Sequential Estimation of Dynamic Discrete Games

Victor Aguirregabiria*   Pedro Mira*
Boston University   CEMFI, Madrid

May 2004

Abstract

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, though 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 feasible for some models and data; 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 are either not available or very imprecise, e.g., models with permanent unobserved heterogeneity. Finally, we illustrate these methods in Montecarlo 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.

JEL Classification: C13, C35, C63, C73.

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

---

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 and 1991b, Berry, 1992, Toivanen and Waterson, 2000), models of spatial competition (Seim, 2000), release timing of motion pictures (Einav, 2003, and Zhang-Foutz and Kadiyali, 2003), intra-family 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 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 in this literature 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 non-linear 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 unfeasible even for simple models. The most common approach to deal with this problem has been to impose restrictions which 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 has been used by Bresnahan and Reiss (1990, 1991) in the context of empirical games of market entry. These authors consider a specification where a firm's profit depends on the number of firms 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 propose an estimator that maximizes a likelihood for the number of entrants. Though 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).[1]

---

[1] 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 the 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 the 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. While 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. And 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. 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 an increasing interest in the estimation of discrete games during the last years, which has generated several methodological papers on this topic. Pesendorfer and Schmidt (2003) propose a two stage method that is equivalent to our two-step estimator when it is initialized with consistent nonparametric estimates. Pakes, Ostrowsky and Berry (2003) consider the same estimator, and combine it also with simulation methods. We compare the performance of this estimator with the NPL in our Monte Carlo experiments. Bajari, Benkard and Levin (2003) show that the simulation-based estimator in Hotz, Miller, Saunders and Smith (1994) can be applied to estimate dynamic models of imperfect competition with both discrete and continuous decision variables. For the case of static games with complete information, Tamer (2003) presents sufficient conditions for the identification of a two-player model and proposes a pseudo maximum likelihood estimation method. Tamer (2004) extends this approach to static games with N players. Bajari, Hong and Ryan (2004) study also 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 of 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 the one in Rust (1994, pp. 154-158). In order to make some of the discussions less abstract, we consider a model where firms competing 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 (non-metropolitan) 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, or Ellickson (2003) for supermarkets.

2.1   Framework and basic assumptions

Each market is characterized by demand conditions which can change over time (e.g., population, income and age distribution, etc). Let d_t be the vector of demand shifters at period t. There are N firms operating in the market, which we index by $i \in I = \{1, 2, ..., N\}$. 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 $A = \{0, 1, ..., J\}$, 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 which 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 which 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 $x_t \equiv (d_t, x_{1t}, x_{2t}, ..., x_{Nt})$ and $\varepsilon_t \equiv (\varepsilon_{1t}, \varepsilon_{2t}, ..., \varepsilon_{Nt})$ be the vectors of common knowledge and private information variables, respectively. A firm's current profits depend on x_t, on its own private information ε_{it}, and on the vector of firms' current decisions, $a_t \equiv (a_{1t}, a_{2t}, ..., a_{Nt})$. 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:

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

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

ASSUMPTION 1 (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 $\Pi_i(.)$ is a real valued function, and $\varepsilon_{it} \equiv \{\varepsilon_{it}(0), \varepsilon_{it}(1), ..., \varepsilon_{it}(J)\} \in R^{J+1}$ is a vector of real valued random variables.

4

ASSUMPTION 2 (Conditional independence): The transition probability $p(.|.)$ factors as: $p(x_{t+1}, \varepsilon_{t+1} | a_t, x_t, \varepsilon_t) = p_\varepsilon(\varepsilon_{t+1}) \, f(x_{t+1} \mid a_t, x_t)$. That is: (1) given firms' decisions at period t, private information variables do not affect the transition of common knowledge variables; and (2) 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, $g_i(.)$ is a density function which is absolutely continuous with respect to the Lebesgue measure.

ASSUMPTION 4 (Discrete common knowledge variables): Common knowledge variables have a discrete and finite support: $x_t \in X \equiv \{x^1, x^2, ..., x^{|X|}\}$ where $|X|$ is a finite number.

Example (Entry and exit in a local retail market): Suppose the 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., $a_{it} \in \{0, 1\}$. If a supermarket does not operate a store, it gets zero profits. Opening a new store requires a sunk entry cost α_2. If the supermarket operates a store, its profits depend on: (1) the state of the local economy (e.g., population, income, unemployment rate, etc), which is captured by the vector of demand shifters d_t; (2) the store's years of experience in this market, represented by the variable c_{it}; (3) the number of firms operating in this market; (4) the average experience of other firms in this market; and (5) a private information shock ε_{it}. Current profits of an active store are:

$$\tilde{\Pi}_{it} = \alpha_0 + \alpha_1 d_t - \alpha_2(1 - a_{i,t-1}) + \alpha_3 \ c_{it} - \delta_1 \ \ln \left(1 + \sum_{j=1}^{N} a_{jt} \right) - \delta_2 \left(\frac{\sum_{j=1}^{N} a_{jt} c_{jt}}{\sum_{j=1}^{N} a_{jt}} \right) + \varepsilon_{it} \quad (2)$$

where α_0, α_1, α_2, α_3, δ_1 and δ_1 are parameters. In particular, the parameters δ_1 and δ_2 capture the existence of strategic interactions. The set of common knowledge state variables consists of the demographic variables in d_t, the decisions at previous period, and the years of experience of all supermarkets.

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 $\sigma_i : X \times R^{J+1} \to A$. 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,

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

where $I\{.\}$ 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 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,

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

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:

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

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:

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

It is convenient to define value functions integrated over private information variables. Let $V_i^\sigma(x)$ be the integrated value function $\int \tilde{V}_i^\sigma(x, \varepsilon_i) \, g(d\varepsilon_i)$. Based on this definition and equation (5), we can obtain the *integrated Bellman equation*:

$$V_i^\sigma(x) = \int \max_{a_i \in A} \left\{ \pi_i^\sigma(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} V_i^\sigma(x') \, f_i^\sigma(x'|x, a_i) \right\} g_i(d\varepsilon_i) \tag{7}$$

The right hand side of equation (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 the equilibrium behavior of other firms. The following definition characterizes equilibrium strategies of all firms as best responses to one another.

