

SEQUENTIAL ESTIMATION OF DYNAMIC DISCRETE GAMES

BY VICTOR AGUIRREGABIRIA AND PEDRO MIRA¹

This paper studies the estimation of dynamic discrete games of incomplete information. Two main econometric issues appear in the estimation of these models: the indeterminacy problem associated with the existence of multiple equilibria and the computational burden in the solution of the game. We propose a class of pseudo maximum likelihood (PML) estimators that deals with these problems, and we study the asymptotic and finite sample properties of several estimators in this class. We first focus on *two-step PML* estimators, which, although they are attractive for their computational simplicity, have some important limitations: they are seriously biased in small samples; they require consistent nonparametric estimators of players' choice probabilities in the first step, which are not always available; and they are asymptotically inefficient. Second, we show that a recursive extension of the two-step PML, which we call *nested pseudo likelihood* (NPL), addresses those drawbacks at a relatively small additional computational cost. The NPL estimator is particularly useful in applications where consistent nonparametric estimates of choice probabilities either are not available or are very imprecise, e.g., models with permanent unobserved heterogeneity. Finally, we illustrate these methods in Monte Carlo experiments and in an empirical application to a model of firm entry and exit in oligopoly markets using Chilean data from several retail industries.

KEYWORDS: Dynamic discrete games, multiple equilibria, pseudo maximum likelihood estimation, entry and exit in oligopoly markets.

1. INTRODUCTION

EMPIRICAL DISCRETE GAMES are useful tools in the analysis of economic and social phenomena whenever strategic interactions are an important aspect of individual behavior. The range of applications includes, among others, models of market entry (Bresnahan and Reiss (1990, 1991b), Berry (1992), Toivanen and Waterson (2000)), models of spatial competition (Seim (2000)), release timing of motion pictures (Einav (2003), Zhang-Foutz and Kadiyali (2003)), intrafamily allocations (Kooreman (1994), Engers and Stern (2002)), and models with social interactions (Brock and Durlauf (2001)). Although dynamic considerations are potentially relevant in some of these studies, most applications of empirical discrete games have estimated static models. Two main econometric

¹We are thankful for comments from a co-editor, four anonymous referees, Pat Bajari, Victor Chernozukov, Phil Haile, Igal Hendel, Han Hong, Ken Judd, Mike Keane, Bob Miller, Ariel Pakes, Martin Pesendorfer, Matt Shum, Steven Stern, Eli Tamer, Ken Wolpin, and seminar participants at Amsterdam (Tinbergen), Universidad Carlos III, Carnegie–Mellon, CEMFI, Duke, Harvard, MIT, Minnesota, New York University, Princeton, Queen's, Stanford, Texas–Austin, Toronto, Toulouse, Virginia, and Wisconsin–Madison. The first author thanks the National Science Foundation for financial support (Grant SES-0241943). The second author acknowledges support from the Ministerio de Educacion y Ciencia (Grant BEC2002-02773). We thank Gustavo Vicentini for his excellent research assistance. We also thank Hernan Roman for providing the data that we used in the empirical application.

issues have limited the scope of applications to relatively simple static games: the computational burden in the solution of dynamic discrete games and the indeterminacy problem associated with the existence of multiple equilibria. This paper studies these issues in the context of a class of dynamic discrete games of incomplete information and develops techniques for the estimation of structural parameters. The rest of this introductory section discusses previous work and describes the contribution of this paper.

The existence of multiple equilibria is a prevalent feature in most empirical games where best response functions are nonlinear in other players' actions. Models with multiple equilibria do not have a unique reduced form, and this incompleteness may pose practical and theoretical problems in the estimation of structural parameters. In particular, maximum likelihood and other extremum estimators require that we obtain all the equilibria for every trial value of the parameters. This can be infeasible even for simple models. The most common approach to dealing with this problem has been to impose restrictions that guarantee equilibrium uniqueness for any possible value of the structural parameters. For instance, if strategic interactions among players have a recursive structure, the equilibrium is unique (see Heckman (1978)). A similar but less restrictive approach was used by Bresnahan and Reiss (1990, 1991a) in the context of empirical games of market entry. These authors considered a specification where a firm's profit depends on the number of firms that are operating in the market but not on the identity of these firms. Under this condition, the equilibrium number of entrants is invariant over the multiple equilibria. Based on this property, Bresnahan and Reiss proposed an estimator that maximizes a likelihood for the number of entrants. Although this can be a useful approach for some applications, it rules out interesting cases like models where firms have heterogeneous production costs or where they produce differentiated products. Notice also that these restrictions are not necessary for the identification of the model (see Tamer (2003)). (In general, a unique reduced form is neither a necessary nor a sufficient condition for identification (Jovanovic (1989)).)

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. Equilibria are fixed points of the system of best response operators, and in dynamic games, each player's best response is itself the solution to a discrete-choice dynamic programming problem. There is a "curse of dimensionality" in the sense that the cost of computing an equilibrium increases exponentially with the number of players. Furthermore, the standard nested fixed-point algorithms used to estimate single-agent dynamic models and static games require repeated solution of the model for each trial value of the vector of parameters to estimate. Therefore, the cost of estimating these models using those algorithms is much larger than the cost of solving the model just once.

This paper considers a class of pseudo maximum likelihood (PML) estimators that deals with these problems and studies the asymptotic and finite sample properties of these estimators. The method of PML was first proposed by Gong and Samaniego (1981) to deal with the problem of incidental parameters. In general, PML estimation consists of replacing all nuisance parameters in a model by estimates and solving a system of likelihood equations for the parameters of interest. This idea has been previously used in estimation of dynamic structural econometric models by Hotz and Miller (1993) and Aguirregabiria and Mira (2002). Here we show that this technique is particularly useful in the estimation of dynamic games of incomplete information with multiple equilibria and large state spaces.

Our PML estimators are based on a representation of Markov perfect equilibria as fixed points of a best response mapping in the space of players' choice probabilities. These probabilities are interpreted as players' beliefs about the behavior of their opponents. Given these beliefs, one can interpret each player's problem as a game against nature with a unique optimal decision rule in probability space, which is the player's best response. Although equilibrium probabilities are not unique functions of structural parameters, the best response mapping is always a unique function of structural parameters and players' beliefs about the behavior of other players. We use these best response functions to construct a pseudo likelihood function and obtain a PML estimator of structural parameters. If the pseudo likelihood function is based on a consistent nonparametric estimator of players' beliefs, we get a two-step PML estimator that is consistent and asymptotically normal. The main advantage of this estimator is its computational simplicity. However, it has three important limitations. First, it is asymptotically inefficient because its asymptotic variance depends on the variance of the initial nonparametric estimator. Second and more important, the nonparametric estimator can be very imprecise in the small samples available in actual applications, and this can generate serious finite sample biases in the two-step estimator of structural parameters. Third, consistent nonparametric estimators of players' choice probabilities are not always feasible for some models and data. These limitations motivate a recursive extension of the two-step PML that we call nested pseudo likelihood estimator (NPL). We show that the NPL estimator addresses these drawbacks of the two-step PML at a relatively small additional computational cost. The NPL estimator is particularly useful in applications where consistent nonparametric estimates of choice probabilities either are not available or are very imprecise, e.g., models with permanent unobserved heterogeneity. We illustrate the performance of these estimators in the context of an actual application and in Monte Carlo experiments based on a model of market entry and exit.

There has been increasing interest in estimation of discrete games during the last few years, which has generated several methodological papers on this topic. Pakes, Ostrovsky, and Berry (2004) considered a two-step method of moments

estimator in the same spirit as the two-step pseudo maximum likelihood estimator in this paper. Pesendorfer and Schmidt-Dengler (2004) defined a general class of minimum distance estimators, i.e., the *asymptotic least squares estimators*. They showed that a number of estimators of dynamic structural models (including our estimators) belong to this class and characterize the efficient estimator within this class. Bajari, Benkard, and Levin (2003) generalized the simulation-based estimator in Hotz, Miller, Sanders, and Smith (1994) to the estimation of dynamic models of imperfect competition with both discrete and continuous decision variables. For the case of static games with complete information, Tamer (2003) presented sufficient conditions for the identification of a two-player model and proposed a pseudo maximum likelihood estimation method. Ciliberto and Tamer (2006) extended this approach to static games with N players. Bajari, Hong, and Ryan (2004) also studied the identification of normal form games with complete information.

The rest of the paper is organized as follows. Section 2 presents the class of models considered in this paper and the basic assumptions. Section 3 explains the problems associated with maximum likelihood estimation, presents the two-step PML and the NPL estimators, and describes their properties. Section 4 presents several Monte Carlo experiments. Section 5 illustrates these methods with the estimation of a model of market entry–exit using actual panel data on Chilean firms. 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 similar to that in Rust (1994, pp. 154–158). To make some of the discussions less abstract, we consider a model where firms that compete in a local retail market decide the number of their outlets. A model of market entry–exit is a particular case of this framework. Although we do not deal with estimation and econometric issues until Section 3, it is useful to anticipate the type of data that we have in mind. We consider a researcher who observes many geographically separate markets such as (nonmetropolitan) small cities or towns. The game is played at the level of individual markets. The number and the identity of the players can vary across markets. Examples of applications with this type of data are Bresnahan and Reiss (1990) for car dealers, Berry (1992) for airlines, Toivanen and Waterson (2000) for fast-food restaurants, De Juan (2001) for banks, Netz and Taylor (2002) for gas stations, Seim (2000) for video rental stores, and Ellickson (2003) for supermarkets.

2.1. Framework and Basic Assumptions

Each market is characterized by demand conditions that can change over time (e.g., population, income and age distribution). Let d_t be the vector of

demand shifters at period t . There are N firms operating in the market, which we index by $i \in I = \{1, 2, \dots, N\}$. At every discrete period, t firms decide simultaneously how many outlets to operate. Profits are bounded from above such that the maximum number of outlets, J , is finite. Therefore, a firm's set of choice alternatives is $A = \{0, 1, \dots, J\}$, which is discrete and finite. We represent the decision of firm i at period t by the variable $a_{it} \in A$.

At the beginning of period t , a firm is characterized by two vectors of state variables that affect its profitability: x_{it} and ε_{it} . Variables in x_{it} are common knowledge for all firms in the market, but the vector ε_{it} is private information of firm i . For instance, some variables that could enter in x_{it} are the firm's number of outlets at the previous period or the years of experience of the firm in the market. Managerial ability at different outlets could be a component of ε_{it} . Let $x_t \equiv (d_t, x_{1t}, x_{2t}, \dots, x_{Nt})$ and $\varepsilon_t \equiv (\varepsilon_{1t}, \varepsilon_{2t}, \dots, \varepsilon_{Nt})$ be the vectors of common knowledge and private information variables, respectively. A firm's current profits depend on x_t , on its own private information ε_{it} , and on the vector of firms' current decisions $a_t \equiv (a_{1t}, a_{2t}, \dots, a_{Nt})$. 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 number of outlets to maximize expected discounted intertemporal profits,

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

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

ASSUMPTION 1—Additive Separability: *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 $\Pi_i(\cdot)$ is a real-valued function and $\varepsilon_{it}(a_{it})$ is the $(a_{it} + 1)$ th component of the $(J + 1) \times 1$ vector ε_{it} with support R^{J+1} .*

ASSUMPTION 2—Conditional Independence: *The transition probability $p(\cdot | \cdot)$ factors as $p(x_{t+1}, \varepsilon_{t+1} | a_t, x_t, \varepsilon_t) = p_\varepsilon(\varepsilon_{t+1}) f(x_{t+1} | a_t, x_t)$. That is, (i) given the firms' decisions at period t , private information variables do not affect the transition of common knowledge variables, and (ii) private information variables are independently and identically distributed over time.*

ASSUMPTION 3—Independent Private Values: *Private information is independently distributed across players, $p_\varepsilon(\varepsilon_t) = \prod_{i=1}^N g_i(\varepsilon_{it})$, where, for any player i ,*

$g_i(\cdot)$ is a density function that is absolutely continuous with respect to the Lebesgue measure.

ASSUMPTION 4—Discrete Common Knowledge Variables: *Common knowledge variables have a discrete and finite support $x_t \in X \equiv \{x^1, x^2, \dots, x^{|X|}\}$, where $|X|$ is a finite number.*

EXAMPLE 1—Entry and Exit in a Local Retail Market: Suppose that players are supermarkets making decisions on whether to open, continuing to operate, or closing their stores. The market is a small city and a supermarket has at most one store in this market, i.e., $a_{it} \in \{0, 1\}$. Every period t , all firms decide whether to operate their store(s) and then the active stores compete in quantities (Cournot competition). This competition is one period or static and it determines the variable profits of incumbent firms at period t . For instance, suppose that (i) the market inverse demand function is linear, $P_t = \alpha_0 - (\alpha_1/S_t)Q_t$, where P_t is the market price, Q_t is the aggregate quantity, S_t represents the market size at period t , and α_0 and α_1 are parameters, and (ii) all firms have the same marginal operating cost c . Under these conditions, it is simple to show that the variable profit function for a symmetric Cournot solution is $\theta_R S_t / (2 + \sum_{j \neq i} a_{jt})^2$, where θ_R is the parameter $(\alpha_0 - c)^2 / \alpha_1$. We obtain the same expression for variable profits if we assume that active firms are spatially differentiated in a circle city and compete in prices, i.e., the Salop (1979) model. In this case, the parameter θ_R is the unit transportation cost. Then current profits of an active firm are

$$(2) \quad \tilde{\Pi}_{it}(1) = \theta_R S_t / \left(2 + \sum_{j \neq i} a_{jt} \right)^2 - \theta_{FC,i} + \varepsilon_{it}(1) - (1 - a_{i,t-1})\theta_{EC},$$

where $\theta_{FC,i} - \varepsilon_{it}(1)$ is the fixed operating cost and it has two components: $\theta_{FC,i}$ is firm specific, time invariant, and common knowledge, and $\varepsilon_{it}(1)$ is private information of firm i and is independently and identically distributed across firms and over time with zero mean. The term $(1 - a_{i,t-1})\theta_{EC}$ is the entry cost, where θ_{EC} is the entry cost parameter. This parameter is multiplied by $(1 - a_{i,t-1})$ because the entry cost is paid only by new entrants. If a supermarket does not operate a store, it can put its capital to other uses. Current profits of a nonactive firm, $\tilde{\Pi}_{it}(0)$, are equal to the value of the best outside opportunity. We assume that $\tilde{\Pi}_{it}(0) = \mu_i + \varepsilon_{it}(0)$, where μ_i is firm specific, time invariant, and common knowledge, and $\varepsilon_{it}(0)$ is private information of firm i and is independently and identically distributed across firms and over time with zero mean. Because the parameter μ_i cannot be identified separately from the average fixed cost $\theta_{FC,i}$, we normalize it to zero. Hereafter, the fixed cost $\theta_{FC,i}$ can be interpreted as including the opportunity cost μ_i . In this model, the vector of common knowledge state variables consists of the market size S_t and the indicators of incumbency status, i.e., $x_t = (S_t, a_{t-1})$, where $a_{t-1} = \{a_{i,t-1} : i = 1, 2, \dots, N\}$.

2.2. Strategies and Bellman Equations

The game has a Markov structure and we assume that firms play (stationary) 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 $\sigma = \{\sigma_i(x, \varepsilon_i)\}$ be a set of strategy functions or decision rules, one for each firm, with $\sigma_i: X \times R^{J+1} \rightarrow A$. Associated with a set of strategy functions σ , we can define a set of *conditional choice probabilities* $P^\sigma = \{P_i^\sigma(a_i|x)\}$ such that

$$(3) \quad P_i^\sigma(a_i|x) \equiv \Pr(\sigma_i(x, \varepsilon_i) = a_i|x) = \int I\{\sigma_i(x, \varepsilon_i) = a_i\} g_i(\varepsilon_i) d\varepsilon_i,$$

where $I\{\cdot\}$ is the indicator function. The probabilities $\{P_i^\sigma(a_i|x) : a_i \in A\}$ 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 σ .

Let $\pi_i^\sigma(a_i, x)$ be firm i 's current expected profit if it chooses alternative a_i and the other firms behave according to their respective strategies in σ .² By the independence of private information,

$$(4) \quad \pi_i^\sigma(a_i, x) = \sum_{a_{-i} \in A^{N-1}} \left(\prod_{j \neq i} P_j^\sigma(a_{-i}[j]|x) \right) \Pi_i(a_i, a_{-i}, x),$$

where a_{-i} is the vector with the actions of all players other than i and where $a_{-i}[j]$ is the j th firm's element in this vector. Let $\tilde{V}_i^\sigma(x, \varepsilon_i)$ 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 σ . By Bellman's principle of optimality, we can write

$$(5) \quad \tilde{V}_i^\sigma(x, \varepsilon_i) = \max_{a_i \in A} \left\{ \pi_i^\sigma(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} \left[\int \tilde{V}_i^\sigma(x', \varepsilon'_i) g_i(\varepsilon'_i) d\varepsilon'_i \right] f_i^\sigma(x'|x, a_i) \right\},$$

where $f_i^\sigma(x'|x, a_i)$ is the transition probability of x conditional on firm i choosing a_i and the other firms behaving according to σ :

$$(6) \quad f_i^\sigma(x'|x, a_i) = \sum_{a_{-i} \in A^{N-1}} \left(\prod_{j \neq i} P_j^\sigma(a_{-i}[j]|x) \right) f(x'|x, a_i, a_{-i}).$$

It is convenient to define value functions integrated over private information variables. Let $V_i^\sigma(x)$ be the integrated value function $\int \tilde{V}_i^\sigma(x, \varepsilon_i) g_i(d\varepsilon_i)$.

²In the terminology of Harsanyi (1995), the profit functions $\Pi_i(a_1, a_2, \dots, a_N, x)$ are the *conditional payoffs* and the expected profit functions $\pi_i^\sigma(a_i, x)$ are the *semiconditional payoffs*.

Based on this definition and (5), we can obtain the *integrated Bellman equation*

$$(7) \quad V_i^\sigma(x) = \int \max_{a_i \in A} \{v_i^\sigma(a_i, x) + \varepsilon_i(a_i)\} g_i(d\varepsilon_i)$$

with $v_i^\sigma(a_i, x) \equiv \pi_i^\sigma(a_i, x) + \beta \sum_{x' \in X} V_i^\sigma(x') f_i^\sigma(x'|x, a_i)$. The functions $v_i^\sigma(a_i, x)$ are called *choice-specific value functions*. The right-hand side of (7) is a contraction mapping in the space of value functions (see Aguirregabiria and Mira (2002)). Therefore, for each firm, there is a unique function $V_i^\sigma(x)$ that solves this functional equation for given σ .

2.3. Markov Perfect Equilibria

So far, σ is arbitrary and does not necessarily describe firms' equilibrium behavior. The following definition characterizes equilibrium strategies of all firms as best responses to one another.³

DEFINITION: A stationary Markov perfect equilibrium (MPE) in this game is a set of strategy functions σ^* such that for any firm i and for any $(x, \varepsilon_i) \in X \times R^{J+1}$,

$$(8) \quad \sigma_i^*(x, \varepsilon_i) = \arg \max_{a_i \in A} \{v_i^{\sigma^*}(a_i, x) + \varepsilon_i(a_i)\}.$$

Following Milgrom and Weber (1985), we can also represent a MPE in probability space.⁴ First, notice that for any set of strategies σ , in equilibrium or not, the functions π_i^σ , V_i^σ , and f_i^σ depend on players' strategies only through the choice probabilities P associated with σ . To emphasize this point and to define a MPE in probability space, we change the notation slightly and use the symbols π_i^P , V_i^P , and f_i^P , respectively, to denote these functions; v_i^P denotes the corresponding choice-specific value functions. Let σ^* be a set of MPE strategies and let P^* be the probabilities associated with these strategies. By definition, $P_i^*(a_i|x) = \int I\{a_i = \sigma_i^*(x, \varepsilon_i)\} g_i(\varepsilon_i) d\varepsilon_i$. Therefore, equilibrium probabilities are a fixed point. That is, $P^* = \Lambda(P^*)$, where, for any vector of probabilities P , $\Lambda(P) = \{\Lambda_i(a_i|x; P_{-i})\}$ and

$$(9) \quad \Lambda_i(a_i|x; P_{-i}) = \int I\left(a_i = \arg \max_{a \in A} \{v_i^{P^*}(a, x) + \varepsilon_i(a)\}\right) g_i(\varepsilon_i) d\varepsilon_i.$$

³In this paper we consider only pure-strategy equilibria because, in the following sense, they are observationally equivalent to mixed-strategy equilibria. Harsanyi's "purification theorem" established that a mixed-strategy equilibrium in a game of complete information can be interpreted as a pure-strategy equilibrium of a game of incomplete information (see Harsanyi (1973) and Fudenberg and Tirole (1991, pp. 230–234)). That is, the probability distribution of players' actions is the same under the two equilibria.

