'}\sqrt{T}(\hat{P}^0_T - P^0) + \sqrt{T}\frac{\partial\tilde{l}_T(P^0,\theta^0)}{\partial\theta}\right\} \quad \text{(A.4.2)}$$

where $\partial^2\tilde{l}_T(\bar{P}_T,\bar{\theta}_T)/\partial(P,\theta)(P',\theta') \rightarrow_p \partial^2\tilde{l}_\infty(P^0,\theta^0)/\partial(P,\theta)(P',\theta')$ (see Amemiya, Theorems 4.2.1 and 4.1.5), and $\left[\sqrt{T}(\hat{P}^0_T - P^0)' \ ; \ \sqrt{T}\ \partial\tilde{l}_T(\theta^0,P^0)/\partial\theta'\right] \rightarrow_d N(0,\Omega^0)$. Therefore, by Mann-Wald Theorem, it is straightforward to show that $\sqrt{T}(\hat{\theta}_T - \theta^0) \rightarrow_d N(0,V^0)$, where:

$$V^0 = (\Omega^0_{\theta\theta})^{-1} + (\Omega^0_{\theta\theta})^{-1}\left[\Delta^0_{\theta P}\ \Omega^0_{PP}\ \Delta^{0\prime}_{\theta P} + \Omega^0_{\theta P}\ \Delta^{0\prime}_{\theta P} + \Delta^0_{\theta P}\ \Omega^0_{P\theta}\right](\Omega^0_{\theta\theta})^{-1}$$

and $\Delta^0_{\theta P} \equiv \partial^2\tilde{l}_\infty(P^0,\theta^0)/\partial P\partial\theta'$.

[A.5] PROOF OF PROPOSITION 3

Consistency of S-PMLE (as T $\rightarrow\infty$ and R $\rightarrow\infty$): First, by a strong law of large numbers and for arbitrary (P,θ), $Q_{T,R}(P,\theta)$ converges to $Q_{\infty,\infty}(P,\theta) = \tilde{l}_\infty(P,\theta)$ as $T\rightarrow\infty$ and $R\rightarrow\infty$. Furthermore, since $Q_{T,R}(P,\theta)$ is continuously differentiable in (P,θ) and bounded, it converges a.s. and uniformly in (P,θ) to $\tilde{l}_\infty(P,\theta)$. Therefore, given that $\tilde{l}_\infty(P,\theta)$ is uniformly continuous and that $\hat{P}^0_T$ converges a.s. to P^0, we have that $Q_{T,R}(\hat{P}^0_T,\theta)$ converges a.s. and uniformly in θ to $\tilde{l}_\infty(P^0,\theta)$. Finally, given that $\tilde{l}_\infty(P^0,\theta)$ has a unique maximum in Θ at θ^0, $\hat{\theta}_{T,R} \equiv \arg\max_{\theta\in\Theta} Q_{T,R}(\hat{P}^0_T,\theta)$ converges a.s. to θ^0.

Asymptotic distribution of S-PMLE (as T $\rightarrow\infty$ and $\sqrt{T}$/R $\rightarrow$ 0): First order conditions of optimality imply that with probability approaching one $\partial Q_{T,R}(\hat{P}^0_T,\theta)/\partial\theta =$

0. Since $Q_{T,R}(P,\theta)$ is twice continuously differentiable, we can apply the stochastic mean value theorem to $\partial Q_{T,R}(P,\theta)/\partial\theta$ between $(\hat{P}_T^0, \hat{\theta}_{T,R})$ and (P^0, θ^0). Let $(\bar{P}_T, \bar{\theta}_{T,R})$ be the vector of mean values. By consistency of $(\hat{P}_T^0, \hat{\theta}_T)$, these mean values also converge a.s. to (P^0, θ^0) as $T \to \infty$ and $R \to \infty$. By the mean value theorem,

$$0 = \frac{\partial Q_{T,R}(\hat{P}_T^0, \hat{\theta}_{T,R})}{\partial\theta} = \frac{\partial Q_{T,R}(P^0,\theta^0)}{\partial\theta} + \frac{\partial Q_{T,R}(\bar{P}_T, \bar{\theta}_{T,R})}{\partial\theta\partial P'}(\hat{P}_T^0 - P^0) + \frac{\partial Q_{T,R}(\bar{P}_T, \bar{\theta}_{T,R})}{\partial\theta\partial\theta'}(\hat{\theta}_T - \theta^0) \tag{A.5.1}$$