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

$$\sigma_i^*(x, \varepsilon_i) = \arg\max_{a_i \in A} \left\{ \pi_i^{\sigma^*}(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x' \in X} V_i^{\sigma^*}(x') \, f_i^{\sigma^*}(x'|x, a_i) \right\} \tag{8}$$

---

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

Following Milgrom and Weber (1985) we can also represent a MPE in probability space.[3] 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 slightly the notation and use the symbols π_i^P, V_i^P and f_i^P, respectively, to denote these 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:

$$\Lambda_i(a_i|x; P_{-i}) = \int I\left(a_i = \arg\max_{j \in A}\left\{ \pi_i^P(j,x) + \varepsilon_i(j) + \beta \sum_{x' \in X} V_i^P(x') \, f_i^P(x'|x,j) \right\}\right) g_i(\varepsilon_i) \, d\varepsilon_i \tag{9}$$

We call the functions Λ_i *best response probability functions.* Given our assumptions on the distribution of private information, best response probability functions are 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 fixed point problems in (7); and given these value functions we obtain best response probabilities using the right hand-side of equation (9).

2.4   An alternative best response mapping

We now provide an alternative best response mapping (in probability space) which avoids the solution of the N dynamic programming problems in (7). 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 $V_1^{P^*}$, $V_2^{P^*}$, ..., $V_N^{P^*}$ be firms' value functions associated with this equilibrium. Since equilibrium probabilities are best responses, we can rewrite the Bellman equation (7) as,

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

where $f^{P^*}(x'|x)$ is the transition probability of x induced by P^*.[4] 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

---

[3]Milgrom and Weber consider 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.*

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

i. By Proposition 1 in Hotz and Miller (1993), this conditional expectation is a function of a_i and $P_i^*(x)$ only.

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:

$$\left(I - \beta \ F^{P^*}\right) V_i^{P^*} = \sum_{a_i \in A} P_i^*(a_i) * \left[\pi_i^{P^*}(a_i) + e_i^{P^*}(a_i)\right] \tag{11}$$

where I is the identity matrix; F^{P^*} is a matrix with transition probabilities $f^{P^*}(x'|x)$; and $P_i^*(a_i)$, $\pi_i^{P^*}(a_i)$ and $e_i^{P^*}(a_i)$ are vectors of dimension $|X|$ which stack the corresponding state-specific elements. 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, $\Gamma_i(.)$ can be interpreted as a *valuation operator:* that is, $\Gamma_i(x; P)$ is the *expected value of firm i if all firms (including firm i) behave today and in the future according to their choice probabilities in P.* Therefore, we can characterize a MPE as a fixed point of a mapping $\Psi(P) \equiv \{\Psi_i(a_i|x; P)\}$ with

$$\Psi_i(a_i|x; P) = \int I \left(a_i = \arg\max_{j \in A} \left\{ \pi_i^P(j, x) + \varepsilon_i(j) + \beta \sum_{x' \in X} \Gamma_i(x'; P) \ f_i^P(x'|x, j) \right\}\right) g_i(\varepsilon_i) \ d\varepsilon_i \tag{12}$$

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 N dynamic programming problems, whereas to obtain Γ_i and Ψ_i one only has to solve N systems 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 (Entry and exit in a local retail market): Consider the example in section 2.1. Expected current profits if not active are $\pi_i^P(0, x) = 0$ and if active they are

$$\pi_i^P(1, x_t) = \alpha_0 + \alpha_1 \ d_t - \alpha_2 \ (1 - a_{i,t-1}) + \alpha_3 \ c_{it} - \delta_1 \ N_{it}^P - \delta_2 \ C_{it}^P \tag{13}$$

with $N_{it}^P = \sum_{a_{-i} \in \{0,1\}^{N-1}} \Pr\left(a_{-i}|x_t\right) \ln\left(1 + \sum_{j=1}^N a_{-i}(j)\right)$, and $C_{it}^P = \sum_{a_{-i} \in \{0,1\}^{N-1}} \Pr\left(a_{-i}|x_t\right)$ $\left(\sum_{j=1}^N a_{-i}(j) \ c_{jt} / \sum_{j=1}^N a_{-i}(j)\right)$. We can also write this expected profit as $\pi_i^P(1, x_t) = z_{it}^P \ \theta$, where $z_{it}^P = \left\{1, d_t, a_{i,t-1} - 1, c_{it}, N_{it}^P, C_{it}^P\right\}$, and θ is the vector of parameters $(\alpha_0, \alpha_1, \alpha_2, \alpha_3, \delta_1, \delta_2)'$. If the private information shock is normal with zero mean and variance σ^2, then $e_i^P(0, x_t) = 0$ and $e_i^P(1, x_t) = \sigma\phi\left(\Phi^{-1}\left(P_i(1|x_t)\right)\right) / P_i(1|x_t)$, where $\phi(.)$ and $\Phi(.)$ are the density and the cumulative distribution of the standard normal, respectively. The multiplicative separability of the parameters θ in expected profits implies that these parameters are also multiplicative separable in the mapping $\Gamma_i(P)$. That is, we can write $\Gamma_i(P) = Z_i^P \ \theta + \sigma \ \lambda_i^P$, where Z_i^P and λ_i^P are a matrix and a vector

which are obtained by solving a system of linear equations as in (11) in order to collect the infinite sum of z^P and $e^P(1,.)$ terms, respectively, along all possible future paths originating from a given state. Finally, the best response functions Ψ_i have the following form:

$$\Psi_i(1|x_t; P) = \Phi\left(\ (z_{it}^P + \beta\ Z_{it}^P)\frac{\theta}{\sigma} + \beta\ \lambda_{it}^P\ \right) \tag{14}$$

where $Z_{it}^P = \sum_{x' \in X} Z_i^P(x')\nabla f_{it}^P(x')$, $\lambda_{it}^P = \sum_{x' \in X} \lambda_i^P(x')\nabla f_{it}^P(x')$, $\nabla f_{it}^P(x') = f_i^P(x'|x_t, 1) - f_i^P(x'|x_t, 0)$, and $Z_i^P(x')$ is the corresponding row of matrix Z_i^P.