⁴Milgrom and Weber considered both discrete-choice and continuous-choice games. In their terminology, $\{P_i^\sigma\}$ are called *distributional strategies* and P^* is an *equilibrium in distributional strategies*.

We call the functions Λ_i *best response probability functions*. Given our assumptions on the distribution of private information, best response probability functions are well defined and continuous in the compact set of players' choice probabilities. By Brower's theorem, there exists at least one equilibrium. In general, the equilibrium is not unique.

Equilibrium probabilities solve the *coupled* fixed-point problems defined by (7) and (9). Given a set of probabilities P , we obtain value functions V_i^P as solutions of the N Bellman equations in (7), and given these value functions, we obtain best response probabilities using the right-hand side of (9).

2.4. An Alternative Best Response Mapping

We now provide an alternative best response mapping (in probability space) that avoids the solution of the N dynamic programming problems in (7).⁵ The evaluation of this mapping is computationally much simpler than the evaluation of the mapping $\Lambda(P)$, and it will prove more convenient for the estimation of the model.

Let P^* be an equilibrium and let $V_1^{P^*}, V_2^{P^*}, \dots, V_N^{P^*}$ be firms' value functions associated with this equilibrium. Because equilibrium probabilities are best responses, we can rewrite the Bellman equation (7) as

$$(10) \quad V_i^{P^*}(x) = \sum_{a_i \in A} P_i^*(a_i|x) [\pi_i^{P^*}(a_i, x) + e_i^{P^*}(a_i, x)] \\ + \beta \sum_{x' \in X} V_i^{P^*}(x') f^{P^*}(x'|x),$$

where $f^{P^*}(x'|x)$ is the transition probability of x induced by P^* .⁶ The term $e_i^{P^*}(a_i, x)$ is the expectation of $\varepsilon_i(a_i)$ conditional on x and on alternative a_i being the optimal response for player i , i.e., $e_i^{P^*}(a_i, x) \equiv E(\varepsilon_i(a_i)|x, \sigma_i^*(x, \varepsilon_i) = a_i)$. This conditional expectation is a function of a_i , $P_i^*(x)$, and the probability distribution g_i only. To see this, note that the event $\{\sigma_i^*(x, \varepsilon_i) = a_i\}$ is equivalent to $\{v_i^{P^*}(a_i, x) + \varepsilon_i(a_i) \geq v_i^{P^*}(a, x) + \varepsilon_i(a) \text{ for any } a \neq a_i\}$. Then

$$(11) \quad e_i^{P^*}(a_i, x) = \frac{1}{P_i^*(a_i|x)} \\ \times \int \varepsilon_i(a_i) I\{\varepsilon_i(a) - \varepsilon_i(a_i) \leq v_i^{P^*}(a_i, x) - v_i^{P^*}(a, x) \\ \forall a \neq a_i\} g_i(\varepsilon_i) d\varepsilon_i.$$

⁵In this subsection, we adapt to this context results from Aguirregabiria and Mira (2002).

⁶That is, $f^{P^*}(x'|x) = \sum_{a \in A^N} (\prod_{j=1}^N P_j^*(a_j|x)) f(x'|x, a)$.

The last expression shows that $e_i^{P^*}(a_i, x)$ depends on the primitives of the model only through the probability distribution of ε_i and the vector of value differences $\tilde{v}_i^{P^*}(x) \equiv \{v_i^{P^*}(a, x) - v_i^{P^*}(0, x) : a \in A\}$. The vector of choice probabilities $P_i^*(x)$ is also a function of g_i and $\tilde{v}_i^{P^*}(x)$, i.e., $P_i^*(a_i|x) = \Pr(\varepsilon_i(a) - \varepsilon_i(a_i) \leq v_i^{P^*}(a_i, x) - v_i^{P^*}(a, x) \ \forall a \neq a_i|x)$. Hotz and Miller showed that this mapping, which relates choice probabilities and value differences, is invertible (see Proposition 1, p. 501, in Hotz and Miller (1993)). Thus, the expectations $e_i^{P^*}(a_i, x)$ are functions of g_i and $P_i^*(x)$ only. The particular functional form of $e_i^{P^*}(a_i, x)$ depends on the probability distribution g_i . A well known case where $e_i^{P^*}(a_i, x)$ has a closed-form expression is when $\varepsilon_i(a)$ is independently and identically distributed with *extreme value* distribution and dispersion parameter σ . In this case,

$$(12) \quad e_i^P(a_i, x) = \text{Euler's constant} - \sigma \ln(P_i(a_i|x)).$$

In a binary choice model with $\varepsilon_i \sim N(0, \Omega)$ we have that⁷

$$(13) \quad e_i^P(a_i, x) = \frac{\text{var}(\varepsilon_i(a_i)) - \text{cov}(\varepsilon_i(0), \varepsilon_i(1))}{\sqrt{\text{var}(\varepsilon_i(1) - \varepsilon_i(0))}} \frac{\phi(\Phi^{-1}(P_i(a_i|x)))}{P_i(a_i|x)},$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the density and the cumulative distribution of the standard normal, respectively, and Φ^{-1} is the inverse function of Φ .

Taking equilibrium probabilities as given, expression (10) describes the vector of values $V_i^{P^*}$ as the solution of a system of linear equations. In vector form,

$$(14) \quad (I - \beta F^{P^*})V_i^{P^*} = \sum_{a_i \in A} P_i^*(a_i) * [\pi_i^{P^*}(a_i) + e_i^{P^*}(a_i)],$$

where I is the identity matrix; F^{P^*} is a matrix with transition probabilities $f^{P^*}(x'|x)$; $P_i^*(a_i)$, $\pi_i^{P^*}(a_i)$, and $e_i^{P^*}(a_i)$ are vectors of dimension $|X|$ that stack the corresponding state-specific elements; and $*$ represents the element-by-element or Hadamard product. Let $\Gamma_i(P^*) \equiv \{\Gamma_i(x; P^*) : x \in X\}$ be the solution to this system of linear equations, such that $V_i^{P^*}(x) = \Gamma_i(x; P^*)$. For arbitrary probabilities P , not necessarily in equilibrium, the mapping $\Gamma_i(P) \equiv (I - \beta F^P)^{-1} \{\sum_{a_i \in A} P_i(a_i) * [\pi_i^P(a_i) + e_i^P(a_i)]\}$ can be interpreted as a *valuation operator*: that is, $\Gamma_i(P) = \{\Gamma_i(x, P) : x \in X\}$, where $\Gamma_i(x, P)$ is the *expected value of firm i if the current state is x and all firms (including firm i) behave today and in the future according to their choice probabilities in P* . Therefore, looking at (9)

⁷To derive this expression, define the random variable $\tilde{\varepsilon} \equiv \varepsilon(1) - \varepsilon(0)$ and let σ^2 be $\text{var}(\tilde{\varepsilon})$. Note that for any constant k , $E(\varepsilon(1)|\tilde{\varepsilon} + k \geq 0) = \text{cov}(\varepsilon(1), \tilde{\varepsilon})/\sigma^2 E(\tilde{\varepsilon}|\tilde{\varepsilon} \geq -k) = \text{cov}(\varepsilon(1), \tilde{\varepsilon})/\sigma \phi(k/\sigma)/\Phi(k/\sigma)$. Similarly, we have that $E(\varepsilon(0)|\tilde{\varepsilon} + k \leq 0) = \text{cov}(\varepsilon(0), \tilde{\varepsilon})/\sigma^2 \times E(\tilde{\varepsilon}|\tilde{\varepsilon} \leq -k) = -\text{cov}(\varepsilon(0), \tilde{\varepsilon})/\sigma \phi(-k/\sigma)/\Phi(-k/\sigma)$.

we can characterize a MPE as a fixed point of a mapping $\Psi(P) \equiv \{\Psi_i(a_i|x; P)\}$ with

$$(15) \quad \Psi_i(a_i|x; P) = \int I\left(a_i = \arg \max_{a \in A} \left\{ \pi_i^P(a, x) + \varepsilon_i(a) + \beta \sum_{x' \in X} \Gamma_i(x'; P) f_i^P(x'|x, a) \right\} \right) g_i(\varepsilon_i) d\varepsilon_i.$$

Clearly, an equilibrium vector P^* is a fixed point of Ψ . The following lemma establishes that the reverse is also true and therefore equilibria can be described as the set of fixed points of the best response mapping Ψ . In particular, uniqueness or multiplicity of equilibria corresponds to uniqueness or multiplicity of the fixed points of Ψ .

REPRESENTATION LEMMA: *Under Assumptions 1–3, the set of fixed points of the best response mappings Λ and Ψ are identical.*

The only difference between best response mappings Λ_i and Ψ_i is that Ψ_i takes firm i 's future actions as given, whereas Λ_i does not. To evaluate Λ_i one has to solve a dynamic programming problem, whereas to obtain $\Gamma_i(P)$ and $\Psi_i(P)$ one has only to solve a system of linear equations. In the context of the estimation of the model, we will see that using mapping Ψ instead of Λ provides significant computational gains.

EXAMPLE 2: Consider Example 1 in Section 2.1. The vector of common knowledge state variables is $x_t = (S_t, a_{t-1})$. Suppose that the private information shocks $\varepsilon_{it}(0)$ and $\varepsilon_{it}(1)$ are normally distributed with zero means, and define $\sigma^2 \equiv \text{var}(\varepsilon_{it}(0) - \varepsilon_{it}(1))$. Under the alternative best response function in (15), a firm will be active if and only if $\{\varepsilon_{it}(0) - \varepsilon_{it}(1)\} \leq \{\pi_i^P(1, x_t) - \pi_i^P(0, x_t) + \beta \sum_{x' \in X} \Gamma_i(x', P) [f_i^P(x'|x_t, 1) - f_i^P(x'|x_t, 0)]\}$. Therefore,

$$(16) \quad \Psi_i(1|x_t; P) = \Phi\left(\frac{1}{\sigma} \left[\pi_i^P(1, x_t) - \pi_i^P(0, x_t) + \beta \sum_{x' \in X} \Gamma_i(x', P) [f_i^P(x'|x_t, 1) - f_i^P(x'|x_t, 0)] \right]\right).$$

The transition probabilities of the state conditional on firm i choosing a_i and the other firms behaving according to P have the form

$$(17) \quad f_i^P(S_{t+1}, a_t|S_t, a_{t-1}, a_i) = f_S(S_{t+1}|S_t) \prod_{j \neq i} P_j(0|x_t)^{1-a_{jt}} P_j(1|x_t)^{a_{jt}} I\{a_{it} = a_i\},$$

where $f_S(S_{t+1}|S_t)$ is the transition probability of market size, which is assumed to follow an exogenous Markov process. Expected current profits if the firm is not active are $\pi_i^P(0, x_t) = 0$ and if the firm is active are

$$(18) \quad \pi_i^P(1, x_t) = \theta_R S_t N_i^P(x_t) - \theta_{FC,i} - \theta_{EC}(1 - a_{i,t-1}),$$

where the expectation $N_i^P(x_t) \equiv \sum_{a_{-i} \in \{0,1\}^{N-1}} \Pr(a_{-i}|x_t)(2 + \sum_{j \neq i} a_{-i}[j])^{-2}$ is taken over other players' choice profiles and $\Pr(a_{-i}|x_t) = \prod_{j \neq i} P_j(0|x_t)^{1-a_{-i}[j]} \times P_j(1|x_t)^{a_{-i}[j]}$.

We now turn to the valuation operator $\Gamma_i(P)$. Given that the private information shocks are normally distributed, the expectation $e_i^P(a_i, x)$ has the form in expression (13). It is convenient to write $\pi_i^P(0, x_t)$ and $\pi_i^P(1, x_t)$ as $z_i^P(0, x_t)\theta_\pi$ and $z_i^P(1, x_t)\theta_\pi$, respectively, where θ_π is the vector of structural parameters $(\theta_R, \theta_{FC,1}, \dots, \theta_{FC,N}, \theta_{EC})'$, $z_i^P(0, x_t)$ is (in this example) a vector of zeros, and $z_i^P(1, x_t) \equiv \{S_t N_i^P(x_t), -D_i, a_{i,t-1} - 1\}$ with D_i a $1 \times N$ vector with a one at column i and zeros elsewhere. Using this notation, the right-hand side of the system of equations that define the valuation operator in (14) can be written as

$$(19) \quad \sum_{a_i \in A} P_i(a_i) * [\pi_i^P(a_i) + e_i^P(a_i)] = Z_i^P \theta_\pi + \lambda_i^P \sigma,$$

where Z_i^P is a matrix with rows $\{P_i(0|x)z_i^P(0, x) + P_i(1|x)z_i^P(1, x)\}$ and λ_i^P is a vector with elements $\phi(\Phi^{-1}(P_i(1|x)))$.⁸ Note that the structural parameters θ_π and σ do not appear as arguments in Z_i^P , λ_i^P , and $(I - \beta F^P)$. Thus multiplicative separability of θ_π and σ in expected profits carries over to the valuation mapping. Substituting (19) into (14) and solving for the value function, we get

$$(20) \quad \Gamma_i(P) = \Gamma_i^Z(P)\theta_\pi + \Gamma_i^\lambda(P)\sigma,$$

where $\Gamma_i^Z(P)$ is the matrix $(I - \beta F^P)^{-1} Z_i^P$ and $\Gamma_i^\lambda(P)$ is the vector $(I - \beta F^P)^{-1} \lambda_i^P$. For each state x , the valuation operator $\Gamma_i(x, P)$ collects the infinite sum of expected current profits and expected private information terms (the elements of $Z_i^P \theta_\pi$ and $\phi_i^P \sigma$), which may occur along all possible future

⁸By the symmetry of the normal density function, we have that $\phi(\Phi^{-1}(P_i(1|x))) = \phi(\Phi^{-1}(1 - P_i(1|x)))$. Also, note that $\text{var}(\varepsilon_i(0)) - \text{cov}(\varepsilon_i(0), \varepsilon_i(1)) + \text{var}(\varepsilon_i(1)) - \text{cov}(\varepsilon_i(0), \varepsilon_i(1)) = \sigma^2$. Therefore, we have

$$\begin{aligned} & P_i(0|x)e_i^P(0|x) + P_i(1|x)e_i^P(1|x) \\ &= \left\{ \frac{1}{\sigma} [\text{var}(\varepsilon_i(0)) - \text{cov}(\varepsilon_i(0), \varepsilon_i(1))] \phi(\Phi^{-1}(P_i(0|x))) \right. \\ & \quad \left. + \frac{1}{\sigma} [\text{var}(\varepsilon_i(1)) - \text{cov}(\varepsilon_i(0), \varepsilon_i(1))] \phi(\Phi^{-1}(P_i(1|x))) \right\} \\ &= \sigma \phi(\Phi^{-1}(P_i(1|x))). \end{aligned}$$

histories that originate from that state. Premultiplication by the corresponding row of $(I - \beta F^P)^{-1}$ means that each term in the sum is discounted and weighted by the probability of the corresponding history. The probabilities of histories, expected current period profits, and expected private information terms are computed under the assumption that all firms (including firm i) behave today and in the future according to the choice probabilities in P . Note that the structural parameters θ_π and σ do not appear as arguments in Z_i^P , λ_i^P , and $(I - \beta F^P)$. Thus multiplicative separability of θ_π and σ in expected profits carries over to the valuation mapping.

Finally, the best response functions Ψ_i have the form

$$(21) \quad \Psi_i(1|x_t; P) = \Phi\left(\tilde{z}_i^P(x_t) \frac{\theta_\pi}{\sigma} + \tilde{\lambda}_i^P(x_t)\right),$$

where

$$(22) \quad \begin{aligned} \tilde{z}_i^P(x_t) &\equiv z_i^P(1, x_t) - z_i^P(0, x_t) \\ &\quad + \beta \sum_{x' \in X} \Gamma_i^Z(x', P) [f_i^P(x'|x_t, 1) - f_i^P(x'|x_t, 0)], \\ \tilde{\lambda}_i^P(x_t) &\equiv \beta \sum_{x' \in X} \Gamma_i^\lambda(x', P) [f_i^P(x'|x_t, 1) - f_i^P(x'|x_t, 0)]. \end{aligned}$$

3. ESTIMATION

3.1. Econometric Model and Data Generating Process

Consider a researcher who observes players' actions and common knowledge state variables across M geographically separate markets over T periods, where M is large and T is small. This is a common sampling framework in empirical applications in industrial organization, which is given as

$$(23) \quad \text{data} = \{a_{mt}, x_{mt} : m = 1, 2, \dots, M; t = 1, 2, \dots, T\},$$

where m is the market subindex and $a_{mt} = (a_{1mt}, a_{2mt}, \dots, a_{Nmt})$. An important aspect of the data is whether players are the same across markets or not. We use the terminology *global players* and *local players*, respectively, to refer to these two cases. In our example of the model of market entry–exit we may have some large firms that—active or not—are potential entrants in every local market and some other firms that are potential entrants in only one local market. For instance, in the fast-food industry McDonald's would be a global player, whereas a family-owned fast-food outlet would be a local player. Our framework can accommodate both cases. However, we can allow for heterogeneity in the structural parameters across players only if those players' decisions are observed across all or most of the markets. To illustrate both cases, the Monte Carlo experiments that we present in Section 4 are for the model with *global*

players only and the empirical application in Section 5 is for a model with *local* players only.

The primitives $\{\Pi_i, g_i, f, \beta : i \in I\}$ are known to the researcher up to a finite vector of structural parameters $\theta \in \Theta \subset R^{|\Theta|}$. Primitives are twice continuously differentiable in θ . We now incorporate θ as an explicit argument in the equilibrium mapping Ψ . Let $\theta^0 \in \Theta$ be the true value of θ in the population. The researcher is interested in the estimation of θ^0 . Under Assumption 2 (i.e., conditional independence), the transition probability function f can be estimated from transition data using a standard maximum likelihood method and without solving the model. We focus on the estimation of the rest of the primitives and hereafter, for the sake of simplicity, we assume that β and the transition probability function are known. We consider the following assumption on the data generating process.

ASSUMPTION 5: Let $P_{mt}^0 \equiv \{\Pr(a_{mt} = a | x_{mt} = x) : (a, x) \in A^N \times X\}$ be the distribution of a_{mt} conditional on x_{mt} in market m at period t . (A) For every observation (m, t) in the sample, $P_{mt}^0 = P^0$ and $P^0 = \Psi(\theta^0, P^0)$. (B) Players expect P^0 to be played in future (out of sample) periods. (C) For any $\theta \neq \theta^0$ and P that solves $P = \Psi(\theta, P)$, it is the case that $P \neq P^0$. (D) The observations $\{a_{mt}, x_{mt} : m = 1, 2, \dots, M; t = 1, 2, \dots, T\}$ are independent across markets and $\Pr(x_{mt} = x) > 0$ for all x in X .

Assumption 5(A) establishes that the data have been generated by only one Markov perfect equilibrium.⁹ Thus even if the model has multiple equilibria, the researcher does not need to specify an equilibrium selection mechanism because the equilibrium that has been selected will be identified from the conditional choice probabilities in the data. Assumption 5(B) is necessary to accommodate dynamic models. Without it, we cannot compute the expected future payoffs of within-sample actions unless we specify the beliefs of players with regard to the probability of switching equilibria in the future. Assumption 5(C) is a standard identification condition.