Rearranging terms,

$$\sqrt{T}(\hat{\theta}_{T,R} - \theta^0) = -\left[\frac{\partial Q_{T,R}(\bar{P}_T, \bar{\theta}_{T,R})}{\partial\theta\partial\theta'}\right]^{-1}\left\{\frac{\partial Q_{T,R}(\bar{P}_T, \bar{\theta}_{T,R})}{\partial\theta\partial P'}\sqrt{T}(\hat{P}_T^0 - P^0) + \sqrt{T}\frac{\partial Q_{T,R}(P^0,\theta^0)}{\partial\theta}\right\} \tag{A.5.2}$$

where $\partial^2 Q_{T,R}(\bar{P}_T, \bar{\theta}_{T,R})/\partial(P,\theta)(P',\theta') \to_p \partial^2\tilde{l}_\infty(P^0,\theta^0)/\partial(P,\theta)(P',\theta')$, as $T \to \infty$ and $R \to \infty$. Therefore, as both T and R go to infinity, the only difference between the asymptotic distribution of the S-PMLE and the one of the PMLE is in the last term of equation (A.5.2). We can write this term as:

$$\begin{aligned}
\sqrt{T}\frac{\partial Q_{T,R}(P^0,\theta^0)}{\partial\theta} &= \sqrt{T}\frac{\partial Q_{T,\infty}(P^0,\theta^0)}{\partial\theta} + \left[\sqrt{T}\frac{\partial Q_{T,R}(P^0,\theta^0)}{\partial\theta} - \sqrt{T}\frac{\partial Q_{T,\infty}(P^0,\theta^0)}{\partial\theta}\right] \\
&= \sqrt{T}\frac{\partial \tilde{l}_T(P^0,\theta^0)}{\partial\theta} + \left[\sqrt{T}\frac{\partial Q_{T,R}(P^0,\theta^0)}{\partial\theta} - \sqrt{T}\frac{\partial \tilde{l}_T(P^0,\theta^0)}{\partial\theta}\right]
\end{aligned}$$

The term between brackets converges in probability to zero as $T \to \infty$ and $\sqrt{T}/R \to 0$ (see for instance Gourieroux and Monfort, 1993). Therefore, it is straightforward that as $T \to \infty$ and $\sqrt{T}/R \to 0$ the asymptotic distributions of S-PMLE and PMLE are the same.

REFERENCES:

Aguirregabiria and Mira (2002): "Swapping the nested fixed point algorithm: A class of estimators for Markov decision models", *Econometrica*, forthcoming.

Ahn and Manski (1993): "Distribution theory for the analysis of binary choice under uncertainty with non-parametric estimation of expectations", *Journal of Econometrics* 56, 291-321

Berry, S. (1992): "Estimation of a model of entry in the airline industry", *Econometrica,* 60(4), 889-917.

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

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium", *Econometrica*, 63, 841-890.

Blundell, R., and R. Smith (1994): "Simultaneous microeconometric models with censored or qualitative dependent variables," in G.S. Maddala, C. R. Rao, and H. D. Vinod (eds.) *Handbook of Statistics,* Vol. 11, 117-143. Elsevier Science Publishers.

Bresnahan, T., and P. Reiss (1990): "Entry in monopoly markets", *Review of Economic Studies*, 57, 531-553.

Bresnahan, T., and P. Reiss (1991): "Empirical models of discrete games", *Journal of Econometrics* 48, 57-81.

Brock, W., and S. Durlauf (2001): "Discrete choice with social interactions", *Review of Economics Studies*, 68, 235-260.

Davis, P. (2000): "Spatial competition in retail markets: Movie theters," mimeo, Sloan School of Business, MIT, Cambridge, MA.

Ericson, R., and A. Pakes (1995): "Markov-Perfect industry dynamics: A framework for empirical work", *Review of Economic Studies,* 62, 53-82.

Gourieroux, C., and A. Monfort: (1993) "Simulation-based inference: A survey with special reference to panel data models", *Journal of Econometrics*, 59, 5-33.

Gourieroux, C., and A. Monfort (1995): "Statistics and Econometric Models: Volume I", Cambridge University Press.

Harsanyi, J. (1995): "Games with incomplete information," *American Economic Review*, 85, 291-303.

Heckman J. (1978): "Dummy endogenous variables in a simultaneous equation system", *Econometrica,* 46, 931-959.