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 IO.

$$Data = \{a_{mt}, x_{mt} : m = 1, 2, ..., M;\ t = 1, 2, ..., T\} \tag{15}$$

where m is the market subindex, and $a_{mt} = (a_{1mt}, a_{2mt}, ..., a_{Nmt})$. 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 who, active or not, are potential entrants in every local market, and some other firms who are potential entrants in only one local market. For instance, in the fast food industry Mac Donald'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 *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 $\theta \in \Theta \subset R^K$. Primitives are continuously differentiable in θ. 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. 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 $P_{mt}^0 = P^0$. (B) Players expect P^0 to be played in future (out of sample) periods. (C) There is a unique $\theta^0 \in \Theta$ such that $P^0 = \Psi(P^0; \theta^0)$ and $P^0 \neq \Psi(P^0; \theta)$ for any $\theta \neq \theta^0$.

Assumption 5A establishes that the data has been generated by only one Markov Perfect equilibrium.[5] Without this assumption, we would need to extend the primitives of the model to include a probability distribution that determines the likelihood with which different equilibria are selected.[6] Assumption 5B is a natural extension which is necessary in order to accommodate dynamic models. Without it, we cannot compute the expected future payoffs of within-sample actions unless we (once again) specify the beliefs of players regarding the probability of switching equilibria in the future. Our assumption avoids the specification of ad-hoc equilibrium selection devices. Assumption 5C is a standard identification condition.

3.2 Maximum likelihood estimation

For some values of the structural parameters the model can have multiple equilibria. Therefore, we have a likelihood correspondence instead of a likelihood function and this makes estimation by maximum likelihood non-standard. To illustrate this issue, let Υ be the set of equilibrium types and suppose that this set is discrete and countable, e.g., $\Upsilon = \{1, 2, 3...\}$. An equilibrium type is a probability function $P^\tau(\theta)$ where $\tau \in \Upsilon$ is the index that represents the type. For any type τ we can define an *equilibrium type-specific* log-likelihood function:

$$l^\tau(\theta) = \frac{1}{M} \sum\nolimits_{m=1}^{M} \sum\nolimits_{t=1}^{T} \sum\nolimits_{i=1}^{N} \ln P_i^\tau(a_{imt} | x_{mt}; \theta) \tag{16}$$

Under Assumption 5 the population probabilities P^0 belong to one and only one equilibrium type. That is, there is a $\tau_0 \in \Upsilon$ and $\theta^0 \in \Theta$ such that $P^0 = P^{\tau_0}(\theta^0)$. If we knew the equilibrium type τ_0, we would maximize $l^{\tau_0}(\theta)$ with respect to θ and obtain the MLE of θ^0. Under standard regularity conditions, this estimator is consistent, asymptotically normal and efficient.

However, we do not know the equilibrium type of P^0, and therefore this MLE is unfeasible. In principle, we could consider an algorithm that searches both for the equilibrium type τ_0 and for the

---

[5] This assumption can be relaxed if we are willing to impose some additional structure on the sampling framework. For instance, suppose there is a finite number of equilibria that are played in the data. The pseudo maximum likelihood methods in this paper can still be applied if the number of observations (markets) of each equilibrium type goes to infinity with sample size, and the researcher knows the type of equilibria that is played in each market and time period, and (if different equilibria are played over time) players do not anticipate the switch from one equilibrium to another.

[6] Moro (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 in order to identify the selected equilibrium from the data.

vector of parameters θ^0. For instance, if we knew and could characterize all the equilibrium types, we would obtain equilibrium type-specific ML estimators: for any $\tau \in \Upsilon$, $\hat{\theta}^\tau = \arg\max_{\theta \in \Theta} l^\tau(\theta)$. Then, we could define the estimator:

$$\hat{\theta} = \hat{\theta}^{\tau^*}, \; where \; \tau^* = \arg\max_{\tau \in \Upsilon} \; l^\tau(\hat{\theta}^\tau). \tag{17}$$

In practice, this estimator can be difficult to implement. First, notice that we need to know all the equilibrium types that the model has on Θ. This is computationally impractical even for relatively simple models. Second, to obtain an *equilibrium type-specific* estimator, say $\hat{\theta}^\tau$, we need an algorithm that guarantees that for different values of θ we always select equilibrium type τ. This can be a very difficult task for some types of equilibria (see McKelvey and McLennan, 1996). And third, 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. Though evaluations of Ψ for different $\theta' s$ can be relatively cheap because we do not have to invert the matrix $(I - \beta F)$ in (11), 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 in models where the dimension of P is relatively large. For instance, that is the case in most models with heterogenous 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. These problems motivate the following pseudo likelihood estimators.

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. We define first the *pseudo likelihood function*:

$$Q_M(\theta, P) = \frac{1}{M} \sum_{m=1}^M \sum_{t=1}^T \sum_{i=1}^N \ln \Psi_i(a_{imt}|x_{mt}; P, \theta) \tag{18}$$

where P is an arbitrary vector of players' choice probabilities. Suppose that we knew the population probabilities P^0, and consider the following PML estimator:

$$\hat{\theta}_U \equiv \arg\max_{\theta \in \Theta} \; Q_M(\theta, P^0) \tag{19}$$

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 unfeasible because P^0 is unknown. Suppose that we can obtain a $\sqrt{M}-$consistent nonparametric estimator of P^0. For instance, if there are not unobservable state variables, we can use a frequency estimator or a kernel method to estimate players' choice probabilities. 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.

PROPOSITION 1: *Suppose that: (1) Assumptions 1 to 5 hold; (2) $\Psi(P, \theta)$ is twice continuously differentiable; (3) Θ is a compact set; (4) $\theta^0 \in int(\Theta)$; and (5) let $\hat{P}^0 = (1/M) \sum_{m=1}^M q_m$ be an estimator of P^0 such that $\sqrt{M}\left(\hat{P}^0 - P^0\right) \longrightarrow_d N(0, \Sigma)$. Then, $\sqrt{M}\left(\hat{\theta}_{2S} - \theta^0\right) \longrightarrow_d N(0, V_{2S})$, where:*

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

and $\Omega_{\theta P} \equiv E(\{\nabla_\theta s_m\}\{\nabla_P s_m\}')$, with ∇_P representing 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 that 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. 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.