Assumption 5(A) does not rule out models with multiple equilibria, but in its strictest version it would seem to rule out the presence of multiple equilibria in the data generating process. However, our notation allows for a flexible interpretation of Assumption 5(A) such that different equilibria are played across subsamples of markets known to the researcher. In this case, subsamples correspond to different “market types,” there is a finite number of types, and the vector x_{mt} of observable state variables is augmented with the time-invariant market type $\bar{x}_m$. Accordingly, the conditional choice probability vec-

⁹Moro (2003) introduced the assumption that only one equilibrium is present in the data in a somewhat different context. In his work the researcher observes a function of the equilibrium strategies rather than the equilibrium object itself; therefore, additional assumptions are needed to identify the selected equilibrium from the data.

tor P stacks Markov perfect equilibria for all market types. For instance, market types could correspond to different regions. Assumption 5(A) says that the data from different markets within the same region are generated by a single equilibrium. However, players in the markets of regions 1 and 2 may be playing different equilibria P_1^0 and P_2^0 that correspond to the same population parameter θ^0 ; i.e., $P_1^0 = \Psi(\theta^0, P_1^0)$ and $P_2^0 = \Psi(\theta^0, P_2^0)$. Furthermore, in Section 3.5 we extend our framework to allow for permanent unobserved heterogeneity and we show that in that case multiple equilibria may be present in the data generating process. Finally, Assumptions 5(A) and (B) can also be relaxed to allow for different equilibria to be played over time as long as the researcher knows (a) the exact time period when the behavior of players switched from one equilibrium to another and (b) the players' expectations about the equilibrium type that would be played in every period; e.g., if players did not anticipate the switch from one equilibrium type to another, the researcher knows this.

3.2. Maximum Likelihood Estimation

Define the *pseudo likelihood function*

$$(24) \quad Q_M(\theta, P) = \frac{1}{M} \sum_{m=1}^M \sum_{t=1}^T \sum_{i=1}^N \ln \Psi_i(a_{imt} | x_{mt}; P, \theta),$$

where P is an arbitrary vector of players' choice probabilities. We call this function a pseudo likelihood because the choice probabilities are not necessarily equilibrium probabilities associated with θ , but just best responses to an arbitrary vector P . Consider first the hypothetical case of a model with a unique equilibrium for each possible value of $\theta \in \Theta$. Then the maximum likelihood estimator (MLE) of θ^0 can be defined from the constrained multinomial likelihood

$$(25) \quad \hat{\theta}_{\text{MLE}} = \arg \max_{\theta \in \Theta} Q_M(\theta, P) \quad \text{subject to} \quad P = \Psi(\theta, P).$$

The computation of this estimator requires one to evaluate the mapping Ψ and the Jacobian matrix $\partial \Psi / \partial P'$ at many different values of P . Although evaluations of Ψ for different θ 's can be relatively cheap because we do not have to invert the matrix $(I - \beta F)$ in (14), evaluations for different P imply a huge cost when the dimension of the state space is large, because this matrix needs to be inverted each time. Therefore, this estimator can be impractical if the dimension of P is relatively large. For instance, that is the case in most models with heterogeneous players, because the dimension of the state space increases exponentially with the number of players. For that type of models, this estimator can be impractical even when the number of players is not too large.

An important complication in the estimation of dynamic games is that for some values of the structural parameters, the model can have multiple equilibria. With multiple equilibria the restriction $P = \Psi(\theta, P)$ does not define a

unique vector P , but a set of vectors. In this case, the MLE can be defined as

$$(26) \quad \hat{\theta}_{\text{MLE}} = \arg \max_{\theta \in \Theta} \left\{ \sup_{P \in (0,1)^{N \times |X|}} Q_M(\theta, P) \text{ subject to } P = \Psi(\theta, P) \right\}.$$

This estimator can be shown to be consistent, asymptotically normal, and efficient. However, in practice, this estimator can be extremely difficult to implement. Notice that for each trial value of θ , we have to compute all the vectors P that are an equilibrium associated with θ and then select the one with the maximum value for $Q_M(\theta, P)$. Finding all the Markov perfect equilibria of a dynamic game can be very difficult even for relatively simple models (see McKelvey and McLennan (1996)). Note also that with multiple equilibria, the number of evaluations of Ψ for different values of P increases very importantly. These problems motivate the pseudo likelihood estimators we develop in the following subsections.

3.3. Pseudo Maximum Likelihood Estimation

The PML estimators try to minimize the number of evaluations of Ψ for different vectors of players' probabilities P . Suppose that we know the population probabilities P^0 and consider the PML estimator¹⁰

$$(27) \quad \hat{\theta} \equiv \arg \max_{\theta \in \Theta} Q_M(\theta, P^0).$$

Under standard regularity conditions, this estimator is root- M consistent and asymptotically normal, and its asymptotic variance is $\Omega_{\theta\theta}^{-1}$, where $\Omega_{\theta\theta}$ is the variance of the pseudo score, i.e., $\Omega_{\theta\theta} \equiv E(\{\nabla_{\theta} s_m\} \{\nabla_{\theta} s_m\}')$, with $s_m \equiv \sum_{t=1}^T \sum_{i=1}^N \ln \Psi_i(a_{imt} | x_{mt}; P^0, \theta^0)$. Notice that to obtain this estimator we have to evaluate the mapping Ψ at only one value of players' choice probabilities.

However, this PML estimator is infeasible because P^0 is unknown. Suppose that we can obtain a $\sqrt{M}$ -consistent nonparametric estimator of P^0 . For instance, if there are no unobservable market characteristics, we can use a frequency estimator or a kernel method to estimate players' choice probabilities.¹¹ Let $\hat{P}^0$ be this nonparametric estimator. Then we can define the feasible two-step PML estimator $\hat{\theta}_{2S} \equiv \arg \max_{\theta \in \Theta} Q_M(\theta, \hat{P}^0)$. Proposition 1 presents the asymptotic properties of this estimator.

¹⁰Aguirregabiria (2004) described this PML estimator in a general class of econometric models, where the distribution of the endogenous variables is implicitly defined as an equilibrium of a fixed-point problem.

¹¹Note that if x includes time-invariant components that describe observable market types, Assumption 5(D) guarantees that consistent estimators of equilibrium choice probabilities can be obtained separately for each market type. However, if we believe that the equilibrium in the data varies across market types, smoothing cannot be used across observations that correspond to different types.

PROPOSITION 1: Suppose that (i) Assumptions 1–5 hold, (ii) $\Psi(\theta, P)$ is twice continuously differentiable, (iii) Θ is a compact set, (iv) $\theta^0 \in \text{int}(\Theta)$, and (v) let $\hat{P}^0 = (1/M) \sum_{m=1}^M q_m$ be an estimator of P^0 such that $\sqrt{M}(\hat{P}^0 - P^0) \rightarrow_d N(0, \Sigma)$. Then $\sqrt{M}(\hat{\theta}_{2S} - \theta^0) \rightarrow_d N(0, V_{2S})$, where

$$V_{2S} = \Omega_{\theta\theta}^{-1} + \Omega_{\theta\theta}^{-1} \Omega_{\theta P} \Sigma \Omega'_{\theta P} \Omega_{\theta\theta}^{-1}$$

and $\Omega_{\theta P} \equiv E(\{\nabla_{\theta} S_m\} \{\nabla_P S_m\}')$, where ∇_P represents the partial derivative with respect to P . Given that $\Omega_{\theta\theta}^{-1} \Omega_{\theta P} \Sigma \Omega'_{\theta P} \Omega_{\theta\theta}^{-1}$ is a positive definite matrix, we have that the feasible PML estimator is less efficient than the PML based on true P^0 , i.e., $V_{2S} \geq \Omega_{\theta\theta}^{-1}$. Furthermore, if $\hat{P}_A^0$ and $\hat{P}_B^0$ are two estimators of P^0 such that $\Sigma_A - \Sigma_B > 0$ (positive definite matrix), then the PML estimator based on $\hat{P}_B^0$ has lower asymptotic variance than the estimator based on $\hat{P}_A^0$.

Root- M consistency and asymptotic normality of $\hat{P}^0$, together with regularity conditions, are sufficient to guarantee root- M consistency and asymptotic normality of this PML estimator.¹² There are several reasons why this estimator is of interest. It deals with the problem of indeterminacy associated with multiple equilibria. Furthermore, repeated solutions of the dynamic game are avoided and this can result in significant computational gains.

EXAMPLE 3: Consider the entry–exit model of Examples 1 and 2. Suppose that we have a random sample of markets where the N firms are potential entrants. We observe market size and entry decisions at two consecutive periods: $\{S_{m1}, a_{m1}, S_{m2}, a_{m2} : m = 1, 2, \dots, M\}$. Nonparametric estimates of choice and transition probabilities can be obtained using sample frequencies or a kernel method. Given these estimates, we can construct the matrices $Z_i^{\hat{P}_0}$ and $\Gamma_i^Z(\hat{P}_0)$, and the vectors $\lambda_i^{\hat{P}_0}$ and $\Gamma_i^\lambda(\hat{P}_0)$, as defined in Example 2. Then, using the formulas in (22) evaluated at the vector of estimates $\hat{P}_0$, we can construct for every sample point the vector $\tilde{z}_i^{\hat{P}_0}(x_{m2})$ and the value $\tilde{\lambda}_i^{\hat{P}_0}(x_{m2})$. The pseudo likelihood function that the two-step PML estimator maximizes is¹³

$$\begin{aligned} (28) \quad Q_M(\theta, \hat{P}_0) &= M^{-1} \sum_{m=1}^M \sum_{i=1}^N a_{im2} \ln \Psi_i(1|x_{m2}; \hat{P}_0) \\ &\quad + (1 - a_{im2}) \ln \Psi_i(0|x_{m2}; \hat{P}_0) \\ &= M^{-1} \sum_{m=1}^M \sum_{i=1}^N \ln \Phi \left([2a_{im2} - 1] \left[\tilde{z}_i^{\hat{P}_0}(x_{m2}) \frac{\theta_\pi}{\sigma} + \tilde{\lambda}_i^{\hat{P}_0}(x_{m2}) \right] \right). \end{aligned}$$

¹²Note that Assumptions 1 and 3 on the distribution of ε and twice continuous differentiability of the primitives with respect to θ imply regularity condition (ii).

¹³Only the second observation for each market is used because incumbency status is not known for the first observation.

This is the log-likelihood of a standard probit model where the coefficient of the explanatory variable $\tilde{\lambda}_i^{\hat{P}_0}(x_{m2})$ is constrained to be equal to one. This pseudo likelihood is globally concave in θ_π/σ . Furthermore, evaluation of $Q_M(\theta, \hat{P}_0)$ at different values of θ is very simple and does not require one to recalculate the matrix $\Gamma_i^Z(\hat{P}_0)$ and the vector $\Gamma_i^\lambda(\hat{P}_0)$. Therefore, the computation of this two-step PML estimator is very straightforward.

However, the two-step PML has some drawbacks. First, its asymptotic variance depends on the variance Σ of the nonparametric estimator $\hat{P}^0$. Therefore, it can be very inefficient when Σ is large. Second, and more important, for the sample sizes available in actual applications, the nonparametric estimator of P^0 can be very imprecise. Note that a curse of dimensionality in estimation may arise from the number of players as well as from the number of state variables. Even when the number of players is not too large (e.g., five players), lack of precision in the first step can generate serious finite sample biases in the estimator of structural parameters as illustrated in our Monte Carlo experiments in Section 4.¹⁴ Third, for some models it is not possible to obtain consistent nonparametric estimates of P^0 ; this is the case in models with unobservable market characteristics.

3.4. Nested Pseudo Likelihood Method

The *nested pseudo likelihood* (NPL) method is a recursive extension of the two-step PML estimator. Let $\hat{P}_0$ be an initial guess of the vector of players' choice probabilities. It is important to emphasize that this guess need not be a consistent estimator of P^0 . Given $\hat{P}_0$, the NPL algorithm generates a sequence of estimators $\{\hat{\theta}_K : K \geq 1\}$, where the *K-stage* estimator is defined as

$$(29) \quad \hat{\theta}_K = \arg \max_{\theta \in \Theta} Q_M(\theta, \hat{P}_{K-1})$$

and the probabilities $\{\hat{P}_K : K \geq 1\}$ are obtained recursively as

$$(30) \quad \hat{P}_K = \Psi(\hat{\theta}_K, \hat{P}_{K-1}).$$

That is, $\hat{\theta}_1$ maximizes the pseudo likelihood $Q_M(\theta, \hat{P}_0)$. Given $\hat{P}_0$ and $\hat{\theta}_1$, we obtain a new vector of probabilities by applying a single iteration in the best response mapping, i.e., $\hat{P}_1 = \Psi(\hat{\theta}_1, \hat{P}_0)$; then $\hat{\theta}_2$ maximizes the pseudo likelihood $Q_M(\theta, \hat{P}_1)$ and so on. This sequence is well defined as long as, for each

¹⁴In our Monte Carlo examples we use frequency estimators in the first step. Replacing these by smooth nonparametric estimators can reduce finite sample bias of the two-step PML estimator. See Pakes, Östrovsky, and Berry (2004) and our discussion in Section 4.

value of P , there is a unique value of θ that maximizes the pseudo likelihood function, which we assume hereafter.

EXAMPLE 4: We follow up on the example of the entry and exit model. For notational simplicity, we use θ to represent the vector of identified structural parameters θ_π/σ . Let $\hat{\theta}_1$ be the two-step PML estimator of θ that we described in Example 3. Then new estimates of choice probabilities can be obtained from the elements of the maximized pseudo likelihood

$$(31) \quad \hat{P}_{i,1}(1|x) = \Phi(\tilde{z}_i^{\hat{P}_0}(x)\hat{\theta}_1 + \tilde{\lambda}_i^{\hat{P}_0}(x)),$$

where $\{\tilde{z}_i^{\hat{P}_0}(x) : x \in X\}$ and $\{\tilde{\lambda}_i^{\hat{P}_0}(x) : x \in X\}$ are the vectors and scalars, respectively, that we used to obtain $\hat{\theta}_1$, i.e., they are based on the initial choice probabilities $\hat{P}_0$. Given the new vector of probabilities $\hat{P}_1$, we calculate new matrices $Z_i^{\hat{P}_1}$ and $\Gamma_i^Z(\hat{P}_1)$, and new vectors $\lambda_i^{\hat{P}_1}$ and $\Gamma_i^\lambda(\hat{P}_1)$, and we use them to construct $\tilde{z}_i^{\hat{P}_1}(x_{m2})$ and $\tilde{\lambda}_i^{\hat{P}_1}(x_{m2})$ according to the formulas in (22). Then we use these values to obtain a new estimator of θ based on a probit model with likelihood function $\sum_m \sum_i \ln \Phi([2a_{im2} - 1][\tilde{z}_i^{\hat{P}_1}(x_{m2})\theta + \tilde{\lambda}_i^{\hat{P}_1}(x_{m2})])$. We apply this procedure recursively. Note that this pseudo likelihood function is globally concave in θ . Therefore, for each value of P , there is a unique value of θ that maximizes the pseudo likelihood.

If the initial guess $\hat{P}_0$ is a consistent estimator, all elements of the sequence of estimators $\{\hat{\theta}_K : K \geq 1\}$ are consistent.¹⁵ Here we focus instead on the estimator we obtain in the limit. If the sequence $\{\hat{\theta}_K, \hat{P}_K\}$ converges, regardless of the initial guess, its limit $(\hat{\theta}, \hat{P})$ satisfies the following two properties: $\hat{\theta}$ maximizes the pseudo likelihood $Q_M(\theta, \hat{P})$ and $\hat{P} = \Psi(\hat{\theta}, \hat{P})$. We call any pair (θ, P) that satisfies these properties a *NPL fixed point*. In this section we show that a NPL fixed point exists in every sample and that if more than one exists, the one with the highest value of the pseudo likelihood is a consistent estimator. Because our method uses NPL iterations to find NPL fixed points, convergence is a concern. Although we have not proved convergence of the NPL algorithm in general, we have always obtained convergence in our Monte Carlo experiments and applications.

It is useful to introduce the following *NPL operator* ϕ_M to describe NPL iterations in (29) and (30):

$$(32) \quad \phi_M(P) \equiv \Psi(\tilde{\theta}_M(P), P) \quad \text{where} \quad \tilde{\theta}_M(P) \equiv \arg \max_{\theta \in \Theta} Q_M(\theta, P).$$

¹⁵This is a straightforward extension of the consistency of the two-step estimator in Proposition 1. See Aguirregabiria and Mira (2002) for the proof in a single-agent context.

A NPL fixed point is a pair (θ, P) such that P is a fixed point of ϕ_M and $\theta = \tilde{\theta}_M(P)$. If the maximizer $\tilde{\theta}_M(P)$ is unique for every P , then the mapping $\tilde{\theta}_M$ is continuous by the theorem of the maximum, and the NPL operator ϕ_M is continuous in the compact and convex set $[0, 1]^{N|X|}$. Thus, for any given sample, Brower's theorem guarantees the existence of at least one NPL fixed point. However, uniqueness need not follow in general. The NPL estimator $(\hat{\theta}_{\text{NPL}}, \hat{P}_{\text{NPL}})$ is defined as the NPL fixed point associated with the maximum value of the pseudo likelihood. In practical terms this means that the researcher should initiate the NPL algorithm with different P guesses and, if different limits are attained, he should select the one that maximizes the value of the pseudo likelihood. Let Y_M be the set of NPL fixed points, i.e., $Y_M \equiv \{(\theta, P) \in \Theta \times [0, 1]^{N|X|} : \theta = \tilde{\theta}_M(P) \text{ and } P = \phi_M(P)\}$. Then the NPL estimator can be defined as

$$(33) \quad (\hat{\theta}_{\text{NPL}}, \hat{P}_{\text{NPL}}) = \arg \max_{(\theta, P) \in Y_M} Q_M(\theta, P).$$

Proposition 2 establishes the large sample properties of the NPL estimator.

PROPOSITION 2: *Let $(\hat{\theta}_{\text{NPL}}, \hat{P}_{\text{NPL}})$ be the NPL estimator, i.e., the NPL fixed point in the sample with the maximum value of the pseudo likelihood. Consider the following population counterparts of the sample functions Q_M , $\tilde{\theta}_M$, and ϕ_M : $Q_0(\theta, P) \equiv E(Q_M(\theta, P))$, $\tilde{\theta}_0(P) \equiv \arg \max_{\theta \in \Theta} Q_0(\theta, P)$, and $\phi_0(P) \equiv \Psi(\tilde{\theta}_0(P), P)$. The set of population NPL fixed points is $Y_0 \equiv \{(\theta, P) \in \Theta \times [0, 1]^{N|X|} : \theta = \tilde{\theta}_0(P) \text{ and } P = \phi_0(P)\}$. Suppose that (i) Assumptions 1–5 hold, (ii) $\Psi(\theta, P)$ is twice continuously differentiable, (iii) Θ is a compact set, (iv) $\theta^0 \in \text{int}(\Theta)$, (v) (θ^0, P^0) is an isolated population NPL fixed point, i.e., it is unique, or else there is an open ball around it that does not contain any other element of Y_0 , (vi) there exists a closed neighborhood of P^0 , $N(P^0)$, such that, for all P in $N(P^0)$, Q_0 is globally concave in θ and $\partial^2 Q_0(\theta, P^0) / \partial \theta \partial \theta'$ is a nonsingular matrix, and (vii) the operator $\phi_0(P) - P$ has a nonsingular Jacobian matrix at P^0 . Then $\hat{\theta}_{\text{NPL}}$ is a consistent estimator and $\sqrt{M}(\hat{\theta}_{\text{NPL}} - \theta^0) \rightarrow_d N(0, V_{\text{NPL}})$, with*

$$(34) \quad V_{\text{NPL}} = [\Omega_{\theta\theta} + \Omega_{\theta P}(I - \nabla_P \Psi')^{-1} \nabla_{\theta} \Psi]^{-1} \\ \times \Omega_{\theta\theta} [\Omega_{\theta\theta} + \nabla_{\theta} \Psi'(I - \nabla_P \Psi)^{-1} \Omega'_{\theta P}]^{-1},$$

where $\nabla_P \Psi$ is the Jacobian matrix $\nabla_P \Psi(P^0, \theta^0)$. Furthermore, if the matrix $\nabla_P \Psi$ has all its eigenvalues between 0 and 1, the NPL estimator is more efficient than the infeasible PML estimator, i.e., $V_{\text{NPL}} < \Omega_{\theta\theta}^{-1} < V_{2S}$.

Nested pseudo likelihood estimation maintains the two main advantages of PML: it is feasible in models with multiple equilibria and it minimizes the number of evaluations of the mapping Ψ for different values of P . Furthermore, it

addresses the three limitations of the two-stage PML that were mentioned previously. First, under some conditions on the Jacobian matrix $\nabla_P \Psi$, the NPL is asymptotically more efficient than the infeasible PML and therefore more efficient than any two-step PML estimator for whatever initial estimator of P^0 we use. In other words, imposing the equilibrium condition in the sample can yield asymptotic efficiency gains relative to the two-step PML estimators. The last part of Proposition 2 provides one set of sufficient conditions for such a result to hold. Second, in small samples the NPL estimator reduces the finite sample bias generated by imprecise estimates of P^0 . This point is illustrated in the Monte Carlo experiments in Section 4. Third, consistency of the NPL estimator does not require that we start the algorithm with a consistent estimator of choice probabilities. A particularly important implication of this is that NPL may be applied to situations in which some time-invariant market characteristics are unobserved by the researcher. We develop this case in some detail in the next subsection.