Hotz, J., and R.A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models", *Review of Economic Studies,* 60, 497-529.

Hotz, J., R.A. Miller, S. Sanders, and J. Smith (1994): "A simulation estimator for dynamic models of discrete choice," *Review of Economic Studies,* 61, 265-89.

Jovanovic, B. (1989): "Observable implications of models with multiple equilibria", *Econometrica*, 57, 1431-1437.

Kooreman, P. (1994): "Estimation of econometric models of some discrete games", *Journal of Applied Econometrics,* 9(3), 255-268.

Manski, C. (1993): "Identification of endogenous social effects: the reflection problem," *Review of Economic Studies*, 60, 531-542.

Manski, C. (1995): "Identification Problems in the Social Sciences", Harvard University Press.

McFadden, D. (1984): "Econometric Analysis of Qualitative Response Models" in Grilliches, Z., and M. Intriligator (eds.) Handbook of Econometrics, Vol 2.

McFadden, D., and K. Train (2000): "Mixed MNL Models for Discrete Response," *Journal of Applied Econometrics*, 15, 447-470.

McKelvey, R., and A. McLennan (1996): "Computation of Equilibria in Finite Games," in R. E. Engle and McFadden (eds.) *Handbook of Econometrics Volume 4*, North-Holland.

Milgrom, P., and R. Weber (1985): "Distributional strategies for games with incomplete information," *Mathematics and Operations Research*, 10,619-632.

Pakes, A. (1994): "Dynamic structural models, problems and prospects", in C. Sims (ed.) *Advances in Econometrics. Sixth World Congress*, Cambridge University Press.

Pakes, A. and P. McGuire (1994): "Computing Markov-Perfect Nash equilibria: Numerical implications of a dynamic differentiated product model, *Rand Journal of*

Economics, 555-589.

Pakes, A. and P. McGuire (2002): "Stochastic algorithms for dynamic models: Markov perfect equilibrium and the curse of dimensionality," *Econometrica*, forthcoming.

Pakes, A. and D. Pollard (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57, 1027-1057,

Reiss, P. (1996): "Empirical models of discrete strategic choices", *American Economic Review,* 86(2), 421-426.

Rust, J. (1994): "Estimation of dynamic structural models, problems and prospects: discrete decision processes", in C. Sims (ed.) *Advances in Econometrics. Sixth World Congress*, Cambridge University Press.

Rust, J. (1997): "Using Randomization to Break the Curse of Dimensionality," *Econometrica*, 65, 487-516.

Rust, J., J. Traub, and H. Wozniakowski (2002): "Is There a Curse of Dimensionality for Contraction Fixed Points in the Worst Case?," *Econometrica,* forthcoming.

Seim, K. (2000): "Geographic Differentiation and Firms' Entry Decisions: The Video Retail Industry", manuscript, GSB, Stanford.

Tamer, E. (2000): "Incomplete simultaneous discrete response model with multiple equilibria," mimeo, Department of Economics, Princeton University.

Toivanen, O., and M. Waterson (2000): "Empirical research on discrete choice game theory models of entry: An illustration", *European Economic Review,* 44, 985-992.

Table 1
Worst-case complexity of different computational tasks in the implementation of the PML estimator

Task	comp(N,J,S,T)
Nonparametric estimation of $\{P_i^0(a_i\|x)\}$	$O(T\ N\ J\ S^N)$
Expected payoffs: $\pi_i(a_i,x)=\sum_{a_{-i}}\left(\prod_{j\neq i}P_j(a_j\|x)\right)\Pi_i(a_i,a_{-i},x)$	$O(J^N\ S^N)$
Transition probabilities: $\bar{f}_i(x'\|x,a_i)=\sum_{a_{-i}}\left(\prod_{j\neq i}P_j(a_j\|x)\right)f(x'\|x,a_i,a_{-i})$	$O(J^N\ S^{2N})$
System of linear equations (LU decomposition)$^{(a)}$: $[I-\beta\ F(P)]\,W_i^j=Y_i^j$	$O(S^{3N}+K\ S^{2N})$
Computation of: $\hat{v}_i(a_i,x)\equiv\pi_i(a_i,x)+\beta\sum_{x'\in X}\Gamma_i(x)\,\bar{f}_i(x'\|x,j)$	$O(T\ K\ J\ S^N)$

Note (a): K is the number of structural parameters in the specification of Π_i