However, the estimator has several important limitations. 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 most importantly, for the sample sizes available in actual applications, the nonparametric estimator of P^0 can be extremely imprecise even when the number of players is not too large (e.g., 5 players). This can generate serious finite sample biases in the estimator of structural parameters. We illustrate this problem with several Monte Carlo experiments in Section 4. And third, for some models it is not possible to obtain consistent nonparametric estimates of P^0. That 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 $\hat{P}_0$, NPL generates a sequence of estimators $\{\hat{\theta}_K : K \geq 1\}$ where the $K - stage$ estimator is defined as:

$$\hat{\theta}_K = \arg\max_{\theta \in \Theta} Q_M(\theta, \hat{P}_{K-1}) \tag{20}$$

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

$$\hat{P}_K = \Psi(\hat{\theta}_K, \hat{P}_{K-1}) \tag{21}$$

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 iterating in the equilibrium 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. A *NPL fixed-point* is the limit of this sequence, if it exists.[7] Clearly, a NPL fixed point $(\hat{\theta}_{NPL}, \hat{P}_{NPL})$ has the following two properties: (a) $\hat{\theta}_{NPL}$ maximizes the pseudo likelihood $Q_M(\theta, \hat{P}_{NPL})$ and (b) $\hat{P}_{NPL} = \Psi(\hat{\theta}_{NPL}, \hat{P}_{NPL})$. For any given sample, Brower's theorem guarantees the existence of at least one NPL fixed-point.[8] However, the set of NPL fixed-points may contain more than one pair (θ, P). The NPL estimator is defined as the NPL fixed-point associated with the maximum value of the pseudo likelihood. Proposition 2 establishes the large sample properties of this estimator.

PROPOSITION 2: Suppose that: (1) Assumptions 1 to 5 hold; (2) $\Psi(P, \theta)$ is twice continuously differentiable; (3) Θ is a compact set; (4) $\theta^0 \in int(\Theta)$; and (5) the NPL estimator is the NPL fixed-point with the maximum value of the pseudo likelihood. Then, $\sqrt{M}\left(\hat{\theta}_{NPL} - \theta^0\right) \longrightarrow_d N(0, V_{NPL})$, where:

$$V_{NPL} = \left[\Omega_{\theta\theta} + \Omega_{\theta P}\left(I - \nabla_P\Psi'\right)^{-1}\nabla_\theta\Psi\right]^{-1}\Omega_{\theta\theta}\left[\Omega_{\theta\theta} + \nabla_\theta\Psi'\left(I - \nabla_P\Psi\right)^{-1}\Omega'_{\theta P}\right]^{-1}$$

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 unfeasible PML estimator, i.e., $V_{NPL} < \Omega_{\theta\theta}^{-1} < V_{2S}$.

NPL 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 above. First, under some conditions on the Jacobian matrix $\nabla_P\Psi$, the NPL is asymptotically more efficient than the unfeasible PML and therefore more efficient than any two-step PML estimator, whatever

---

[7]Although we have not proved convergence of the NPL algorithm in general, we have always obtained convergence in our Monte Carlo experiments and applications.

[8]The pseudo score $\partial Q_M(\theta, P)/\partial\theta$ and the equilibrium mapping $P - \Psi(\theta, P)$ are continuous mappings in the compact set of (θ, P).

the initial estimator of P^0 that 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 of section 4. And third, consistency of the NPL estimator does not require that we start the algorithm with a consistent estimator of choice probabilities. If the initial guess $\hat{P}_0$ is a consistent estimator, consistency of a NPL fixed-point is straightforward because all elements of the sequence of estimators $\{\hat{\theta}_K, \hat{P}_K : K \geq 1\}$ obtained from the NPL algorithm are consistent (see Aguirregabiria and Mira 2002 for a recursive proof in the single agent context). If the researcher uses an initial guess which is not consistent, such as a reduced form parametric approximation with unknown probability limit, the NPL estimator will still converge in probability to an NPL fixed-point of the $Q_\infty(\theta, P)$ function, the probability limit of the sample criterion $Q_M(\theta, P)$. If the population function has more than one fixed-point local maximum, a 'poorly behaved' initial guess $\hat{P}_0$ might identify a NPL fixed point which is not (θ^0, P^0). Condition (5) in Proposition 2 rules this out. In practical terms this means that the researcher should initiate the NPL with different P guesses and, if different limits are attained he should choose the one which maximizes the value of the pseudo likelihood. 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. This case is treated in some detail in the Appendix.

4   Monte Carlo experiments

This section presents the results from several Monte Carlo experiments based on a dynamic game of market entry and exit with heterogeneous firms. The specification of the profit function of firm i is:

$$\tilde{\Pi}_{imt} = \alpha_{0i} + \alpha_1 \ d_{mt} - \alpha_2 \ (1 - a_{im,t-1}) - \delta \ \ln \left(1 + \sum_{j=1}^{N} a_{jmt}\right) + \varepsilon_{imt} \tag{22}$$

d_{mt} represents the size of market m at period t (e.g., population in the market). The parameters to estimate are $\{\alpha_{0i} : i = 1, 2, ...N\}$, α_1, α_2 and δ. The parameters α_2 and δ represent entry cost and the magnitude of strategic interactions, respectively, and they are particularly important for the dynamics of market structure in this model. We consider a sampling framework in which the same N firms are the potential entrants over M separate markets. The following primitives are invariant across the different experiments: $N = 5$, $\alpha_{01} = -1.9$, $\alpha_{02} = -1.8$, $\alpha_{03} = -1.7$, $\alpha_{04} = -1.6$, $\alpha_{05} = -1.5$, $\alpha_1 = 1$, $\beta = 0.95$, and $\{\varepsilon_{imt}\}$ are iid extreme value with zero mean and unit dispersion. Also, the variable d_{mt} has a discrete support with 5 points and it follows a first

order Markov process which is homogeneous across markets.[9] For each experiment, we computed one MPE by iterating in the best response probability mapping starting with a vector of choice probabilities $P_i(a_i = 1|x) = 0.5$ for every i and x. We have implemented experiments with sample size $M = 200$ and $M = 400$ markets, but the results are very similar and we report here only results for $M = 400$. For each experiment we use 1000 Monte Carlo simulations to approximate the finite sample distribution of the estimators. The transition probability of the variable d_{mt} is considered as known in these experiments.