In the [proof](#) of Proposition 2 in the [Appendix](#) we show that sample NPL fixed points converge in probability to population NPL fixed points. An implication of this result is that if the researcher is able to establish uniqueness of the NPL fixed point in the population, any sample NPL fixed point is a consistent estimator and there is no need to search for—and compare—multiple NPL fixed points. However, if the population function has more than one NPL fixed point, a “poorly behaved” initial guess $\hat{P}_0$ might identify a NPL fixed point that is not (θ^0, P^0) . Thus a comparison between all NPL fixed points in the sample is needed to guarantee consistency. We show that Assumption 5(A) yields identification because it implies that (θ^0, P^0) uniquely maximizes the pseudo likelihood $Q_0(\theta, P)$ in the set of population NPL fixed points. Therefore, the estimator $(\hat{\theta}_{\text{NPL}}, \hat{P}_{\text{NPL}})$ is to the sample function Q_M what the true parameter (θ^0, P^0) is to the limiting function Q_0 . However, proving consistency of the NPL estimator is less straightforward than for standard extremum estimators because the definition of the NPL estimator combines both maximization and fixed-point conditions. If the population has multiple NPL fixed points, we require additional regularity conditions that are listed as (v)–(vii) in Proposition 2. Condition (vii) is the more substantive. We need this condition in order to guarantee that there exists a sample NPL fixed point close to (θ^0, P^0) . The condition states that the Jacobian of the population NPL operator does not have any eigenvalues in the unit circle. It is the vector-valued equivalent of a requirement that a scalar differentiable function with a fixed point should cross the 45° line at the fixed point, rather than being tangent. Condition (v) rules out the possibility of a continuum of NPL fixed points around (θ^0, P^0) . This condition is actually implied by (vii), but we state it separately in spite of its redundancy because it is used independently in the proof. Condition (vi) states that the population pseudo likelihood is globally concave in θ for all P in a neighborhood of P^0 and that this concavity is strict at P^0 . The proof of consistency requires that $\tilde{\theta}_0(P)$, the maximizer of the population pseudo likelihood

in θ , be a single-valued and continuous function of P in a neighborhood of P^0 . Condition (vi) is not necessary for this, but it is a sufficient condition on the primitives of the model, and it holds in our example and in the models that we estimate in Sections 4 and 5.

Finding all NPL fixed points may not be a simple task. This is similar to the problem that arises in extremum estimators when the criterion function has more than one local extremum. Clearly, the problem does not arise if the sample NPL fixed point is unique and stable, but we do not have general enough sufficient conditions on the primitives of the model that guarantee this property. However, it should be stressed that multiplicity of NPL fixed points does not follow from multiplicity of equilibria of the model. A fixed point of the NPL operator is not just a solution to $P = \Psi(\theta, P)$. The NPL fixed points are special in that $\hat{\theta}$ maximizes $Q_M(\theta, \hat{P})$. A model may have multiple equilibria yet only one NPL fixed point. For instance, Example 5 shows that when the number of structural parameters to estimate is equal to the number of free probabilities in P , then the NPL fixed point is always unique and stable.

EXAMPLE 5: Returning to the entry–exit model of our previous examples, we consider specifications in which the number of structural parameters in θ is exactly equal to the number of free probabilities in P ; i.e., the model is just identified. A concrete example can be described as follows. There are only two firms in the M markets ($N = 2$). Market size is constant (i.e., $S_{mt} = 1$ for all (m, t)) and therefore the vector of state variables x_t is $(a_{1,t-1}, a_{2,t-1})$, which belongs to the set $X = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. The entry cost is $(1 - a_{i,t-1})(\theta_{EC,1} + a_{-i,t-1}\theta_{EC,2})$, where $\theta_{EC,1}$ and $\theta_{EC,2}$ are parameters. The parameter $\theta_{EC,2}$ captures the notion that entry in a market where there is an incumbent can be more costly than entry in a market with no active firms. Firms are homogeneous in their fixed operating costs. Suppose the equilibrium in the data generating process (DGP) is symmetric between the two firms. Then the model has four structural parameters, $\theta = (\theta_R/\sigma, \theta_{FC}/\sigma, \theta_{EC,1}/\sigma, \theta_{EC,2}/\sigma)'$, and four free probabilities, $P(1|0, 0)$ (i.e., probability of entry in a market with no active firms), $P(1|0, 1)$ (i.e., probability of entry in a market with an incumbent monopolist), $P(1|1, 0)$ (i.e., probability that the incumbent monopolist will stay), and $P(1|1, 1)$ (i.e., probability that a duopolist will stay). In general, the model will have multiple equilibria for the true parameter θ^0 .¹⁶

In a probit model with discrete explanatory variables, when the number of parameters is the same as the number of free conditional choice probabilities, one can show that the maximum likelihood estimator of the structural parameters is the value that makes the predicted probabilities equal to the

¹⁶For a particular parameterization of this example, which has $\sigma = 1$, $\theta_R = 17.28$, $\theta_{FC} = 3.22$, $\theta_{EC,1} = 0.10$, and $\theta_{EC,2} = 0$, Pesendorfer and Schmidt-Dengler (2004) showed that at least five different equilibria exist.

corresponding sample frequencies.¹⁷ In our case, this implies that the PML estimator $\tilde{\theta}_M(P)$ is the value of θ that solves the system of $N|X|$ equations with $N|X|$ unknowns,¹⁸

$$(35) \quad \Phi(\tilde{z}_i^P(x)\theta + \tilde{\lambda}_i^P(x)) = \hat{P}_i^{\text{Freq}}(1|x)$$

for any $x \in X$ and any $i \in \{1, 2, \dots, N\}$,

where $\hat{P}_i^{\text{Freq}}(1|x) \equiv (\sum_{m=1}^M I\{x_m = x\}a_{im})/(\sum_{m=1}^M I\{x_m = x\})$. Solving for θ , we have that

$$(36) \quad \tilde{\theta}_M(P) = \begin{pmatrix} \tilde{z}_1^P(x^1) \\ \tilde{z}_1^P(x^2) \\ \vdots \\ \tilde{z}_N^P(x^{|X|}) \end{pmatrix}^{-1} \left[\begin{pmatrix} \Phi^{-1}(\hat{P}_1^{\text{Freq}}(1|x^1)) \\ \Phi^{-1}(\hat{P}_1^{\text{Freq}}(1|x^2)) \\ \vdots \\ \Phi^{-1}(\hat{P}_N^{\text{Freq}}(1|x^{|X|})) \end{pmatrix} - \begin{pmatrix} \tilde{\lambda}_1^P(x^1) \\ \tilde{\lambda}_1^P(x^2) \\ \vdots \\ \tilde{\lambda}_N^P(x^{|X|}) \end{pmatrix} \right],$$

where the rank condition for identification implies that the matrix $(\tilde{z}_1^P(x^1)', \dots, \tilde{z}_N^P(x^{|X|})')$ is nonsingular for P in a neighborhood of P^0 . Solving the expression for $\tilde{\theta}_M(P)$ into the best response mapping Ψ in (21), we obtain the NPL operator (ϕ_M)

$$(37) \quad \phi_M(P) = \begin{pmatrix} \Phi(\tilde{z}_1^P(x^1)\tilde{\theta}_M(P) + \tilde{\lambda}_1^P(x^1)) \\ \Phi(\tilde{z}_1^P(x^2)\tilde{\theta}_M(P) + \tilde{\lambda}_1^P(x^2)) \\ \vdots \\ \Phi(\tilde{z}_N^P(x^{|X|})\tilde{\theta}_M(P) + \tilde{\lambda}_N^P(x^{|X|})) \end{pmatrix} = \begin{pmatrix} \hat{P}_1^{\text{Freq}}(1|x^1) \\ \hat{P}_1^{\text{Freq}}(1|x^2) \\ \vdots \\ \hat{P}_N^{\text{Freq}}(1|x^{|X|}) \end{pmatrix},$$

where the last equality follows from Ψ being the left-hand side of the first order conditions in (35). That is, for any vector P , we have that $\phi_M(P) = \hat{P}^{\text{Freq}}$. The

¹⁷The result follows from the form of the first order conditions, which in our case are the pseudo likelihood equations

$$\sum_{m=1}^M \sum_{i=1}^N W_i(x_m, P, \theta)(a_{im} - \Phi(\tilde{z}_i^P(x)\theta + \tilde{\lambda}_i^P(x))) = 0,$$

where $W_i(x_m, P, \theta) \equiv [\Phi(\tilde{z}_i^P(x)\theta + \tilde{\lambda}_i^P(x))\Phi(-\tilde{z}_i^P(x)\theta - \tilde{\lambda}_i^P(x))]^{-1}\phi(\tilde{z}_i^P(x)\theta + \tilde{\lambda}_i^P(x))\tilde{z}_i^P(x)'$.

¹⁸If the state space has $|X|$ points, there are N players, and the equilibrium is not symmetric, there are $2N|X|$ conditional choice probabilities. However, the number of free probabilities is $N|X|$ because $P_i(0|x) + P_i(1|x) = 1$ for all (i, x) .

model can have multiple equilibria, but the NPL mapping is constant and has a unique fixed point that is stable: $\hat{P}_{\text{NPL}} = \hat{P}_M^{\text{Freq}}$ and

$$(38) \quad \hat{\theta}_{\text{NPL}} = \begin{pmatrix} \tilde{z}_1^{\hat{P}^{\text{Freq}}}(x^1) \\ \tilde{z}_1^{\hat{P}^{\text{Freq}}}(x^2) \\ \vdots \\ \tilde{z}_N^{\hat{P}^{\text{Freq}}}(x^{|X|}) \end{pmatrix}^{-1} \left[\begin{pmatrix} \Phi^{-1}(\hat{P}_1^{\text{Freq}}(1|x^1)) \\ \Phi^{-1}(\hat{P}_1^{\text{Freq}}(1|x^2)) \\ \vdots \\ \Phi^{-1}(\hat{P}_N^{\text{Freq}}(1|x^{|X|})) \end{pmatrix} - \begin{pmatrix} \tilde{\lambda}_1^{\hat{P}^{\text{Freq}}}(x^1) \\ \tilde{\lambda}_1^{\hat{P}^{\text{Freq}}}(x^2) \\ \vdots \\ \tilde{\lambda}_N^{\hat{P}^{\text{Freq}}}(x^{|X|}) \end{pmatrix} \right].$$

Notice that regardless of the value of P that we use to initialize the NPL procedure, we always converge to this estimator in two NPL iterations. Also note that the two-step PML is inconsistent if P is inconsistent, but the NPL estimator is consistent.

Although the example illustrates the difference between multiple equilibria and multiple NPL fixed points, assuming a just identified model is too restrictive for most relevant applications of dynamic games. However, our limited experience with the NPL algorithm in the Monte Carlo experiments (in Section 4) and in our empirical application (in Section 5) suggests that the uniqueness of the NPL fixed point in the presence of multiple equilibria may be more general than the just identified case. For every one of the 6,000 Monte Carlo samples and for the actual sample in our application we always converged to the same NPL fixed point, regardless of the initial values that we considered for P .¹⁹

3.5. NPL with Permanent Unobserved Heterogeneity

The econometric model in Section 3.1 allowed for the case in which the vector of common knowledge state variables x_{mt} includes time-invariant market characteristics $\bar{x}_m$ with discrete and finite support. We now consider the estimation of models in which some of the time-invariant common knowledge characteristics are unobservable. For instance, in the entry–exit model, we may have a profit function

$$(39) \quad \tilde{\Pi}_{imt}(1) = \theta_R S_{mt} / \left(2 + \sum_{j \neq i} a_{jmt} \right)^2 - \theta_{\text{FC},i} - \theta_{\text{EC}}(1 - a_{im,t-1}) + \omega_m + \varepsilon_{imt},$$

where ω_m is a random effect interpreted as a time-invariant market characteristic that affects firms' profits, which is common knowledge to the players but

¹⁹In those cases where the NPL algorithm has multiple fixed points, the combination of the NPL with a stochastic algorithm is a promising approach, which we have explored in other work (Aguirregabiria and Mira (2005)). See Rust (1997) and Pakes and McGuire (2001) for earlier applications of stochastic algorithms to dynamic programming and equilibrium problems, the latter in the context of the Ericson and Pakes (1995) model.

unobservable to the econometrician. We make the following assumptions on the distribution of observable and unobservable market characteristics:

ASSUMPTION 6: (A) *The vector of unobservable common knowledge market characteristics ω_m has a discrete and finite support $\Omega = \{\omega^1, \omega^2, \dots, \omega^L\}$. (B) Conditional on $\bar{x}_m$, ω_m is independently and identically distributed across markets with probability mass function $\varphi_l(\bar{x}_m) \equiv \Pr(\omega_m = \omega^l | \bar{x}_m)$. (C) The vector ω_m does not enter into the conditional transition probability of x_{mt} , i.e., $\Pr(x_{m,t+1} | a_{mt}, x_{mt}, \omega_m) = f(x_{m,t+1} | a_{mt}, x_{mt})$.*

Assumption 6(B) allows the unobserved component of ω_m to be correlated with observable (fixed) market characteristics. Assumption 6(C) states that all markets are homogeneous with respect to (exogenous) transitions and it implies that the transition probability function f can still be estimated from transition data without solving the model.

The vector of structural parameters to be estimated, θ , now includes the parameters in the conditional distributions of ω . The vector P now stacks the distributions of players' actions conditional on all values of observable and unobservable common knowledge state variables. We use the notation $P = \{P_1, P_2, \dots, P_L\}$ with $P_l \equiv \{\Pr(a_{mt} = a | x_{mt} = x, \omega_m = \omega^l) : (a, x) \in A^N \times X\}$. We adapt Assumptions 5(A), (B), (D) on the *data generating process* as follows:

ASSUMPTION 5' (A), (B), (D): *Let $P_{mt}^0 \equiv \{\Pr(a_{mt} = a | x_{mt} = x, \omega_m = \omega) : (a, x, \omega) \in A^N \times X \times \Omega\}$ be the distributions of a_{mt} conditional on x_{mt} and ω_m in market m at period t . (A) For every observation (m, t) in the sample, $P_{mt}^0 = P^0$ and $P_l^0 = \Psi(\theta^0, P_l^0, \omega^l)$ for any l . (B) Players expect P^0 to be played in future (out of sample) periods. (D) The observations $\{a_{mt}, x_{mt} : m = 1, 2, \dots, M; t = 1, 2, \dots, T\}$ are independent across markets and $\Pr(x_{mt} = x) > 0$ for all x in X .*

Assumption 5' still states that only one equilibrium is played in the data *conditional on market type* $(\bar{x}_m, \omega_m)$, which is partly unobservable to the econometrician but known to players. This has the important implication that the data generating process may correspond to multiple equilibria. In our entry–exit example, markets that are observationally equivalent to the econometrician may have different probabilities of entry and exit because the random effect component of profits ω is different. Furthermore, differences across unobserved market types need not even be payoff relevant. That is, a “sunspot” mechanism may sort otherwise identical markets into two or more unobserved market types that select different equilibria. Our estimation framework allows for such a data generating process, as long as the econometrician knows the number of

unobserved types and the parameter θ is identified (see Assumption 5'(C) in the subsequent text).²⁰

To obtain the pseudo likelihood function, we integrate the best response probabilities over the conditional distribution of unobservable market characteristics. We have that

$$(40) \quad \ln \Pr(\text{data}|\theta, P) = \sum_{m=1}^M \ln \Pr(\tilde{a}_m, \tilde{x}_m|\theta, P) \\ = \sum_{m=1}^M \ln \left(\sum_{l=1}^L \varphi_l(\bar{x}_m) \Pr(\tilde{a}_m, \tilde{x}_m|\omega^l; \theta, P) \right),$$

where $\tilde{a}_m = \{a_{mt} : t = 1, 2, \dots, T\}$ and $\tilde{x}_m = \{x_{mt} : t = 1, 2, \dots, T\}$. Applying the Markov structure of the model, and Assumption 6(C), we get

$$(41) \quad \Pr(\tilde{a}_m, \tilde{x}_m|\omega^l; \theta, P) \\ = \left(\prod_{t=1}^T \Pr(a_{mt}|x_{mt}, \omega^l; \theta, P) \right) \left(\prod_{t=2}^T \Pr(x_{mt}|a_{m,t-1}, x_{m,t-1}, \omega^l; \theta, P) \right) \\ \times \Pr(x_{m1}|\omega^l; \theta, P) \\ = \left(\prod_{t=1}^T \prod_{i=1}^N \Psi_i(a_{imt}|x_{mt}, \omega^l; P_l, \theta) \right) \left(\prod_{t=2}^T f(x_{mt}|a_{m,t-1}, x_{m,t-1}; \theta) \right) \\ \times \Pr(x_{m1}|\omega^l; \theta, P).$$

Therefore,

$$(42) \quad \ln \Pr(\text{data}|\theta, P) \\ = \sum_{m=1}^M \ln \left(\sum_{l=1}^L \varphi_l(\bar{x}_m) \left(\prod_{t=1}^T \prod_{i=1}^N \Psi_i(a_{imt}|x_{mt}, \omega^l; P_l, \theta) \right) \right. \\ \left. \times \Pr(x_{m1}|\omega^l; \theta, P) \right) \\ + \sum_{m=1}^M \sum_{l=1}^L \ln f(x_{mt}|a_{m,t-1}, x_{m,t-1}; \theta).$$

²⁰With unobserved heterogeneity and multiple equilibria in the DGP, identification places a heavier burden on the data. In this sense, the strict version of Assumption 5 (“only one equilibrium is played in the data”) can be thought of as an identifying assumption.

The first component on the right-hand side is the pseudo likelihood function $Q_M(\theta, P)$. The second component is the part of the likelihood associated with transition data. As previously mentioned, the transition probability functions f can still be estimated separately from transition data without solving the model. Likewise, we are not considering estimation of the marginal distribution of the observed market type $\bar{x}_m$. Therefore, all probability statements and the pseudo likelihood have been implicitly conditioned on each observation's market type $\bar{x}_m$.

Given our sampling framework, the observed state vector at the first observation for each market x_{m1} is not exogenous with respect to unobserved market type: $\Pr(x_{m1}|\omega_m) \neq \Pr(x_{m1})$. This is the *initial conditions problem* in the estimation of dynamic discrete models with autocorrelated unobservables (Heckman (1981)). Under the assumption that the time-varying component of x_{m1} is drawn from the stationary distribution induced by the Markov perfect equilibrium, we may implement a computationally tractable solution to this problem. Let $p^*(x|f, P)$ denote the steady-state probability distributions of state variables x under transition probability f and Markov perfect equilibrium P .²¹ Therefore, our pseudo likelihood function is

$$(43) \quad Q_M(\theta, P, f) = \frac{1}{M} \sum_{m=1}^M \ln \left(\sum_{l=1}^L \varphi_l(\bar{x}_m) \left[\prod_{t=1}^T \prod_{i=1}^N \Psi_i(a_{imt}|x_{mt}, \omega^l; P_l, \theta) \right] \times p^*(x_{m1}|f, P_l) \right).$$

Given this pseudo likelihood function, the NPL estimator is defined as in Section 3.4. To obtain consistency, the identification condition in Assumption 5 is suitably modified:

ASSUMPTION 5'(C): *There is a unique $\theta^0 \in \Theta$ such that $\theta^0 = \arg \max_{\theta} Q_0(\theta, P^0, f)$, where*

$$(44) \quad Q_0(P, \theta, f) \equiv E \left(\ln \left(\sum_{l=1}^L \varphi_l(\bar{x}_m) \left[\prod_{t=1}^T \prod_{i=1}^N \Psi_i(a_{imt}|x_{mt}, \omega^l; P_l, \theta) \right] \times p^*(x_{m1}|f, P_l) \right) \right).$$

²¹There is a slight abuse of notation here because $p^*(x|\cdot)$ is the steady-state distribution of the time-varying component of x conditional on the fixed component of x and the other conditioning variables. Also note that the function f stacks all transition probabilities for the time-varying components of x , for all values of the fixed components $\bar{x}$.