Entry costs, α_2, and the magnitude of strategic interactions, δ, vary over the experiments. We consider values of α_2 between 0 and 4 and values of δ between 0 and 2. To give an idea of the magnitude of these values, notice that the expected one-period profit of firm $i = 5$ (i.e., the most efficient firm) is:

$$
\begin{array}{llll}
Under\ monopoly: & \Pr ofit(firm\ 5) = \alpha_{05} + \alpha_1 E(d_{mt}) & = & 1.5 \\
Under\ duopoly: & \Pr ofit(firm\ 5) = \alpha_{05} + \alpha_1 E(d_{mt}) - \delta\ \ln(2) & = & 1.5 - 0.69\ \delta \\
With\ three\ firms: & \Pr ofit(firm\ 5) = \alpha_{05} + \alpha_1 E(d_{mt}) - \delta\ \ln(3) & = & 1.5 - 1.10\ \delta
\end{array}
$$

Therefore, $\delta = 1$ implies that profits of this firm decrease by 54% when we go from a monopoly to a duopoly, and by 73% when we go from monopoly to three firms. With $\delta = 2$, these percentages are 92% and 147%, respectively. An entry cost $\alpha_2 = 1$ implies 67% of firm 5's profit as a monopolist, and 124% of its profit as a duopolist (with $\delta = 1$).

Table 1 presents the values of α_2 and δ in the different experiments, as well as some descriptive statistics associated with the Markov perfect equilibrium of each experiment.[10] An increase in δ reduces firms' profits and 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 parameters of the model.[11] In Table 1, we can see that for larger values of δ 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 heterogenous potential entrants. That is, it

---

[9]The support of d_{mt} is $\{1, 2, 3, 4, 5\}$, and the transition probability matrix is:

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

[10]The descriptive statistics in Table 1 were obtained using a large sample of 50,000 markets where the initial values of state variables were drawn from their steady-state distribution.

[11]Notice that the number of exits is equal to the number of active firms times the probability of exit. While a higher δ increases the probability of exit, it also reduces the number of active firms, and therefore its effect on the number of exits is ambiguous.

tends to increase the probability of being active of relatively more efficient firms, and reduces that probability for the more inefficient firms.

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 probabilities: (1) the true vector of equilibrium probabilities P^0; (2) nonparametric frequency estimates; (3) logit models, one for each firm, with explanatory variables $\{a_{m,t-1}, d_{mt}\}$; and (4) random draws from a $U(0,1)$. The first case is the unfeasible or PML estimator (which we label '2-true') and we will use it as a benchmark for comparison with the other estimators. The estimator initiated with logit estimates ('2-logit') is not consistent but it has lower variance than the estimator initiated with nonparametric frequency estimates ('2-freq') and therefore it can have better properties in small samples. The random values for $\hat{P}_0$ represent an extreme case of inconsistent initial estimates of choice probabilities. Notice that with 400 observations an a state space with 160 points, the frequency estimator is also very imprecise, i.e., most estimates are zeros or ones.

Tables 2, 3 and 4 summarize the results from these experiments. Table 2 presents the median number of iterations it takes the NPL algorithm to obtain a NPL fixed-point. Table 3 shows the empirical mean and standard deviations of the estimators based on the 1000 replications. Table 4 compares the Mean Squared Error (MSE) of the 2-freq, 2-logit and NPL estimators by showing the ratio of the MSE of each of them to the MSE of the 'benchmark' 2-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 property only specific of our functional form assumptions (e.g., logit, multiplicative separability of parameters) or of the equilibria we considered (e.g., stable equilibrium). But it is encouraging to see that, at least for this particular 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 2 shows that with $\delta = 1$, we need a relatively small number of iterations to obtain the NPL estimator. With $\delta = 2$, the number of NPL iterations is significantly larger. In general, the algorithm converges faster when we initialize it with the logit estimates.

Remark 3: The two-freq estimator has a very large bias in all the experiments, though its variance is similar to, and sometimes even smaller than, the variances of NPL and two-true estimators. Therefore, it seems that the main limitation of two-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 two-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 2-true estimator. In fact, the NPL estimator can have lower MSE than the 2-true estimator. This was always the case in experiments where the parameter δ is relatively large, as in experiment 3.

Remark 5: The two-logit performs very well for this simple model. In fact, it has very similar bias and variance as the NPL estimator. Only in experiment 4, with $\delta = 2$, we find very significant gains in term of lower bias and variance of using NPL instead of two-logit estimator. In general, the stronger the strategic interactions the more important the gains of 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, α_2

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 paying sales taxes. Sales taxes are 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 to 1999. The variables in the data set at the firm level are: (1) firm identification number; (2) firm industry at the 5 digit level; (3) annual sales, discretized in twelve cells; and (4) the *comuna* (i.e., county) where the firm is located. We combine these data with information on population from the 1992 and 2002 censuses at the level of comunas.