We obtain this NPL estimator using an iterative procedure that is similar to the one without unobserved heterogeneity. The main difference is that now we have to calculate the steady-state distributions $p^*(\cdot|f, P_l)$ to deal with the initial conditions problem. However, our pseudo likelihood approach also reduces very significantly the cost of dealing with the initial conditions problem. The reason is that given the probabilities (f, P_l) , the steady-state distributions $p^*(\cdot|f, P_l)$ do not depend on the structural parameters in θ . Therefore, the distributions $p^*(\cdot|f, P_l)$ remain constant during any pseudo maximum likelihood estimation and they are updated only between two pseudo maximum likelihood estimations when we obtain new choice probabilities. This implies a very significant reduction in the computational cost associated with the initial conditions problem. We now describe our algorithm in detail.

At iteration 1, we start with L vectors of players' choice probabilities, one for each market type: $\hat{P}_0 = \{\hat{P}_{0l} : l = 1, 2, \dots, L\}$. Then we perform the following steps:

Step 1: For every market type $l \in \{1, 2, \dots, L\}$, we obtain the steady-state distribution of x_{m1} as the unique solution to the system of linear equations,

$$(45) \quad p^*(x|f, \hat{P}_{0l}) = \sum_{x_0 \in X} f^{\hat{P}_{0l}}(x|x_0) p^*(x_0|f, \hat{P}_{0l}) \quad \text{for any } x \in X,$$

where $f^{\hat{P}_{0l}}(\cdot|\cdot)$ is the transition probability for x induced by the conditional transition probability $f(\cdot|\cdot, \cdot)$ and by the choice probabilities in $\hat{P}_{0l}$. That is,

$$(46) \quad f^{\hat{P}_{0l}}(x|x_0) = \sum_{a \in A} \left(\prod_{i=1}^N \hat{P}_{l,i}^0(a_i|x_0) \right) f(x|x_0, a).$$

Step 2: Given the steady-state probabilities $\{p^*(\cdot|f, \hat{P}_{0l}) : l = 1, 2, \dots, L\}$, construct the pseudo likelihood function $Q_M(\theta, \hat{P}_0)$ and obtain the pseudo maximum likelihood estimator of θ as $\hat{\theta}_1 = \arg \max_{\theta \in \Theta} Q_M(\theta, \hat{P}_0)$.

Step 3: For every market type l , update the vector of players' choice probabilities using the best response probability mapping associated with market type l . That is, $\hat{P}_{1l} = \Psi(\hat{\theta}_1, \hat{P}_{0l}, \omega^l)$.

Step 4: If $\|\hat{P}_1 - \hat{P}_0\|$ is smaller than a predetermined small constant, then stop the iterative procedure and choose $(\hat{\theta}_1, \hat{P}_1)$ as the NPL estimator. Otherwise, replace $\hat{P}_0$ by $\hat{P}_1$ and repeat Steps 1–4.

3.6. One-Step Maximum Likelihood Estimation

The pseudo likelihood approach we are advocating in this paper can be combined with a “one-step” MLE that uses consistent estimates of θ^0 and P^0 from NPL estimators to construct consistent estimates of the likelihood score and the information matrix. The likelihood score evaluated at the consistent estimates $\hat{\theta}$ is

$$(47) \quad \widetilde{\nabla}_{\theta} l \equiv \frac{\partial Q_M(\hat{\theta}, \hat{P})}{\partial \theta} + \frac{\partial \hat{P}(\hat{\theta})'}{\partial \theta} \frac{\partial Q_M(\hat{\theta}, \hat{P})}{\partial P},$$

where $\hat{P} = \Psi(\hat{\theta}, \hat{P})$ and $\partial \hat{P}(\hat{\theta})' / \partial \theta$ can be obtained using Assumption 5 and the implicit function theorem: $\partial \Psi(\hat{\theta}, \hat{P})' / \partial \theta (I - \partial \Psi(\hat{\theta}, \hat{P}) / \partial P')^{-1}$. The information matrix can be estimated as

$$\tilde{I} = \sum_{m=1}^M \left\{ \nabla_{\theta} \hat{s}_m + \frac{\partial \hat{P}(\hat{\theta})'}{\partial \theta} \nabla_P \hat{s}_m \right\} \left\{ \nabla_{\theta} \hat{s}_m + \frac{\partial \hat{P}(\hat{\theta})'}{\partial \theta} \nabla_P \hat{s}_m \right\}',$$

i.e., the sum of the outer products of individual score observations in (47) evaluated at $(\hat{\theta}, \hat{P})$. The one-step MLE is then $\tilde{\theta} = \hat{\theta} - \tilde{I}^{-1} \widetilde{\nabla}_{\theta} l$. A consistent estimate of the population fixed-point pair is immediately available from the NPL estimator: $(\hat{\theta}, \hat{P}) = (\hat{\theta}_{\text{NPL}}, \hat{P}_{\text{NPL}})$. Furthermore, most of the terms needed to compute the estimates of the likelihood score and the information matrix are also available as side products from the calculation of the asymptotic variance of the NPL estimator. The one-step MLE is asymptotically efficient and therefore improves asymptotically on the NPL estimator from which it is derived. Its performance in finite samples as well as the trade-offs between computational simplicity and finite sample precision across pseudo likelihood and one-step maximum likelihood (ML) estimators are topics for further research.²²

4. MONTE CARLO EXPERIMENTS

4.1. Data Generating Process and Simulations

This section presents the results from several Monte Carlo experiments based on a dynamic game of market entry and exit with heterogeneous firms. The model is similar to those in our examples, but we consider a log-linear specification of the variable profit function, $\theta_{RS} \ln(S_{mt}) - \theta_{RN} \ln(1 + \sum_{j \neq i} a_{jmt})$, where θ_{RS} and θ_{RN} are parameters. Therefore, the profit function of an active

²²One might consider using the two-step PML estimator as a starting point for one-step efficient estimation. However, to implement the one-step ML, we would first need to iterate in the best response mapping to find the equilibrium probabilities that satisfy $\hat{P}_{2S} = \Psi(\hat{\theta}_{2S}, \hat{P}_{2S})$, and elements of the score would not be available as side products of two-step PML estimation. Also note that in models with permanent unobserved heterogeneity, where nonparametric estimates of P^0 are not available, the one-step ML estimator can be obtained only from the NPL estimates.

firm i is

$$(48) \quad \tilde{\Pi}_{imt}(1) = \theta_{RS} \ln(S_{mt}) - \theta_{RN} \ln\left(1 + \sum_{j \neq i} a_{jmt}\right) \\ - \theta_{FC,i} - \theta_{EC}(1 - a_{im,t-1}) + \varepsilon_{imt}.$$

The parameters we estimate are the fixed operating costs $\{\theta_{FC,i} : i = 1, 2, \dots, N\}$, the entry cost θ_{EC} , and the parameters in the variable profit function, θ_{RS} and θ_{RN} . The logarithm of market size S_{mt} is discrete and it follows a first order Markov process that is known by the researcher and is homogeneous across markets. We consider a sampling framework in which the same N firms are the potential entrants in M separate markets, i.e., all firms are *global* players.

The following parameters are invariant across the different experiments. The number of potential entrants is $N = 5$. Fixed operating costs are $\theta_{FC,1} = -1.9$, $\theta_{FC,2} = -1.8$, $\theta_{FC,3} = -1.7$, $\theta_{FC,4} = -1.6$, and $\theta_{FC,5} = -1.5$, such that firm 5 is the most efficient and firm 1 is the least efficient. The support of the logarithm of market size has five points: $\{1, 2, 3, 4, 5\}$.²³ The private information shocks $\{\varepsilon_{imt}\}$ are independent and identically distributed extreme value type I, with zero mean and unit dispersion. The parameter θ_{RS} is equal to 1. The discount factor equals 0.95 and is known by the researcher. The space of common knowledge state variables (S_{mt}, a_{t-1}) has $2^5 * 5 = 160$ cells. There is a different vector of choice probabilities for each firm. Therefore, the dimension of the vector of choice probabilities for all firms is $5 * 160 = 800$.

We present results from six experiments. The only difference between the data generating processes in these experiments is in the values of the parameters θ_{EC} and θ_{RN} . Table I presents the values of these parameters for the six experiments. The table also presents two ratios that provide a measure of the magnitude of these parameter values relative to variable profits. The term $\theta_{EC}/(\theta_{RS} \ln(3))$ is the ratio between entry costs and the variable profit of a monopolist in a market of average size, and $100(\theta_{RN} \ln(2))/(\theta_{RS} \ln(3))$ represents the percentage reduction in variable profits when we go from a monopoly to a duopoly in a market of average size.

For each experiment we compute a MPE. The equilibrium is obtained by iterating in the best response probability mapping starting with a 800×1 vector of choice probabilities with all probabilities equal to 0.5, i.e., $P_i(a_i = 1|x) = 0.5$ for every i and x . Given the equilibrium probabilities and the transition prob-

²³The transition probability matrix for market size is

$$\begin{pmatrix} 0.8 & 0.2 & 0.0 & 0.0 & 0.0 \\ 0.2 & 0.6 & 0.2 & 0.0 & 0.0 \\ 0.0 & 0.2 & 0.6 & 0.2 & 0.0 \\ 0.0 & 0.0 & 0.2 & 0.6 & 0.2 \\ 0.0 & 0.0 & 0.0 & 0.2 & 0.8 \end{pmatrix}.$$

The steady-state distribution implied by this transition probability is $\Pr(S = j) = 0.2$ for any $j \in \{1, 2, 3, 4, 5\}$. The average market size is 3.

TABLE I
MONTE CARLO EXPERIMENTS: PARAMETERS θ_{EC} AND θ_{RN}

	Exp. 1	Exp. 2	Exp. 3	Exp. 4	Exp. 5	Exp. 6
Parameter ^a	$\theta_{EC} = 1.0$ $\theta_{RN} = 0.0$	$\theta_{EC} = 1.0$ $\theta_{RN} = 1.0$	$\theta_{EC} = 1.0$ $\theta_{RN} = 2.0$	$\theta_{EC} = 0.0$ $\theta_{RN} = 1.0$	$\theta_{EC} = 2.0$ $\theta_{RN} = 1.0$	$\theta_{EC} = 4.0$ $\theta_{RN} = 1.0$
$\frac{\theta_{EC}}{\theta_{RS} \ln(3)}$	0.91	0.91	0.91	0.00	1.82	3.64
$100 \frac{\theta_{RN} \ln(2)}{\theta_{RS} \ln(3)}$	0.0%	63.1%	126.2%	63.1%	63.1%	63.1%

^aThe parameter $\theta_{EC}/(\theta_{RS} \ln(3))$ represents the ratio between entry costs and the annual variable profit of a monopolist in a market of average size, i.e., $S_{mt} = 3$. The parameter $100(\theta_{RN} \ln(2))/(\theta_{RS} \ln(3))$ represents the percentage reduction in annual variable profits when we go from a monopoly to a duopoly in an average size market, i.e., $S_{mt} = 3$.

abilities for market size, we obtain the steady-state distribution of the state, $p^*(S_t, a_{t-1})$, which is the unique solution to the system of equations

$$(49) \quad p^*(s, a) = \sum_{a_0 \in \{0,1\}^5} \sum_{s_0=1}^5 f_s(S_{t+1} = s | S_t = s_0) \\ \times \Pr(a_t = a | S_t = s_0, a_{t-1} = a_0) p^*(S_0, a_0)$$

for all $(s, a) \in \{1, 2, 3, 4, 5\} \times \{0, 1\}^5$, where $\Pr(a_t | S_t, a_{t-1}) = \prod_{i=1}^N P_i(a_{it} | S_t, a_{t-1})$, where $P_i(\cdot | \cdot, \cdot)$ is the equilibrium probabilities and $f_s(S_{t+1} | S_t)$ is the transition probability function of market size.

Steady-state probabilities and equilibrium choice probabilities are used to generate all the Monte Carlo samples of an experiment. These samples are $\{S_{m1}, a_{m0}, a_{m1} : m = 1, 2, \dots, M\}$. The initial state values $\{S_{m1}, a_{m0} : m = 1, 2, \dots, M\}$ are random draws from the steady-state distribution of these variables. Then, given a draw (S_{m1}, a_{m0}) , a_{m1} is obtained by drawing a single choice for each firm from the equilibrium choice probabilities $P_i(\cdot | S_{m1}, a_{m0})$. Because markets are homogeneous and all firms are global players in this design, the time dimension of the data is not important and firm fixed effects are identified from multiple market observations. We have implemented experiments with sample sizes $M = 200$ and $M = 400$. The results for the two sample sizes are qualitatively very similar and therefore we only report results for $M = 400$. For each experiment we use 1,000 Monte Carlo simulations to approximate the finite sample distribution of the estimators.

Table II presents some descriptive statistics associated with the Markov perfect equilibrium of each experiment. These descriptive statistics were obtained using a large sample of 50,000 markets where the initial values of state variables were drawn from their steady-state distribution. An increase in θ_{RN} reduces firms' profits; therefore it reduces the number of firms in the market and the probability of entry, and it increases the probability of exit. The effect on the number of exits (or entries) is ambiguous and depends on the other pa-

TABLE II
MONTE CARLO EXPERIMENTS: DESCRIPTION OF THE MARKOV PERFECT
EQUILIBRIUM IN THE DGP^a

	Exp. 1	Exp. 2	Exp. 3	Exp. 4	Exp. 5	Exp. 6
	$\theta_{EC} = 1.0$	$\theta_{EC} = 1.0$	$\theta_{EC} = 1.0$	$\theta_{EC} = 0.0$	$\theta_{EC} = 2.0$	$\theta_{EC} = 4.0$
Descriptive Statistics	$\theta_{RN} = 0.0$	$\theta_{RN} = 1.0$	$\theta_{RN} = 2.0$	$\theta_{RN} = 1.0$	$\theta_{RN} = 1.0$	$\theta_{RN} = 1.0$
Number of active firms						
Mean	3.676	2.760	1.979	2.729	2.790	2.801
Std. dev.	1.551	1.661	1.426	1.515	1.777	1.905
AR(1) for number of active firms (autoregressive parameter)	0.744	0.709	0.571	0.529	0.818	0.924
Number of entrants (or exits) per period	0.520	0.702	0.748	0.991	0.463	0.206
Excess turnover ^b (in # of firms per period)	0.326	0.470	0.516	0.868	0.211	0.029
Correlation between entries and exits	-0.015	-0.169	-0.220	-0.225	-0.140	-0.110
Prob. of being active						
Firm 1	0.699	0.496	0.319	0.508	0.487	0.455
Firm 2	0.718	0.527	0.356	0.523	0.521	0.501
Firm 3	0.735	0.548	0.397	0.547	0.556	0.550
Firm 4	0.753	0.581	0.434	0.564	0.592	0.610
Firm 5	0.770	0.607	0.475	0.586	0.632	0.686

^aFor all these experiments, the values of the rest of the parameters are $N = 5$, $\theta_{FC,1} = -1.9$, $\theta_{FC,2} = -1.8$, $\theta_{FC,3} = -1.7$, $\theta_{FC,4} = -1.6$, $\theta_{FC,5} = -1.5$, $\theta_{RS} = 1.0$, $\sigma_e = 1$, and $\beta = 0.95$.

^bExcess turnover is defined as $(\#entrants + \#exits) - \text{abs}(\#Entrants - \#Exits)$.

rameters of the model.²⁴ In Table II, we can see that for larger values of θ_{RN} , we get fewer active firms but more exits and entries. We can also see that in markets with higher entry costs, we have lower turnover and more persistence in the number of firms. Interestingly, increasing the cost of entry has different effects on the heterogeneous potential entrants. That is, it tends to increase the probability of being active of relatively more efficient firms and reduces that probability for the more inefficient firms.

4.2. Results²⁵

For each of these six experiments we have obtained the two-step PML and the NPL estimators under the following choices of the initial vector of prob-

²⁴Notice that the number of exits is equal to the number of active firms times the probability of exit. Although a higher θ_{RN} increases the probability of exit, it also reduces the number of active firms; therefore, its effect on the number of exits is ambiguous.

²⁵The estimation programs that implement these Monte Carlo experiments, as well as those for the empirical application in Section 5, have been written in the Gauss language. These programs can be downloaded from Victor Aguirregabiria's web page at <http://individual.utoronto.ca/vaguirre>.

TABLE III
MONTE CARLO EXPERIMENTS: MEDIAN NUMBER
OF ITERATIONS OF THE NPL ALGORITHM

Experiment	Initial Probabilities		
	Frequencies	Logits	Random
1	8	4	6
2	11	7	9
3	27	19	23
4	16	8	11
5	12	7	9
6	13	9	10

abilities: (i) the true vector of equilibrium probabilities P^0 ; (ii) nonparametric frequency estimates; (iii) logit models, one for each firm, with explanatory variables the logarithm of market size and the indicators of incumbency status for all the firms; and (iv) independent random draws from a Uniform(0, 1) random variable. The first estimator, which we label 2S-True, is the infeasible PML estimator. We use this estimator as a benchmark for comparison with the other estimators. The estimator initiated with logit estimates (which we label 2S-Logit) is not consistent, but has lower variance than the estimator initiated with nonparametric frequency estimates (labeled 2S-Freq) and therefore may have better properties in small samples. The random values for $\hat{P}_0$ represent an extreme case of inconsistent initial estimates of choice probabilities. The two-step estimator based on these initial random draws is called 2S-Random.

Tables III, IV, and V summarize the results from these experiments. Table III presents the median number of iterations it takes the NPL algorithm to obtain a NPL fixed point. Table IV shows the empirical mean and standard deviations of the estimators based on the 1,000 replications. Table V compares the mean squared error (MSE) of the 2S-Freq, 2S-Logit, and NPL estimators by showing the ratio of the MSE of each to the MSE of the benchmark 2S-True estimator.

REMARK 1: The NPL algorithm always converged and, more importantly, it always converged to the same estimates regardless of the value of $\hat{P}_0$ (true, nonparametric, logit, or random) that we used to initialize the procedure. This was the case not only for the 6,000 data sets generated in the six experiments presented here, but also for other similar experiments that we do not report here (e.g., 6,000 data sets with 200 observations). Of course, this may be a consequence of our functional form assumptions (e.g., logit, multiplicative separability of parameters) or of the equilibrium that we have selected for the DGP (e.g., stable equilibrium), but it is encouraging to see that, at least for this par-

TABLE IV
MONTE CARLO EXPERIMENTS: EMPIRICAL MEANS AND EMPIRICAL STANDARD
DEVIATIONS OF ESTIMATORS

Experiment	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	True values	-1.900	1.000	1.000	0.000
	2S-True	-1.915 (0.273)	1.007 (0.152)	1.002 (0.139)	0.002 (0.422)
	2S-Freq	-0.458 (0.289)	0.374 (0.141)	1.135 (0.190)	0.200 (0.364)
	2S-Logit	-1.929 (0.279)	1.006 (0.153)	0.997 (0.138)	-0.009 (0.431)
	NPL	-1.902 (0.279)	1.018 (0.157)	0.994 (0.139)	0.036 (0.439)
2	True values	-1.900	1.000	1.000	1.000
	2S-True	-1.894 (0.212)	1.002 (0.186)	1.007 (0.118)	1.007 (0.583)
	2S-Freq	-0.919 (0.208)	0.351 (0.119)	0.886 (0.123)	0.095 (0.337)
	2S-Logit	-1.920 (0.226)	0.977 (0.197)	1.000 (0.122)	0.915 (0.597)
	NPL	-1.893 (0.232)	1.016 (0.220)	0.998 (0.121)	1.050 (0.681)
3	True values	-1.900	1.000	1.000	2.000
	2S-True	-1.910 (0.183)	1.006 (0.209)	1.000 (0.112)	2.008 (0.783)
	2S-Freq	-1.126 (0.189)	0.286 (0.094)	0.792 (0.107)	0.027 (0.311)
	2S-Logit	-1.919 (0.248)	1.022 (0.305)	0.985 (0.145)	2.070 (1.110)
	NPL	-1.920 (0.232)	0.950 (0.189)	1.007 (0.116)	1.792 (0.667)
4	True values	-1.900	1.000	0.000	1.000
	2S-True	-1.890 (0.516)	1.020 (0.329)	0.001 (0.119)	1.063 (1.345)
	2S-Freq	-0.910 (0.243)	0.337 (0.104)	0.239 (0.113)	0.127 (0.354)
	2S-Logit	-2.070 (0.436)	0.903 (0.262)	0.000 (0.119)	0.571 (1.061)
	NPL	-1.891 (0.482)	1.014 (0.291)	0.001 (0.115)	1.047 (1.186)
5	True values	-1.900	1.000	2.000	1.000
	2S-True	-1.912 (0.178)	1.007 (0.142)	2.008 (0.132)	1.006 (0.359)
	2S-Freq	-0.840 (0.218)	1.379 (0.130)	1.591 (0.143)	0.181 (0.302)
	2S-Logit	-1.921 (0.204)	0.997 (0.167)	2.002 (0.138)	0.971 (0.405)
	NPL	-1.924 (0.203)	1.018 (0.178)	2.000 (0.137)	1.027 (0.435)
6	True values	-1.900	1.000	4.000	1.000
	2S-True	-1.899 (0.206)	1.003 (0.132)	4.050 (0.203)	1.006 (0.238)
	2S-Freq	-0.558 (0.228)	0.332 (0.128)	2.745 (0.211)	0.206 (0.238)
	2S-Logit	-1.895 (0.240)	0.996 (0.147)	4.048 (0.208)	0.992 (0.277)
	NPL	-1.918 (0.239)	1.009 (0.152)	4.044 (0.207)	1.009 (0.285)

ticular class of models, the NPL works even when initial probabilities are random. We obtained the same result when using actual data in the application in Section 5.

REMARK 2: Table II shows that with $\theta_{RN} = 1$, we need a relatively small number of iterations to obtain the NPL estimator. With $\theta_{RN} = 2$, the number of NPL iterations is significantly larger. In general, the algorithm converges faster when we initialize it with the logit estimates.

TABLE V
SQUARE-ROOT MEAN SQUARE ERROR RELATIVE TO THE ONE-STAGE PML WITH TRUE P^0

Experiment	Estimator	Parameters			
		$\theta_{FC,1}$	θ_{RS}	θ_{EC}	θ_{RN}
1	2S-Freq	5.380	4.222	1.676	0.983
	2S-Logit	1.027	1.006	1.002	1.022
	NPL	1.019	1.040	0.996	1.044
2	2S-Freq	4.736	3.553	1.415	1.655
	2S-Logit	1.070	1.066	1.029	1.034
	NPL	1.098	1.188	1.020	1.171
3	2S-Freq	4.347	3.440	2.095	2.549
	2S-Logit	1.357	1.462	1.301	1.419
	NPL	1.268	0.935	1.038	0.892
4	2S-Freq	1.977	2.035	2.228	0.699
	2S-Logit	0.906	0.848	1.000	0.850
	NPL	0.935	0.884	0.969	0.881
5	2S-Freq	6.054	4.459	3.279	2.429
	2S-Logit	1.146	1.176	1.043	1.130
	NPL	1.143	1.250	1.037	1.210
6	2S-Freq	6.591	5.589	6.072	3.487
	2S-Logit	1.162	1.209	1.020	1.166
	NPL	1.158	1.248	1.010	1.197

REMARK 3: The 2S-Freq estimator has a very large bias in all the experiments, although its variance is similar to, and sometimes even smaller than, the variances of NPL and 2S-True estimators. Therefore, it seems that the main limitation of 2S-Freq is not its larger asymptotic variance (relative to NPL), but its large bias in small samples.

REMARK 4: The NPL estimator performs very well relative to the 2S-True estimator both in terms of variance and bias. The square-root MSE of the NPL estimator is never more than 27% larger than that of the 2S-True estimator. In fact, the NPL estimator can have lower MSE than the 2S-True estimator. This was always the case in experiments where the parameter θ_{RN} is relatively large, as in Experiment 3. When strategic interactions are stronger, the NPL estimator, which imposes the equilibrium condition, has better asymptotic and finite sample properties than an estimator that does not impose this restriction, such as the two-step PML.

REMARK 5: The 2S-Logit performs very well for this simple model. In fact, it has bias and variance very similar to the NPL estimator. Only in Experiment 4, with $\theta_{RN} = 2$, do we find very significant gains in terms of lower bias and variance by using NPL instead of the 2S-Logit estimator. Again, the stronger the

strategic interactions, the more important the gains from iterating in the NPL procedure.

REMARK 6: In all the experiments, the most important gains associated with the NPL estimator occur for the entry cost parameter θ_{EC} .

In drawing conclusions about the relative merits of two-step PML and NPL estimators, one final word of caution is warranted. Notice that with 400 observations and a state space of 160 points, the frequency estimator in this example, although consistent, is very imprecise; i.e., most estimates are zeros or ones. Using a smooth nonparametric estimator in the first step may reduce finite sample bias in the second step. Pakes, Ostrovsky, and Berry (2004) do this, and in the second step they also consider an alternative method of moments estimator based on the moment conditions $E(W_{it}[I\{a_{it} = a\} - \Psi_i(a|x_{it}; \theta, P)]) = 0$, where the “instruments” W_{it} do not depend on θ . In a Monte Carlo experiment using a very simple entry model with homogeneous firms, one state variable, and two structural parameters, they find that this estimator performs much better than the two-step PML. On the other hand, Pakes, Ostrovsky, and Berry’s two-step method of moments estimator was first proposed by Hotz and Miller (1993) in models with no strategic interactions. Several Monte Carlo studies, including Hotz, Miller, Sanders, and Smith (1994), have shown that this estimator can still perform poorly when sample sizes are small relative to the size of the state space. Furthermore, the Monte Carlo experiments for single-agent models in Aguirregabiria and Mira (2002) show that, even when smoothing is used in the first step, there can be significant improvements in finite sample properties when we iterate in the NPL procedure.

5. AN APPLICATION

5.1. *Data and Descriptive Evidence*

This section presents an empirical application of a dynamic game of firm entry and exit in local retail markets. The data come from a census of Chilean firms created for tax purposes by the Chilean Servicio de Impuestos Internos (Internal Revenue Service). This census contains the whole population of Chilean firms that pay the sales tax (Impuesto de Ventas y Servicios). The sales tax is mandatory for any firm in Chile regardless of its size, industry, region, etc. The data set has a panel structure; it has annual frequency and covers the years 1994–1999. The variables in the data set at the firm level are (i) firm identification number, (ii) firm industry at the five digit level, (iii) annual sales, discretized in 12 cells, and (iv) the *comuna* (i.e., county) where the firm is located. We combine these data with annual information on population at the level of comunas for every year between 1990 and 2003.

We consider five retail industries and estimate a separate model for each. The industries are restaurants, bookstores, gas stations, shoe shops, and fish

shops. Competition in these retail industries occurs at the local level and we consider comunas as local markets. There are 342 comunas in Chile. To have a sample of independent local markets, we exclude those comunas in the metropolitan areas of the larger towns: Santiago (52 comunas), Valparaiso (9 comunas), Rancagua (17 comunas), Concepcion (11 comunas), Talca (10 comunas), and Temuco (20 comunas). We also exclude 34 comunas with populations larger than 50,000 because it is likely that they have more than one market for some of the industries we consider. Our working sample contains 189 comunas. In 1999, the median population of a comuna in our sample was 10,400, and the first and third quartiles were 5,400 and 17,900, respectively.

Table VI presents descriptive statistics on the structure and the dynamics of these markets. There are some significant differences in the structure of the five industries. The number of restaurants (15 firms per 10,000 people) is

TABLE VI
DESCRIPTIVE STATISTICS: 189 MARKETS; YEARS 1994–1999

Descriptive Statistics	Restaurants	Gas Stations	Bookstores	Shoe Shops	Fish Shops
Number of firms per 10,000 people	14.6	1.0	1.9	0.9	0.7
Markets with					
0 firms	32.2%	58.6%	49.5%	67.1%	74.1%
1 firm	1.3%	15.3%	15.8%	10.8%	9.6%
2 firms	1.2%	7.8%	8.0%	6.7%	5.0%
3 firms	0.5%	5.2%	6.9%	3.8%	3.4%
4 firms	1.2%	4.0%	3.6%	2.7%	2.0%
More than 4 firms	63.5%	9.2%	16.2%	8.9%	5.9%
Herfindahl index (median)	0.169	0.738	0.663	0.702	0.725
Annual revenue per firm (in thousand \$)	17.6	67.7	23.3	67.2	124.8
Regression log(1 + # firms) on log(market size) ^a	0.383 (0.043)	0.133 (0.019)	0.127 (0.024)	0.073 (0.020)	0.062 (0.018)
Regression log(firm size) on log(market size) ^b	−0.019 (0.034)	0.153 (0.082)	−0.066 (0.050)	0.223 (0.081)	0.097 (0.111)
Entry rate (%) ^c	9.8	14.6	19.7	12.8	21.3
Exit rate (%) ^d	9.9	7.4	13.5	10.4	14.5
Survival rate (hazard rate)					
1 year (%) ^e	86.2 (13.8)	89.5 (10.5)	84.0 (16.0)	86.8 (13.2)	79.7 (20.3)
2 years (%)	69.5 (19.5)	88.5 (1.1)	70.0 (16.6)	71.1 (18.2)	58.1 (27.2)
3 years (%)	60.1 (14.9)	84.6 (4.3)	60.0 (14.3)	52.6 (25.1)	44.6 (23.3)

^aMarket size = population. Regression included time dummies. Standard errors are given in parentheses.

^bFirm size = revenue per firm. Regression included time dummies. Standard errors are given in parentheses.

^cEntry rate = entrants / incumbents.

^dExit rate = exits / incumbents.

^eSurvival and hazard rates are calculated using the subsample of new entrants in years 1995 and 1996.

much larger than the number of gas stations, bookstores, fish shops, or shoe shops (between 1 and 2 firms per 10,000 people). Market concentration, measured by the Herfindahl index, and firm size (i.e., revenue per firm) are also smaller in the restaurant industry. Turnover rates are very high in all these retail industries. It is difficult to survive during the first three years after entry. However, survival is more likely in gas stations than in the other industries.

There are at least three factors that could explain why the number of restaurants is much larger than the number of gas stations or bookstores. First, differences in economies of scale are potentially important. The proportion of fixed costs in total operating costs may be smaller for restaurants. Second, differences in entry sunken costs might also be relevant. Although the creation of a new gas station or a new bookstore requires an important investment in industry-specific capital, this type of irreversible investment may be less important for restaurants. Third, strategic interactions could be smaller between restaurants than between other retail businesses. For instance, product differentiation might be more important among restaurants than among gas stations. To analyze how these three factors contribute to explain the differences between these industries, we estimate a model of entry and exit that incorporates these elements.

5.2. Specification

The specification of the profit function is similar to that in our Monte Carlo experiments but with two differences: we assume that firms are homogeneous in their fixed operating costs and we incorporate the variable ω_m that represents time-invariant market characteristics that are common knowledge to the players but are unobservable to us. The profit of a nonactive firm is zero and the profit of an active firm is

$$(50) \quad \tilde{\Pi}_{imt}(1) = \theta_{RS} \ln(S_{mt}) - \theta_{RN} \ln\left(1 + \sum_{j \neq i} a_{jmt}\right) - \theta_{FC} - \theta_{EC}(1 - a_{im,t-1}) + \omega_m + \varepsilon_{imt}.$$

The inclusion of unobserved market heterogeneity has important implications in our estimation results. We found that for some industries the model without heterogeneity provides a significantly negative estimate for the parameter θ_{RN} , i.e., a firm's current expected profit increases with the expected number of active firms in the market. This result is not economically plausible. It may reflect the existence of positive correlation between the expected value of $\ln(1 + \sum_{j \neq i} a_{jmt})$ and some unobserved market characteristics that affect firms' profits. If this unobserved heterogeneity is not accounted for, the estimates of θ_{RN} will be biased downward. As we subsequently show, this conjecture is confirmed for all industries.

We assume that the current payoff of a nonactive firm is zero regardless of its incumbency status. Therefore, we are implicitly normalizing the exit value to zero. A nice feature of this normalization is that the estimate of θ_{EC} is an estimate of the sunken cost, i.e., entry cost minus exit value. However, this normalization is not innocuous for the interpretation of other parameter estimates. In particular, our estimate of θ_{FC} is an estimate of the fixed operating cost plus a term that is zero only if the exit value is zero.

Our measure of market size S_{mt} is the population in comuna m at year t . We assume that the logarithm of market size follows an AR(1) process, where the autoregressive parameter is homogeneous across markets but the mean varies over markets:

$$(51) \quad \ln(S_{mt}) = \eta_m + \rho \ln(S_{m,t-1}) + u_{mt}.$$

We use the method in Tauchen (1986) to discretize this AR(1) process and obtain the transition matrix of the discretized variable with a 10-point support. For the unobserved market effect ω_m , we assume it has a discrete distribution with $L = 21$ points of support and independent of our observed measure of market size. More precisely, the distribution of ω_m is a discretized version of a Normal(0, σ_ω^2) with support points $\{\omega^l : \omega^l = \sigma_\omega c^l, l = 1, 2, \dots, L\}$, where c^l is the expected value of a standard normal random variable between percentile $100((l-1)/L)$ and percentile $100(l/L)$. Let p_l be the percentile $100(l/L)$ of a standard normal such that $p_l = \Phi^{-1}(l/L)$, where $\Phi^{-1}(\cdot)$ is the inverse function of the cumulative distribution function of a standard normal. Then

$$(52) \quad c^l = \frac{-\phi(p_l) + \phi(p_{l-1})}{\Phi(p_l) - \Phi(p_{l-1})} = -(\phi(p_l) - \phi(p_{l-1}))/L,$$

where ϕ and Φ are the probability distribution function and the cumulative distribution function of the standard normal, respectively. By construction, all the probabilities $\varphi_l \equiv \Pr(\omega_m = \omega^l)$ are equal to $1/L$.

Notice that all firms are ex ante identical and, therefore, we consider symmetric Markov perfect equilibria. That is, every incumbent firm has the same probability of exit and every potential entrant has the same probability of entry. Second, a firm's profit depends on the number of competitors but not on the identity of the competitors. Taking into account these two features of the model, it is simple to show that all the information in $\{a_{im,t-1} : i = 1, 2, \dots, N\}$ that is relevant to predict a firm's current and future profits is contained in just two variables: the firm's own incumbency status, $a_{im,t-1}$, and the number of incumbent firms, $n_{m,t-1}$. The number of possible states associated with these two variables is $2N$. Therefore, the size of the full state space including market size S_{mt} is $(20)N$.

The distribution of the private information variables is assumed normal with mean zero and variance σ_ε^2 . Therefore, the log pseudo likelihood function is

$$(53) \quad \sum_{m=1}^M \ln \left(\sum_{l=1}^L \left(\prod_{t=1}^T \prod_{i=1}^{N_m} \Phi \left([2a_{imt} - 1] \left[\tilde{z}_i^P(x_{mt}) \frac{\theta}{\sigma_\varepsilon} + \tilde{\lambda}_i^P(x_{mt}) + \frac{\sigma_\omega}{\sigma_\varepsilon} c^l \right] \right) \right) \right. \\ \left. \times p^*(x_{m1}|f, P_l) \right),$$

where $\theta = (\theta_{RS}, \theta_{RN}, \theta_{FC}, \theta_{EC})'$. This function is globally concave in the parameters $\theta/\sigma_\varepsilon$ and $\sigma_\omega/\sigma_\varepsilon$. This is another nice feature of the pseudo likelihood approach to deal with the initial conditions problem, because the probabilities $p^*(x_{m1}|f, P_l)$ do not depend on the structural parameters $\theta/\sigma_\varepsilon$ and $\sigma_\omega/\sigma_\varepsilon$. We estimate these parameters following the procedure that we described at the end of Section 3.5.

All the parameters in the model are identified even in a myopic (not forward looking) version of this game. The identification of the parameters $\theta_{RN}/\sigma_\varepsilon$ and $\theta_{EC}/\sigma_\varepsilon$ deserves some explanation. The entry cost parameter is essentially identified from the difference between incumbent firms and potential entrants in the probability of being active and from turnover behavior. The strategic interaction parameter $\theta_{RN}/\sigma_\varepsilon$ is identified because entry–exit probabilities in the data vary with the number of incumbent firms, and this variable enters the profit function only through the expected value of $\ln(1 + \sum_{j \neq i} a_{jmt})$ conditional on x_{mt} .

5.3. Estimation Results

The parameters of the AR(1) process for the logarithm of population are estimated by full maximum likelihood using data for the period 1990–2003. The estimate of the autoregressive coefficient is 0.9757 (s.e. = 0.0008). Other estimation methods provide very similar estimates.²⁶ As previously mentioned, we follow Tauchen (1986) to discretize this variable and to obtain the matrix of transition probabilities. Given that the intercept of the AR(1) process is market specific, the discretization and the transition matrix vary over markets.

We treat the number of potential entrants in each market as an estimable parameter and we assume that it varies across markets and industries but is constant over time. Our estimate of the number of potential entrants in market–industry m is

$$(54) \quad N_m = \max \left\{ \max_{t \in \{1, 2, \dots, T\}} \{n_{m,t-1} + en_{mt}\}; 2 \right\},$$

²⁶The within groups (or fixed effects) estimator is 0.9766 (s.e. = 0.0008). The ordinary least squares estimator in first differences is: 0.9739 (s.e. = 0.0032). The instrumental variables estimator in first differences using the population at $t - 2$ as the instrument is 0.9706 (s.e. = 0.0128).

TABLE VII
DISTRIBUTION OF THE ESTIMATED NUMBER OF POTENTIAL ENTRANTS

N_m	Restaurants	Gas Stations	Bookstores	Shoe Shops	Fish Shops
2	63 (33.3%)	146 (77.3%)	123 (65.1%)	153 (81.0%)	158 (83.6%)
3	1 (0.5%)	9 (4.8%)	14 (7.4%)	6 (3.2%)	6 (3.2%)
4	3 (1.6%)	8 (4.2%)	10 (5.3%)	8 (4.2%)	9 (4.8%)
5	1 (0.5%)	8 (4.2%)	5 (2.7%)	5 (2.7%)	2 (1.1%)
6	1 (0.5%)	3 (1.6%)	5 (2.7%)	4 (2.1%)	4 (2.1%)
Maximum	105	17	48	16	20

where $n_{m,t-1}$ is the number of firms active at period $t - 1$, en_{mt} is the number of new entrants at period t , and we assume that there are at least two potential entrants in each market. Table VII presents the distribution of the number of potential entrants for each industry. We have also obtained estimates of the models under two alternative scenarios about the number of potential entrants: (a) the same N for different markets within an industry but different N 's across industries, and (b) the same N for every market and every industry. The qualitative estimation results that we describe subsequently are very similar regardless of which of these three approaches is used to estimate the number of potential entrants.

Table VIII presents NPL estimates of this model for the five industries. The discount factor is fixed at $\beta = 0.95$. As in the case of the Monte Carlo experiments, we initialized the NPL algorithm with different vectors of probabilities and we always converged to the same NPL fixed point. In spite of the parsimonious specification of the model, with only five parameters, the measures of goodness of fit are high. Both for the number of entrants and for the number of exits, the R -squared coefficients are always larger than 0.19. All the parameters have the expected signs. It is important to note that in the estimation of a version of the model without unobserved market characteristics, we obtained much smaller estimates of the parameter θ_{RN} for every industry. In fact, this estimate was negative for the gas station and the shoe shop industries.

As is common in discrete-choice models, the parameters in the profit function are identified only up to scale. Given that the dispersion of the unobservable ε 's may change across industries, we cannot obtain the relative magnitude of fixed costs, entry costs, or strategic interactions by just comparing the values of $\theta_{FC}/\sigma_\varepsilon$, $\theta_{EC}/\sigma_\varepsilon$, or $\theta_{RN}/\sigma_\varepsilon$ for different industries. For this reason, Table IX reports the following normalized coefficients:

- The parameter $\theta_{FC}/(\theta_{RS} \ln(S_{Med}))$ is the ratio between fixed operating costs and the variable profit of a monopolist in a market of median size (i.e., 10,400 consumers).
- The parameter $\theta_{EC}/(\theta_{RS} \ln(S_{Med}))$ is the ratio between sunken entry costs and the variable profit of a monopolist in a market of median size.