We consider five retail industries and estimate a separate model for each of them. 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. In order 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 comunas with populations larger than $50,000$ because it is likely that they have more than one market (34 comunas). 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 5 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 (20 firms per 10,000 people) is much larger than the number of gas stations, bookstores or shoe shops (between 1 and 4 firms per 10,000 people). Market concentration, measured by the Herfindahl index, is smaller in the restaurant industry. Firm size (i.e., revenue per firm) is almost four times larger in gas stations than in restaurants. 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 sunk costs might also be relevant. While 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. And 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 profit function if the firm operates in the market is:

$$\tilde{\Pi}_{imt}(1) = \alpha_0 + \alpha_1 \ln(POP_{mt}) + \alpha_2 (1 - a_{im,t-1}) - \delta \ln\left(1 + \sum_{j \neq i} a_{jmt}\right) + \omega_m + \varepsilon_{imt} \quad (23)$$

POP_{mt} is the population in market m and year t. The variable ω_m represents time-invariant market characteristics that are common knowledge to the players but are unobservable to us. Appendix B describes the NPL estimator for this model with unobserved time-invariant market characteristics. We assume that ω_m is i.i.d. over markets $N(0, \sigma_\omega^2)$.[12] The economic interpretation of the parameters is the following: $-\alpha_0$ is a fixed operating cost; α_2 is an entry sunk cost; α_1 measures how the variable profit of a monopolist increases with market size; and δ captures the effect of the number of competitors on a firm's profit, i.e., strategic interactions. We assume that the

---

[12] In fact, both $\ln(POP_{mt})$ and ω_m have discrete distributions. For ω_m we consider a discretized version of a normal distribution with zero mean. The support is symmetric around zero with 21 points. We use the method in Tauchen (1986) to discretize the AR(1) process for $\ln(POP_{mt})$. We consider 10 points in the support of $\ln(POP_{mt})$.

logarithm of POP_{mt} follows an AR(1) process where the autoregressive parameter is homogeneous across markets but the mean varies over markets:

$$\ln(POP_{mt}) = \eta_m + \rho \ln(POP_{m,t-1}) + u_{mt} \tag{24}$$

The vector of state variables in this model includes the incumbent indicator of each firm at the beginning of the year (i.e., $a_{im,t-1}$ for $i = 1, 2, ..., N$). The number of states associated with these state variables is 2^N, which is intractable. However, the structure of this model is such that we can reduce the number of states to $2N$. First, 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. And 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 the all the information in $\{a_{im,t-1} : i = 1, 2, ..., N\}$ that is relevant to predict a firms' current and future profits is contained in just two variables: the firm's own incumbent status, $a_{im,t-1}$, and the number of incumbent firms, $n_{m,t-1}$. The number of states associated with these two variables is $2N$.

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.[13] To obtain the matrices of transition probabilities associated with a discretization of these AR(1) processes we use the method in Tauchen (1986).

The number of potential entrants is constant over time but varies over markets and industries. We consider the number of potential entrants in each market as parameters to estimate. Our estimate of the number of potential entrants in market m is:

$$N_m = \max\left\{ \max_{t \in (1,2,...,T)} \{n_{m,t-1} + en_{mt}\} \; ; \; 2 \right\} \tag{25}$$

where $n_{m,t-1}$ is the number of firms active in market m 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. It is straightforward to show that this estimator is consistent as $T \rightarrow \infty$. Table 6 presents the distribution of the number of potential entrants for each industry.

---

[13] The within-groups (or fixed effects) estimator is 0.9766 (s.e. = 0.0008). OLS in first differences: 0.9739 (s.e. = 0.0032). And the IV in first differences using population at $t-2$ as instrument is 0.9706 (s.e. = 0.0128).

Table 7 presents NPL estimates of this model for the five industries.[14] 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-square coefficients are always larger than 0.19. All the parameters have the expected signs. It is important to notice that in the estimation of a version of the model without unobserved market characteristics we obtained negative estimates of δ in the gas station and the shoe shop industries.

As it 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 $\varepsilon's$ may change across industries, we cannot obtain the relative magnitude of fixed costs, entry costs or strategic interactions by just comparing the values of α_0/σ, α_2/σ or δ/σ for different industries. For this reason, we also report three normalized coefficients at the bottom of Table 7. The coefficient $exp\left(-\alpha_0/\alpha_1\right)$ represents the minimum market population such that variable profits of a monopolist can cover fixed operating costs. We can see that fixed operating costs, relative to variable profits, are smaller in restaurants than in the other four industries. This can be a major factor to explain the relatively large number of firms in the restaurant industry. Bookstores are the retailers with the larger proportion of fixed costs in total operating costs. The coefficient $exp\left(\alpha_2/\alpha_1\right)$ represents the minimum market population such that variable profits of a monopolist can cover entry sunk costs. The estimates of this coefficient are significantly smaller than for the coefficient $exp\left(-\alpha_0/\alpha_1\right)$ associated with fixed operating costs. It seems that for these retail industries sunk entry costs are small relative to fixed operating costs. Gas stations are the retailers with larger sunk costs. However, the inter-industry differences in sunk costs explain little of the differences in the number of firms. The importance of strategic interactions can be measured by the coefficient $\delta\ln(2)/\alpha_1$. It represents the percentage increase in market population such that profits of a duopolist in the larger market are equal to profits of a monopolist in the smaller market. According to this coefficient, restaurants and bookstores are the retailers with the smallest strategic interactions. This might be due to product differentiation in these two industries.

Based on these estimations the main differences between these retail industries can be summarized as follows. First, economies of scale are very small in the restaurant industry, and this is a main factor to explain the relatively large number of restaurants. Second, strategic interactions are particularly small among restaurants and among bookstores, what might be due to more product differentiation in these industries. Third, economies of scale are particularly important in the bookstore industry. However, the number of bookstores is in fact larger than the number of

---

[14]The discount factor is fixed at $\beta = 0.95$. As in the case of the Monte Carlo experiments, we initialize the NPL algorithm with different vectors of probabilities and we always converged to the same NPL fixed point.

gas stations or the number of shoe shops. The reason is that negative strategic interactions are very small in this industry. And fourth, industry specific investments, i.e., sunk costs, are small in all these industries. Gas stations is the industry with larger sunk costs, but the magnitude of these costs does not explain very much of the small number of firms in this industry. However, it contributes to explain the higher survival probabilities 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 (PML) methods that deal with these issues: two-step PML and nested PML. 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 a empirical application. In particular, the two-step PML tend to have a larger finite sample bias than the NPL.