TABLE VIII
NPL ESTIMATION OF ENTRY-EXIT MODEL^a

Parameters	Restaurants	Gas Stations	Bookstores	Shoe Shops	Fish Shops
Variable profit:					
$\frac{\theta_{RS}}{\sigma_\varepsilon}$	1.743 (0.045)	1.929 (0.127)	2.029 (0.076)	2.030 (0.121)	0.914 (0.125)
$\frac{\theta_{RN}}{\sigma_\varepsilon}$	1.643 (0.176)	2.818 (0.325)	1.606 (0.201)	2.724 (0.316)	1.395 (0.234)
Fixed operating cost:					
$\frac{\theta_{FC}}{\sigma_\varepsilon}$	9.519 (0.478)	12.769 (1.251)	15.997 (0.141)	14.497 (1.206)	6.270 (1.233)
Entry cost:					
$\frac{\theta_{EC}}{\sigma_\varepsilon}$	5.756 (0.030)	10.441 (0.150)	5.620 (0.081)	5.839 (0.145)	4.586 (0.121)
$\frac{\sigma_\omega}{\sigma_\varepsilon}$	1.322 (0.471)	2.028 (1.247)	1.335 (0.133)	2.060 (1.197)	1.880 (1.231)
Number of observations	945	945	945	945	945
R-squared:					
Entries	0.298	0.196	0.442	0.386	0.363
Exits	0.414	0.218	0.234	0.221	0.298

^aStandard errors are given in parentheses. These standard errors are computed from the formulae in Section 4, which do not account for the error in the estimation of the parameters in the autoregressive process of market size.

- The parameter $100(\theta_{RN} \ln(2))/(\theta_{RS} \ln(S_{\text{Med}}))$ is the percentage reduction in variable profits per firm when we go from a monopoly to a duopoly in a market of median size.

TABLE IX
NORMALIZED PARAMETERS

Parameters ^a	Restaurants	Gas Stations	Bookstores	Shoe Shops	Fish Shops
$\frac{\theta_{FC}}{\theta_{RS} \ln(S_{\text{Med}})}$	0.590	0.716	0.852	0.772	0.742
$\frac{\theta_{EC}}{\theta_{RS} \ln(S_{\text{Med}})}$	0.357	0.585	0.299	0.311	0.542
$100 \frac{\theta_{RN} \ln(2)}{\theta_{RS} \ln(S_{\text{Med}})}$	7.1%	10.9%	5.9%	10.1%	11.4%
$\frac{\sigma_\omega^2}{\theta_{RS}^2 \text{var}(\ln(S)) + \sigma_\omega^2 + 1}$	0.278	0.436	0.235	0.423	0.642

^a $\theta_{FC}/(\theta_{RS} \ln(S_{\text{Med}}))$ is the ratio between fixed operating costs and variable profits of a monopolist in a market of median size. $S_{\text{Med}} = 10,400$ individuals. $\theta_{EC}/(\theta_{RS} \ln(S_{\text{Med}}))$ is the ratio between entry costs and variable profits of a monopolist in a market of median size. $100\theta_{RN} \ln(2)/(\theta_{RS} \ln(S_{\text{Med}}))$ is the percentage reduction in variable profits per firm when we go from a monopoly to a duopoly in a market of median size. $\sigma_\omega^2/(\theta_{RS}^2 \text{var}(\ln(S)) + \sigma_\omega^2 + 1)$ is the proportion of the cross-sectional variability in monopoly profits that is explained by the unobserved market type ω_m . Note that $\text{var}(\ln(S)) = 1.16$.

- The parameter $\sigma_\omega^2 / (\theta_{RS}^2 \text{var}(\ln(S)) + \sigma_\omega^2 + 1)$ is the proportion of the cross-market variability in monopoly profits that is explained by the unobserved market type ω_m .

Fixed operating costs are a very important component of total profits. These costs range between 59% (in restaurants) and 85% (in bookstores) of the variable profit of a monopolist in a median market. The relatively small degree of economies of scale in restaurants seems to be a major factor to explain the large number of firms in the restaurant industry. Sunken entry costs are statistically significant in the five industries. They range between 31% (in shoe shops) and 58% (in gas stations) of monopolist variable profits in a median market. However, it seems that for these retail industries, sunken entry costs are smaller than annual fixed operating costs. Gas stations are the retailers with largest sunken costs. However, the interindustry differences in sunken costs explain little of the differences in the number of firms.

The strategic interaction parameter is statistically significant for all five industries. The normalized coefficient measures the percentage reduction in variable profits when we go from a monopoly to a duopoly in a medium size market. According to this parameter, restaurants and bookstores are the retailers with the smallest strategic interactions. This might be due to product differentiation in these two industries. However, strategic interactions do not seem very strong as measured by the normalized parameter because its value ranges between 5.9% and 11.4%.

In this model there are four sources of cross-sectional variation in average profits across markets: market size, the unobserved variable ω_m , the private information shocks, and the number of incumbent firms, which is endogenous. The coefficient $\sigma_\omega^2 / (\theta_{RS}^2 \text{var}(\ln(S)) + \sigma_\omega^2 + 1)$ measures the contribution of the unobserved market type to the cross-market variability in monopoly profits. This contribution varies over industries, but is always very important. For the case of gas stations, this contribution is just as important as market size (i.e., 44%); for the fish shops industries, it is much more important (i.e., 64%).

Based on these estimations, the main differences between these retail industries can be summarized as follows. First, economies of scale are smaller in the restaurant industry and this is the main factor that explains the large number of restaurants. Second, strategic interactions are particularly small among restaurants and among bookstores, which might be due to more product differentiation in those industries. This also contributes to explain the large number of restaurants. Third, economies of scale seem very important in the bookstore industry. However, the number of bookstores is, in fact, larger than the number of gas stations or the number of shoe shops. The reason is that negative strategic interactions are weak in this industry. Fourth, industry-specific investments, i.e., sunken entry costs, are significant in the five industries. However, these costs are smaller than annual fixed operating costs. Gas stations is the industry with largest sunken costs, but the magnitude of these costs does not result in a particularly small number of firms in this industry. However, it does contribute to explain the lower turnover for gas stations.

6. CONCLUSIONS

This paper presents a class of empirical dynamic discrete games and studies the estimation of structural parameters in these models. We are particularly concerned with two estimation problems: the computational burden in the solution of the game and the problem of multiple equilibria. We proposed two different pseudo maximum likelihood methods that successfully deal with these issues: two-step pseudo maximum likelihood and nested pseudo maximum likelihood (NPL). We argue that the second method has several potential advantages relative to the first. These advantages are illustrated in our Monte Carlo experiments and in an empirical application to a model of firm entry and exit in oligopoly markets. In particular, the NPL estimator has a smaller finite sample bias than the two-step PML estimator. Furthermore, it can be implemented in models with permanent unobserved heterogeneity.

Dept. of Economics, University of Toronto, 150 St. George Street, Toronto, Ontario M5S 3G, Canada; victor.aguirregabiria@utoronto.ca

and

CEMFI, Casado del Alisal 5, 28028 Madrid, Spain; mira@cemfi.es.

Manuscript received March, 2005; final revision received September, 2006.

APPENDIX: PROOFS

PROOF OF THE REPRESENTATION LEMMA: The best response operator Ψ was derived in Section 2.4 to provide an alternative characterization of an equilibrium so that any fixed point of Λ is obviously a fixed point of Ψ . To prove the converse, partition a vector of choice probabilities P as (P_i, P_{-i}) . Given the choice probabilities of other players P_{-i} , player i faces a “game against nature,” which is a single-agent Markov decision process as defined in Aguirregabiria and Mira (2002). Note that Ψ_i (in P_i , with P_{-i} fixed) is that paper’s “policy iteration operator” in the single agent’s game against nature. It follows from Proposition 1(a) and (c) in Aguirregabiria and Mira (2002) that a fixed point of Ψ_i is unique and the “smoothed” value function associated with it is the V_i^P function that solves the integrated Bellman equation in (7) of Section 2.2. Repeating this reasoning for all i establishes that a fixed point of Ψ is a fixed point of Λ . *Q.E.D.*

PROOFS OF PROPOSITIONS 1 AND 2.: To prove the propositions, we consider various attributes.

Uniform Convergence of Q_M to Q_0 . Under Assumption 1, the probabilities $\Psi(\theta, P)$ are bounded away from zero and one for every value of (θ, P) , and this condition implies that $Q_M(\theta, P)$ converges almost surely and uniformly in

$(\theta, P) \in \Theta \times [0, 1]^{N|X|}$ to a continuous nonstochastic function $Q_0(\theta, P)$, where $Q_0(\theta, P) \equiv E(\sum_{t=1}^T \sum_{i=1}^N \ln \Psi_i(a_{imt}|x_{mt}; \theta, P))$. Continuity and differentiability of sample and population criterion functions follow from conditions (i) and (ii) in Propositions 1 and 2.

Consistency of the Two-Step PML Estimator. Notice that (a) $Q_0(\theta, P)$ is uniformly continuous, (b) $Q_M(\theta, P)$ converges almost surely and uniformly in (θ, P) to $Q_0(\theta, P)$, and (c) $\hat{P}^0$ converges almost surely to P^0 . Under (a)–(c), $Q_M(\theta, \hat{P}^0)$ converges almost surely and uniformly in θ to $Q_0(\theta, P^0)$ (Lemma 24.1 in Gourieroux and Monfort (1995)). By the identification Assumption 5(C), θ^0 is the only vector in Θ such that $\Psi(\theta, P^0) = P^0$. Therefore, by the information inequality, $Q_0(\theta, P^0)$ has a unique maximum in Θ at θ^0 . It follows that $\hat{\theta}_{2S} \equiv \arg \max_{\theta \in \Theta} Q_M(\theta, \hat{P}^0)$ converges almost surely to θ^0 (Property 24.2 in Gourieroux and Monfort (1995)).

Consistency of the NPL Estimator. Throughout the proof we use the superscript index c to denote the complement of a set. We use Y_0 to denote the set of population NPL fixed points: $Y_0 \equiv \{(\theta, P) \in \Theta \times [0, 1]^{N|X|} : \theta = \tilde{\theta}_0(P) \text{ and } P = \phi_0(P)\}$.

We begin with an outline of the proof, which proceeds in five steps:

- Step 1. The term (θ^0, P^0) uniquely maximizes $Q_0(\theta, P)$ in the set Y_0 .
- Step 2. With probability approaching 1, every element of Y_M belongs to the union of a set of arbitrarily small open balls around the elements of Y_0 .
- Step 3. The function ϕ_M converges to ϕ_0 in probability uniformly in $P \in N(P^0)$.
- Step 4. With probability approaching 1, there exists an element (θ_M^*, P_M^*) of Y_M in any open ball around (θ^0, P^0) .
- Step 5. With probability approaching 1, the NPL estimator is the element of Y_M that belongs to an open ball around (θ^0, P^0) .

When Y_0 contains only one element, then Steps 1 and 2 prove the consistency of the NPL estimator. When there are multiple NPL fixed points in the population, Steps 3–5 are needed to prove consistency.

STEP 1—The term (θ^0, P^0) uniquely maximizes $Q_0(\theta, P)$ in the set Y_0 : Given the identification Assumption 5, we can show that (θ^0, P^0) is the unique pair that satisfies the conditions (a) $\theta^0 = \arg \max_{\theta \in \Theta} Q_0(\theta, P^0)$, (b) $P^0 = \Psi(\theta^0, P^0)$, and (c) for any (θ, P) that verifies (a) and (b), then $Q_0(\theta^0, P^0) \geq Q_0(\theta, P)$. Conditions (a) and (c) result from the application of Assumption 5(A) and (B) and the Kullback–Leibler information inequality. Condition (b) follows from Assumption 5(B) and uniqueness follows from Assumption 5(C). Note that (θ^0, P^0) may be the only pair that satisfies (a) and (b), i.e., Y_0 may be a singleton.

STEP 2—With probability approaching 1, every element of Y_M belongs to the union of a set of arbitrarily small open balls around the elements of Y_0 : Define the function

$$Q_0^*(\theta, P) \equiv \max_{c \in \Theta} \{Q_0(c, P)\} - Q_0(\theta, P).$$

Because $Q_0(\theta, P)$ is continuous and $\Theta \times [0, 1]^{N|X|}$ is compact, Berge's maximum theorem establishes that $Q_0^*(\theta, P)$ is a continuous function in $\Theta \times [0, 1]^{N|X|}$. By construction, $Q_0^*(\theta, P) \geq 0$ for any $(\theta, P) \in \Theta \times [0, 1]^{N|X|}$. Let E be the set of vectors (θ, P) that are fixed points of the equilibrium mapping Ψ , i.e.,

$$E \equiv \{(\theta, P) \in \Theta \times [0, 1]^{N|X|} : P - \Psi(\theta, P) = 0\}.$$

Given that $\Theta \times [0, 1]^{N|X|}$ is compact and the function $P - \Psi(\theta, P)$ is continuous, then E is a compact set. By definition, the set Y_0 is included in E . For each element of Y_0 , consider an (arbitrarily small) open ball that contains it. Let $\mathfrak{Z}$ be the union of these open balls that contain the elements of Y_0 . Because the sets E and $\mathfrak{Z}^c$ are compact, then $\mathfrak{Z}^c \cap E$ is also compact. Define the constant

$$\varepsilon = \min_{(\theta, P) \in \mathfrak{Z}^c \cap E} Q_0^*(\theta, P).$$

By construction, we have that $\varepsilon > 0$. To see this, note that $\varepsilon = 0$ implies that there is a (θ, P) such that $(\theta, P) \notin Y_0$ but $(\theta, P) \in E$ and $Q_0(\theta, P) = \max_{c \in \Theta} \{Q_0(c, P)\}$, a contradiction. Define the event

$$A_M \equiv \{|Q_M(\theta, P) - Q_0(\theta, P)| < \varepsilon/2 \text{ for all } (\theta, P) \in \Theta \times [0, 1]^{N|X|}\}.$$

Let (θ_M^*, P_M^*) be an element of Y_M . Then we have that (a) A_M implies $Q_0(\theta_M^*, P_M^*) > Q_M(\theta_M^*, P_M^*) - \varepsilon/2$ and (b) for any $\theta \in \Theta$, A_M implies $Q_M(\theta, P_M^*) > Q_0(\theta, P_M^*) - \varepsilon/2$. Furthermore, given that (θ_M^*, P_M^*) is a NPL fixed point, we have that (c) $Q_M(\theta_M^*, P_M^*) \geq Q_M(\theta, P_M^*)$ for any $\theta \in \Theta$. Combining inequalities (a) and (c), we get that A_M implies $Q_0(\theta_M^*, P_M^*) > Q_M(\theta, P_M^*) - \varepsilon/2$ for any $\theta \in \Theta$. Adding up this inequality to (b), we get

$$\begin{aligned} A_M &\Rightarrow \{Q_0(\theta_M^*, P_M^*) > Q_0(\theta, P_M^*) - \varepsilon \text{ for any } \theta \in \Theta\} \\ &\Rightarrow \left\{Q_0(\theta_M^*, P_M^*) > \max_{\theta \in \Theta} \{Q_0(\theta, P_M^*)\} - \varepsilon\right\} \\ &\Rightarrow \{\varepsilon > Q_0^*(\theta_M^*, P_M^*)\} \\ &\Rightarrow \left\{\min_{(\theta, P) \in \mathfrak{Z}^c \cap E} Q_0^*(\theta, P) > Q_0^*(\theta_M^*, P_M^*)\right\} \\ &\Rightarrow \{(\theta_M^*, P_M^*) \in \mathfrak{Z}\} \end{aligned}$$

because $(\theta_M^*, P_M^*) \in E$. Therefore, $\Pr(A_M) \leq \Pr((\theta_M^*, P_M^*) \in \mathfrak{S})$. Thus, $\Pr((\theta_M^*, P_M^*) \in \mathfrak{S})$ converges to 1 as M goes to infinity, i.e., any element of the set Y_M belongs to an open ball around an element of Y_0 with probability approaching 1.

As a corollary, note that if Y_0 is a singleton, consistency follows and there is no need to search for multiple NPL fixed points.

STEP 3—The function ϕ_M converges to ϕ_0 in probability uniformly in $P \in N(P^0)$: First, we show that $\tilde{\theta}_M(P)$ converges to $\tilde{\theta}_0(P)$ in probability uniformly in $P \in N(P^0)$. By condition (vi), $\tilde{\theta}_0(P)$ is an interior singleton for $P \in N(P^0)$. Let $N_\varepsilon(\tilde{\theta}_0(P)) \subset \Theta$ be an open ball around $\tilde{\theta}_0(P)$ with radius $\varepsilon > 0$. By condition (vi) in Proposition 2, for all ε small enough, there is a constant $\delta(\varepsilon) > 0$ such that

$$\max_{\theta \in \Theta \cap N_\varepsilon(\tilde{\theta}_0(P))^c} Q_0(\theta, P) \leq Q_0(\tilde{\theta}_0(P), P) - \delta(\varepsilon)$$

for all P in $N(P^0)$. Define the event

$$A_M \equiv \left\{ |Q_M(\theta, P) - Q_0(\theta, P)| < \delta(\varepsilon)/2 \right. \\ \left. \text{for all } (\theta, P) \in \Theta \times [0, 1]^{N|X|} \right\}.$$

For any $P \in N(P^0)$ we have that (a) A_M implies $Q_0(\tilde{\theta}_M(P), P) > Q_M(\tilde{\theta}_M(P), P) - \delta(\varepsilon)/2$ and (b) A_M implies $Q_M(\tilde{\theta}_0(P), P) > Q_0(\tilde{\theta}_0(P), P) - \delta(\varepsilon)/2$. Furthermore, by definition of $\tilde{\theta}_M(P)$, we have (c) $Q_M(\tilde{\theta}_M(P), P) > Q_M(\tilde{\theta}_0(P), P)$. Combining inequalities (a) and (c), we get that A_M implies $Q_0(\tilde{\theta}_M(P), P) > Q_M(\tilde{\theta}_0(P), P) - \delta(\varepsilon)/2$. Adding up this inequality to (b), we get

$$A_M \Rightarrow \left\{ \text{for any } P \in N(P^0), \text{ we have that} \right. \\ \left. Q_0(\tilde{\theta}_0(P), P) - Q_0(\tilde{\theta}_M(P), P) < \delta(\varepsilon) \right\}.$$

Given the definition of $\delta(\varepsilon)$, the previous expression implies that

$$A_M \Rightarrow \left\{ \text{for any } P \in N(P^0), \text{ we have that} \right. \\ \left. Q_0(\tilde{\theta}_M(P), P) > \max_{\theta \in \Theta \cap N_\varepsilon(\tilde{\theta}_0(P))^c} Q_0(\theta, P) \right\} \\ \Rightarrow \left\{ \text{for any } P \in N(P^0), \text{ we have that } \tilde{\theta}_M(P) \in N_\varepsilon(\tilde{\theta}_0(P)) \right\} \\ \Rightarrow \left\{ \sup_{P \in N(P^0)} \|\tilde{\theta}_M(P) - \tilde{\theta}_0(P)\| < \varepsilon \right\}.$$

Therefore,

$$\Pr(A_M) \leq \Pr\left(\sup_{P \in N(P^0)} \|\tilde{\theta}_M(P) - \tilde{\theta}_0(P)\| < \varepsilon\right).$$

We previously showed that $Q_M(\theta, P)$ converges in probability to $Q_0(\theta, P)$ uniformly in (θ, P) , i.e., $\Pr(A_M)$ goes to 1 as M goes to infinity. Thus, $\Pr(\sup_{P \in N(P^0)} \|\tilde{\theta}_M(P) - \tilde{\theta}_0(P)\| < \varepsilon)$ also converges to 1 as M goes to infinity, i.e., $\tilde{\theta}_M(P)$ converges to $\tilde{\theta}_0(P)$ in probability uniformly in $P \in N(P^0)$. Finally, $\phi_M(P) \equiv \Psi(\tilde{\theta}_M(P), P)$ converges uniformly in probability to $\phi_0(P) \equiv \Psi(\tilde{\theta}_0(P), P)$ because Ψ is uniformly continuous and bounded in the compact space $\Theta \times [0, 1]^{N|X|}$.

STEP 4—With probability approaching 1, there exists an element (θ_M^*, P_M^*) of Y_M in an open ball around (θ^0, P^0) : This part of the proof is constructive and goes beyond the result in Step 2, because it establishes that there exists an element of Y_M in the open ball around a particular population NPL fixed point. Condition (vii) in Proposition 2 is the key sufficient condition here. The proof is based on an argument by Manski (1988).²⁷ Because NPL iteration can be described in terms of P only, with θ as a by-product, we focus here on the choice probability component of the NPL operator. Let $N(P^0)$ be a neighborhood of P^0 that is small enough that P^0 is the only P with $\phi_0(P) = P$. Consider the estimator

$$P_M^* = \arg \min_{P \in N(P^0)} \|\phi_M(P) - P\|^2.$$

By Step 3, we have that $\sup_{P \in N(P^0)} \{\|\phi_M(P) - P\|^2 - \|\phi_0(P) - P\|^2\} \rightarrow_p 0$. Also, $\|\phi_0(P) - P\|^2$ has a unique minimum at P^0 in $N(P^0)$. Therefore, by the usual extremum estimator consistency argument, $P_M^* \rightarrow_p P^0$. Because P_M^* is interior to $N(P^0)$ with probability approaching 1, then the first order conditions

$$2\left(\frac{\partial \phi_M(P_M^*)'}{\partial P} - I\right)(\phi_M(P_M^*) - P_M^*) = 0$$

must be satisfied. By $P_M^* \rightarrow_p P^0$ and the fact that $(\partial \phi_0(P^0)'/\partial P - I)$ is nonsingular, we have that $(\partial \phi_M(P_M^*)'/\partial P - I)$ is nonsingular with probability approaching 1, so, by the first order conditions, $\phi_M(P_M^*) = P_M^*$.

STEP 5—with probability approaching 1, the NPL estimator is the element of Y_M that belongs to an open ball around (θ^0, P^0) (or another one in the same ball): Let $\mathfrak{S}_0$ be an open ball around (θ^0, P^0) and let $\mathfrak{S}_1$ be the union of a set of

²⁷We thank a co-editor for pointing out this argument, which is simpler than our original proof.

open balls around the elements of Y_0 other than (θ^0, P^0) . Define the constant

$$\varepsilon \equiv \left\{ \inf_{(\theta, P) \in \mathfrak{S}_0} Q_0(\theta, P) \right\} - \left\{ \sup_{(\theta, P) \in \mathfrak{S}_1} Q_0(\theta, P) \right\}.$$

Given the continuity of Q_0 , the result of Step 1 above, and the assumption that (θ^0, P^0) is isolated, we can always take small enough neighborhoods $\mathfrak{S}_0$ and $\mathfrak{S}_1$ to guarantee that $\varepsilon > 0$. That is, we can construct a ball around (θ^0, P^0) such that every point in this ball has a higher value of Q_0 than the points in the balls around the other population NPL fixed points. Note that $\mathfrak{S}_0$ and $\mathfrak{S}_1$ are disjoint sets. From the result in Step 4, let (θ_M^*, P_M^*) be the sample NPL fixed point that belongs to $\mathfrak{S}_0$ with probability approaching 1. We show here that with probability approaching 1, the vector (θ_M^*, P_M^*) maximizes Q_M in the set Y_M , i.e., (θ_M^*, P_M^*) is the NPL estimator. Consider the event

$$A_M \equiv \{|Q_M(\theta, P) - Q_0(\theta, P)| < \varepsilon^*/2 \text{ for all } (\theta, P) \in \Theta \times [0, 1]^{N|X|}\},$$

where ε^* is the smallest of the ε 's defined here and in Steps 2–4. Then we have that (a) A_M implies $Q_M(\theta_M^*, P_M^*) > Q_0(\theta_M^*, P_M^*) - \varepsilon^*/2$ and (b) A_M implies $Q_0(\theta, P) > Q_M(\theta, P) - \varepsilon^*/2$ for any (θ, P) . Furthermore, given that (θ_M^*, P_M^*) belongs to $\mathfrak{S}_0$ with probability approaching 1, we have that (c) for any $(\theta, P) \in Y_M \cap \mathfrak{S}_1$, $Q_0(\theta_M^*, P_M^*) \geq Q_0(\theta, P) + \varepsilon^*$. Combining inequalities (a) and (c), we get that A_M implies that $Q_M(\theta_M^*, P_M^*) > Q_0(\theta, P) + \varepsilon^*/2$ for any $(\theta, P) \in Y_M \cap \mathfrak{S}_1$. Adding up this inequality to (b), we get

$$A_M \Rightarrow \{Q_M(\theta_M^*, P_M^*) > Q_M(\theta, P) \text{ for any } (\theta, P) \in Y_M \cap \mathfrak{S}_1\}.$$

By Step 2, with probability approaching 1, Y_M is contained in $\mathfrak{S}_0 \cup \mathfrak{S}_1$. Therefore, with probability approaching 1, the NPL estimator is either (θ_M^*, P_M^*) or another element of $Y_M \cap \mathfrak{S}_0$. Thus, with probability approaching 1 the NPL estimator belongs to any ball around (θ^0, P^0) . The NPL estimator converges in probability to (θ^0, P^0) .

Asymptotic Distribution of the Two-Step PML Estimator. For notational simplicity we consider that $T = 1$, and we omit the time subindex. We use $P_{(a,x)}^0$ to denote the vector of dimension $NJ|X| \times 1$ with the joint distribution of a_m and x_m in the population. The vector $\hat{P}_{(a,x)}^0$ is the sample counterpart of $P_{(a,x)}^0$, i.e., the frequency estimator of $P_{(a,x)}^0$. Using this notation, we can write expectations and sample means in matrix form. For instance,

$$(A.1) \quad E \left(\sum_{i=1}^N \ln \Psi_i(a_{im}|x_m; \theta, P) \right) = \ln \Psi(\theta, P)' P_{(a,x)}^0,$$

$$(1/M) \sum_{m=1}^M \sum_{i=1}^N \ln \Psi_i(a_{im}|x_m; \theta, P) = \ln \Psi(\theta, P)' \hat{P}_{(a,x)}^0.$$

We use also $\nabla_\theta \Psi(\theta, P)$ and $\nabla_P \Psi(\theta, P)$ to denote the Jacobian matrices $\partial \Psi(\theta, P)/\partial \theta'$ and $\partial \Psi(\theta, P)/\partial P'$, respectively.

Let $\nabla_\theta s_m$ and $\nabla_P s_m$ be the pseudo scores (for observation m) evaluated at the true parameter values, i.e., $\nabla_\theta s_m = \sum_{i=1}^N \nabla_\theta \ln \Psi_i(a_{im}|x_m; P^0, \theta^0)$ and $\nabla_P s_m = \sum_{i=1}^N \nabla_P \ln \Psi_i(a_{im}|x_m; P^0, \theta^0)$. Define $\Omega_{\theta\theta} \equiv E(\nabla_\theta s_m \nabla_\theta s_m')$ and $\Omega_{\theta P} \equiv E(\nabla_\theta s_m \nabla_P s_m')$. Given conditions (i) and (ii), the best response mapping $\Psi(\theta, P)$, which defines probability distributions for the discrete choice, is continuously differentiable and choice probabilities are bounded away from zero. Therefore, the regularity condition specified in McFadden and Newey (1994, p. 2164) is satisfied (i.e., the square root of the likelihood is continuously differentiable at the true parameters θ_0, P_0) and the *generalized information matrix equality* holds (see McFadden and Newey (1994, p. 2163)): i.e., we have that $E((q_m - P^0) \nabla_\theta s_m') = 0$ and $E((q_m - P^0) \nabla_P s_m') = I$, where I is the identity matrix. Therefore,

$$(A.2) \quad \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M \nabla_\theta s_m \right) - \Omega_{\theta P} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M (q_m - P^0) \right) \\ \rightarrow_d N(0, \Omega_{\theta\theta} + \Omega_{\theta P} \Sigma \Omega_{\theta P}').$$

The first order conditions that define this estimator are $\nabla_\theta Q_M(\hat{P}^0, \hat{\theta}_{FU}) = 0$. A mean value theorem between (θ^0, P^0) and $(\hat{\theta}_{2S}, \hat{P}^0)$, together with consistency of $(\hat{\theta}_{2S}, \hat{P}^0)$, implies that

$$(A.3) \quad 0 = \nabla_\theta Q_M(P^0, \theta^0) + \nabla_{\theta\theta} Q_M(P^0, \theta^0)(\hat{\theta}_{2S} - \theta^0) \\ + \nabla_{\theta P} Q_M(P^0, \theta^0)(\hat{P}^0 - P^0) + o_p(1).$$

By the central limit theorem and the *information matrix inequality*, we have that $\nabla_{\theta\theta} Q_M(P^0, \theta^0) \rightarrow_p -\Omega_{\theta\theta}$ and $\nabla_{\theta P} Q_M(P^0, \theta^0) \rightarrow_p -\Omega_{\theta P}$. Then

$$(A.4) \quad \sqrt{M}(\hat{\theta}_{2S} - \theta^0) \\ = \Omega_{\theta\theta}^{-1} \left\{ -\Omega_{\theta P} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M (q_m - P^0) \right) + \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M \nabla_\theta s_m \right) \right\} \\ + o_p(M^{-1/2}).$$

By the Mann–Wald theorem, $\sqrt{M}(\hat{\theta}_{2S} - \theta^0)$ converges in distribution to a vector of normal random variables with zero means and variance matrix:

$$(A.5) \quad V_{2S} = \Omega_{\theta\theta}^{-1}(\Omega_{\theta\theta} + \Omega_{\theta P} \Sigma \Omega_{\theta P}') \Omega_{\theta\theta}^{-1} = \Omega_{\theta\theta}^{-1} + \Omega_{\theta\theta}^{-1} \Omega_{\theta P} \Sigma \Omega_{\theta P}' \Omega_{\theta\theta}^{-1}.$$

Asymptotic Distribution of the NPL Estimator. The marginal conditions that define the NPL estimator are

$$(A.6) \quad (1/M) \sum_{m=1}^M \nabla_{\theta} s_m(\hat{\theta}, \hat{P}) = 0,$$

$$\hat{P} - \Psi(\hat{\theta}, \hat{P}) = 0.$$

A stochastic mean value theorem between (θ^0, P^0) and $(\hat{\theta}, \hat{P})$, together with consistency of $(\hat{\theta}, \hat{P})$ implies that

$$(A.7) \quad (1/\sqrt{M}) \sum_{m=1}^M \nabla_{\theta} s_m - \Omega_{\theta\theta} \sqrt{M}(\hat{\theta} - \theta^0) - \Omega_{\theta P} \sqrt{M}(\hat{P} - P^0) = o_p(\sqrt{M}),$$

$$(I - \nabla_P \Psi) \sqrt{M}(\hat{P} - P^0) - \nabla_{\theta} \Psi \sqrt{M}(\hat{\theta} - \theta^0) = o_p(\sqrt{M}).$$

Solving the second set of equations into the first set, we get

$$(A.8) \quad [\Omega_{\theta\theta} + \Omega_{\theta P}(I - \nabla_P \Psi)^{-1} \nabla_{\theta} \Psi] \sqrt{M}(\hat{\theta} - \theta^0)$$

$$= (1/\sqrt{M}) \sum_{m=1}^M \nabla_{\theta} s_m + o_p(\sqrt{M}).$$

By the Mann–Wald theorem, we have that $\sqrt{M}(\hat{\theta} - \theta^0) \rightarrow_d N(0, V_{\text{NPL}})$, where

$$(A.9) \quad V_{\text{NPL}} = [\Omega_{\theta\theta} + \Omega_{\theta P}(I - \nabla_P \Psi)^{-1} \nabla_{\theta} \Psi]^{-1}$$

$$\times \Omega_{\theta\theta} [\Omega_{\theta\theta} + \nabla_{\theta} \Psi' (I - \nabla_P \Psi')^{-1} \Omega'_{\theta P}]^{-1}.$$

Relative Efficiency of NPL and Infeasible Two-Step PML. The asymptotic variance of the infeasible two-step PML is $\Omega_{\theta\theta}^{-1}$. Taking into account that $\Omega_{\theta P} = \nabla_{\theta} \Psi' \text{diag}(P^0)^{-1} \nabla_P \Psi$, we can write the variance of the NPL estimator as

$$(A.10) \quad V_{\text{NPL}} = [(I + \nabla_{\theta} \Psi' S \nabla_{\theta} \Psi \Omega_{\theta\theta}^{-1}) \Omega_{\theta\theta} (I + \Omega_{\theta\theta}^{-1} \nabla_{\theta} \Psi' S' \nabla_{\theta} \Psi)]^{-1},$$

where $S \equiv (I - \nabla_P \Psi')^{-1} \nabla_P \Psi \text{diag}(P^0)^{-1}$. Then $\Omega_{\theta\theta}^{-1} - V_{\text{NPL}}$ is positive definite if

$$(A.11) \quad \Delta = (I + \nabla_{\theta} \Psi' S \nabla_{\theta} \Psi \Omega_{\theta\theta}^{-1}) \Omega_{\theta\theta} (I + \Omega_{\theta\theta}^{-1} \nabla_{\theta} \Psi' S' \nabla_{\theta} \Psi) - \Omega_{\theta\theta}$$

is positive definite. Operating in the previous expression, we can get that

$$(A.12) \quad \Delta = \nabla_{\theta} \Psi' (S + S') \nabla_{\theta} \Psi + (\nabla_{\theta} \Psi' S \nabla_{\theta} \Psi) \Omega_{\theta\theta}^{-1} (\nabla_{\theta} \Psi' S' \nabla_{\theta} \Psi)'$$

It is clear that Δ is positive definite if S is positive definite. Because $\text{diag}(P^0)^{-1}$ is a positive definite diagonal matrix, Δ is positive definite if $(I - \nabla_P \Psi')^{-1} \nabla_P \Psi'$ is positive definite. Finally, a sufficient condition for $(I - \nabla_P \Psi')^{-1} \nabla_P \Psi'$ to be positive definite is that all the eigenvalues of $\nabla_P \Psi'$ are between 0 and 1. Q.E.D.

REFERENCES

- AGUIRREGABIRIA, V. (2004): "Pseudo Maximum Likelihood Estimation of Structural Models Involving Fixed-Point Problems," *Economics Letters*, 84, 335–340. [16]
- AGUIRREGABIRIA, V., AND P. MIRA (2002): "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Markov Decision Models," *Econometrica*, 70, 1519–1543. [3,8,9,19,36,44]
- (2005): "A Genetic Algorithm for the Structural Estimation of Games with Multiple Equilibria," Unpublished Manuscript, Department of Economics, Boston University. [24]
- BAJARI, P., L. BENKARD, AND J. LEVIN (2003): "Estimating Dynamic Models of Imperfect Competition," Unpublished Manuscript, Stanford University. [4]
- BAJARI, P., H. HONG, AND S. RYAN (2004): "Identification and Estimation of Normal Form Games," Unpublished Manuscript, Duke University. [4]
- BERRY, S. (1992): "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 60, 889–917. [1,4]
- BRESNAHAN, T., AND P. REISS (1990): "Entry in Monopoly Markets," *Review of Economic Studies*, 57, 531–553. [1,2,4]
- (1991a): "Empirical Models of Discrete Games," *Journal of Econometrics*, 48, 57–81. [2]
- (1991b): "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 99, 977–1009. [1]
- BROCK, W., AND S. DURLAUF (2001): "Discrete Choice with Social Interactions," *Review of Economics Studies*, 68, 235–260. [1]
- CILIBERTO, F., AND E. TAMER (2006): "Market Structure and Multiple Equilibria in the Airline Industry," Unpublished Manuscript, Department of Economics, University of Virginia. [4]
- DE JUAN, R. (2001): "Entry and Exit in Independent Markets: An Application to the Spanish Retail Banking Market," Working Paper 0105, Fundacion Empresa Publica. [4]
- EINAV, L. (2003): "Not All Rivals Look Alike: An Empirical Model for Discrete Games with Asymmetric Rivals," Working Paper, Graduate School of Business, Stanford University. [1]
- ELLICKSON, P. (2003): "Competition in the Supermarket Industry: Sunk Costs and Market Structure," Unpublished Manuscript, Department of Economics, Duke University. [4]
- ENGERS, M., AND S. STERN (2002): "Long-Term Care and Family Bargaining," *International Economic Review*, 43, 73–114. [1]
- ERICSON, R., AND A. PAKES (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53–82. [24]
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. Cambridge, MA: MIT Press. [8]
- GONG, G., AND F. SAMANIEGO (1981): "Pseudo Maximum Likelihood Estimation: Theory and Applications," *The Annals of Statistics*, 9, 861–869. [3]
- GOURIEROUX, C., AND A. MONFORT (1995): *Statistics and Econometric Models*, Vol. I. Cambridge, U.K.: Cambridge University Press. [45]
- HARSANYI, J. (1973): "Games with Randomly Perturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1–23. [8]
- (1995): "Games with Incomplete Information," *American Economic Review*, 85, 291–303. [7]
- HECKMAN, J. (1978): "Dummy Endogenous Variables in a Simultaneous Equation System," *Econometrica*, 46, 931–959. [2]

- (1981): “The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time–Discrete Data Stochastic Process,” in *Structural Analysis of Discrete Data with Econometric Applications*, ed. by C. Manski and D. McFadden. Cambridge, MA: MIT Press. [27]
- HOTZ, J., AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529. [3,10,36]
- 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, 265–289. [4,36]
- JOVANOVIĆ, B. (1989): “Observable Implications of Models with Multiple Equilibria,” *Econometrica*, 57, 1431–1437. [2]
- KOOREMAN, P. (1994): “Estimation of Econometric Models of Some Discrete Games,” *Journal of Applied Econometrics*, 9, 255–268. [1]
- MANSKI, C. (1988): *Analog Estimation Methods in Econometrics*. New York: Chapman and Hall. [48]
- McFADDEN, D., AND W. NEWEY (1994): “Large Sample Estimation and Hypothesis Testing,” in *Handbook of Econometrics*, Vol. 4, ed. by R. E. Engle and McFadden. Amsterdam: North-Holland. [50]
- MCKELVEY, R., AND A. MCLENNAN (1996): “Computation of Equilibria in Finite Games,” in *Handbook of Econometrics*, Vol. 4, ed. by R. E. Engle and McFadden. Amsterdam: North-Holland. [16]
- MILGROM, P., AND R. WEBER (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics and Operations Research*, 10, 619–632. [8]
- MORO, A. (2003): “The Effect of Statistical Discrimination on Black–White Wage Inequality: Estimating a Model with Multiple Equilibria,” *International Economic Review*, 44, 467–500. [14]
- NETZ, J., AND B. TAYLOR (2002): “Maximum or Minimum Differentiation? Location Patterns or Retail Outlets,” *Review of Economics and Statistics*, 84, 162–175. [4]
- PAKES, A., AND P. MCGUIRE (2001): “Stochastic Algorithms, Symmetric Markov Perfect Equilibrium and the Curse of Dimensionality,” *Econometrica*, 69, 1261–1282. [24]
- PAKES, A., M. OSTROVSKY, AND S. BERRY (2004): “Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples),” Working Paper 10506, NBER. [3,18,36]
- PESENDORFER, M., AND SCHMIDT-DENGLER (2004): “Least Squares Estimators of Dynamic Games,” Working Paper, London School of Economics. [4,22]
- RUST, J. (1994): “Estimation of Dynamic Structural Models, Problems and Prospects: Discrete Decision Processes,” in *Advances in Econometrics. Sixth World Congress*, ed. by C. Sims. Cambridge, U.K.: Cambridge University Press, 119–170. [4]
- (1997): “Using Randomization to Break the Curse of Dimensionality,” *Econometrica*, 65, 487–516. [24]
- SALOP, S. (1979): “Monopolistic Competition with Outside Goods,” *Bell Journal of Economics*, 10, 141–156. [6]
- SEIM, K. (2000): “Geographic Differentiation and Firms’ Entry Decisions: The Video Retail Industry,” Unpublished Manuscript, Graduated School of Business, Stanford. [1,4]
- TAMER, E. (2003): “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria,” *Review of Economic Studies*, 70, 147–165. [2,4]
- TAUCHEN, G. (1986): “Finite State Markov Chain Approximation to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177–181. [39,40]
- TOIVANEN, O., AND M. WATERSON (2000): “Empirical Research on Discrete Choice Game Theory Models of Entry: An Illustration,” *European Economic Review*, 44, 985–992. [1,4]
- ZHANG-FOUTZ, N., AND V. KADIYALI (2003): “Competitive Dynamics in Optimal Release Timing of Motion Pictures,” Unpublished Manuscript, S. C. Johnson Graduate School of Management, Cornell University. [1]