APPENDIX A: PROOFS OF PROPOSITIONS

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

$$
\begin{aligned}
E\left(\sum_{i=1}^N \ln \Psi_i(a_{im}|x_m; \theta, P)\right) &= \ln \Psi(\theta, P)' P^0_{(a,x)} \\
(1/M) \sum_{m=1}^M \sum_{i=1}^N \ln \Psi_i(a_{im}|x_m; \theta, P) &= \ln \Psi(\theta, P)' \hat{P}^0_{(a,x)}
\end{aligned}
\tag{A.1}
$$

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

Proof of Proposition 1:

Consistency of two-step PML: Define $Q_\infty(P, \theta) \equiv E(\sum_i \ln \Psi_i(a_{im}|x_m; P, \theta))$. Notice that: (a) $Q_\infty(P, \theta)$ is uniformly continuous; (b) $Q_M(P, \theta)$ converges a.s. and uniformly in (P, θ) to $Q_\infty(P, \theta)$; and (c) $\hat{P}^0$ converges a.s. to P^0. Under (a)-(c), $Q_M(\hat{P}^0, \theta)$ converges a.s. and uniformly in θ to $Q_\infty(P^0, \theta)$ (Lemma 24.1 in Gourieroux and Monfort). By the identification assumption 5C, θ^0 is the only vector in Θ such that $\Psi(\theta, P^0) = P^0$ Therefore, by the information inequality $Q_\infty(P^0, \theta)$ has a unique maximum in Θ at θ^0. It follows that $\hat{\theta}_{2S} \equiv \arg\max_{\theta \in \Theta} Q_M(\hat{P}^0, \theta)$ converges a.s. to θ^0 (Property 24.2 in Gourieroux and Monfort).

Asymptotic distribution of two-step PML: 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 $\Omega_{\theta P} \equiv E(\nabla_\theta s_m \nabla_P s'_m)$. By the *generalized information matrix inequality* (see McFadden and Newey, 1994, p. 2163) 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,

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

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:

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

By the CLT and the *information matrix inequality*, we have that $\nabla_{\theta\theta} Q_M(P^0, \theta^0) \to_p -\Omega_{\theta\theta}$, and

$\nabla_{\theta P} Q_M(P^0, \theta^0) \to_p -\Omega_{\theta P}$. Then,

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

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

$$V_{2S} = \Omega_{\theta\theta}^{-1}\left(\Omega_{\theta\theta} + \Omega_{\theta P}\ \Sigma\ \Omega'_{\theta P}\right)\Omega_{\theta\theta}^{-1} = \Omega_{\theta\theta}^{-1} + \Omega_{\theta\theta}^{-1}\ \Omega_{\theta P}\ \Sigma\ \Omega'_{\theta P}\ \Omega_{\theta\theta}^{-1} \tag{A.5}$$

Proof of Proposition 2:

Consistency of the NPL: The NPL estimator $(\hat{\theta}_{NPL}, \hat{P}_{NPL})$ is defined by conditions: (a) $\hat{\theta}_{NPL}$ maximizes in $\theta \in \Theta$ the pseudo likelihood $Q_M(\theta, \hat{P}_{NPL})$; (b) $\hat{P}_{NPL} = \Psi(\hat{\theta}_{NPL}, \hat{P}_{NPL})$; and (c) for any (θ, P) satisfying conditions (a) and (b), $Q_M(\hat{\theta}_{NPL}, \hat{P}_{NPL}) \geq Q_M(\theta, P)$. Since $Q_\infty(\theta, P)$ is uniformly continuous and $Q_M(\theta, P)$ converges a.s. and uniformly in (θ, P) to $Q_\infty(\theta, P)$, we have that $(\hat{\theta}_{NPL}, \hat{P}_{NPL})$ converges a.s. to a point (θ^*, P^*) such that: (a') θ^* maximizes in $\theta \in \Theta$ the population pseudo likelihood $Q_\infty(\theta, P^*)$; (b') $P^* = \Psi(\theta^*, P^*)$; and (c') for any (θ, P) satisfying conditions (a') and (b'), $Q_\infty(\theta^*, P^*) \geq Q_\infty(\theta, P)$. Now, we show that this point (θ^*, P^*) should be (θ^0, P^0). By the identification assumption 5C, (θ^0, P^0) satisfies conditions (a') and (b'). That is, $P^0 = \Psi(\theta^0, P^0)$ and by the Kullback-Leibler information inequality we have that for any $\theta \in \Theta$, $Q_\infty(\theta, P^0) \leq Q_\infty(P^0, \theta^0)$. For any other point (θ, P) that satisfies conditions (a') and (b'), we have that $P \neq P^0$ (again by assumption 5C). Therefore, Kullback-Leibler information inequality implies that for any $(\theta, P) \neq (\theta^0, P^0)$ satisfying conditions (a') and (b'), we have that $Q_\infty(\theta, P) < Q_\infty(P^0, \theta^0)$. We conclude that (θ^0, P^0) is the only pair that satisfies conditions (a'), (b') and (c'), and therefore the NPL estimator converges a.s. to θ^0.

Asymptotic distribution of the NPL: The marginal conditions that define the NPL estimator are:

$$\begin{aligned}(1/M)\sum_{m=1}^M \nabla_\theta s_m(\hat{P}, \hat{\theta}) &= 0 \\ \hat{P} - \Psi(\hat{P}, \hat{\theta}) &= 0\end{aligned} \tag{A.6}$$

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:

$$\begin{aligned}(1/\sqrt{M})\sum_{m=1}^M \nabla_\theta s_m - \Omega_{\theta\theta}\ \sqrt{M}\left(\hat{\theta} - \theta^0\right) - \Omega_{\theta P}\ \sqrt{M}\left(\hat{P} - P^0\right) &= o_p(\sqrt{M}) \\ (I - \nabla_P \Psi)\ \sqrt{M}\left(\hat{P} - P^0\right) - \nabla_\theta \Psi\ \sqrt{M}\left(\hat{\theta} - \theta^0\right) &= o_p(\sqrt{M})\end{aligned} \tag{A.7}$$

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

$$\left[\Omega_{\theta\theta} + \Omega_{\theta P}\left(I - \nabla_P \Psi\right)^{-1}\nabla_\theta \Psi\right]\sqrt{M}\left(\hat{\theta} - \theta^0\right) = (1/\sqrt{M})\sum_{m=1}^M \nabla_\theta s_m + o_p(\sqrt{M}) \tag{A.8}$$

By Mann-Wald theorem, we have that $\sqrt{M}\left(\hat{\theta} - \theta^0\right) \longrightarrow_d N\left(0, V_{NPL}\right)$ where:

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

Relative efficiency of NPL and Unrestricted PML: The asymptotic variance of $\hat{\theta}_U$ is $\Omega_{\theta\theta}^{-1}$. Taking into account that $\Omega_{\theta P} = \nabla_\theta\Psi'diag(P^0)^{-1}\nabla_P\Psi$, we can write the variance of the NPL estimator as:

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

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

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

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

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

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

APPENDIX B: MODEL WITH PERMANENT UNOBSERVED HETEROGENEITY

Let x_{mt} be the observable state variables, and suppose that there is also a vector of time invariant common knowledge unobservable market characteristics, ω_m. For instance, in the entry-exit model, we may have a profit function:

$$\tilde{\Pi}_{imt}(1) = \alpha_0 + \alpha_1 \ d_{mt} + \alpha_2 \ (1 - a_{im,t-1}) - \delta \ \ln\left(1 + \sum_{j=1}^{N} a_{jmt}\right) + \omega_m + \varepsilon_{imt} \tag{B.1}$$

where ω_m represents time-invariant market characteristics affecting firms' profits, which are common knowledge to the players but are unobservable to the econometrician. We make the following assumptions on the distribution of these unobservables.

ASSUMPTION 6: The vector of unobservable variables ω_m is such that: (A) it has a discrete and finite support $\Omega = \left\{\omega^1, \omega^2, \ldots, \omega^L\right\}$; (B) it is independently and identically distributed over markets with probability mass function $\varphi_l \equiv \Pr(\omega_m = \omega^l)$; and (C) 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 6C states that all markets are homogenous with respect to (exogenous) transitions, and it implies that the transition probability functions f can still be estimated from transition data without solving the model. The other parameters θ to be estimated now include the support and the distribution of the unobservables ω.

The vector P now stacks the distributions of players' actions conditional on all values of observable and unobservable common knowledge state variables. P_l is the subvector describing the equilibrium in a market with unobservable ω^l (i.e., a 'type l 'market). We adapt assumptions 5AB on the *data generating process* as follows:

ASSUMPTION 5': Let $P_{mt}^0 \equiv \left\{\Pr(a_{mt} = a|x_{mt} = x, \omega_m = \omega) : (a, x, \omega) \in A^N \times X \times \Omega\right\}$ 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 $P_{mt}^0 = P^0$. (B) Players expect P^0 to be played in future (out of sample) periods.

Assumption 5 still states that only one equilibrium is played in the data *conditional on market type*, which is unobservable to the econometrician but not to players. Now, to obtain the pseudo likelihood function we integrate the best response probabilities over the distribution of unobservable market characteristics. We have that:

$$\ln\Pr(Data|\theta, P) = \sum_{m=1}^{M} \ln\Pr(\tilde{a}_m, \tilde{x}_m|\theta, P) = \sum_{m=1}^{M} \ln\left(\sum_{l=1}^{L} \varphi_l \ \Pr(\tilde{a}_m, \tilde{x}_m|\omega^l; \theta, P)\right) \tag{B.2}$$

where $\tilde{a}_m = \{a_{mt} : t = 1, 2, ..., T\}$ and $\tilde{x}_m = \{x_{mt} : t = 1, 2, ..., T\}$. Applying the chain rule, the

Markov structure of the model, and assumption 6C, we get:

$$\Pr(\tilde{a}_m, \tilde{x}_m | \omega^l; \theta, P) = \left(\prod_{t=1}^{T} \Pr(a_{mt} | x_{mt}, \omega^l; \theta, P) \right) \left(\prod_{t=2}^{T} \Pr(x_{mt} | a_{m,t-1}, x_{m,t-1}, \omega^l; \theta, P) \right) \Pr(x_{m1} | \omega^l; \theta, P)$$

$$= \left(\prod_{t=1}^{T} \prod_{i=1}^{N} \Psi_i(a_{imt} | x_{mt}, \omega^l; P_l, \theta) \right) \left(\prod_{t=2}^{T} f(x_{mt} | a_{m,t-1}, x_{m,t-1}; \theta) \right) \Pr(x_{m1} | \omega^l; \theta, P) \tag{B.3}$$

Therefore,

$$\ln \Pr(Data | \theta, P) = \sum_{m=1}^{M} \ln \left(\sum_{l=1}^{L} \varphi_l \left(\prod_{t=1}^{T} \prod_{i=1}^{N} \Psi_i(a_{imt} | x_{mt}, \omega^l; P_l, \theta) \right) \Pr(x_{m1} | \omega^l; \theta, P) \right)$$

$$+ \sum_{m=1}^{M} \sum_{t=2}^{T} \ln f(x_{mt} | a_{m,t-1}, x_{m,t-1}; \theta) \tag{B.4}$$

The first component in 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 we have mentioned above, the transition probability functions f can still be estimated from transition data without solving the model.

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, so called, *initial conditions problem* in the estimation of dynamic discrete models with autocorrelated unobservables (Heckman, 1981). Under the assumption that x_{m1} is drawn from the stationary distribution induced by the Markov perfect equilibrium, we may implement a computationally tractable solution of this problem. Let $p^*(x; f, P)$ be the steady state probability of state x under transition probability f and Markov perfect equilibrium P. Therefore, our pseudo likelihood function is:

$$Q_M(\theta, P, f) = \frac{1}{M} \sum_{m=1}^{M} \ln \left(\sum_{l=1}^{L} \varphi_l \left(\prod_{t=1}^{T} \prod_{i=1}^{N} \Psi_i(a_{imt} | x_{mt}, \omega^l; P_l, \theta) \right) p^*(x_{m1}; f, P_l) \right) \tag{B.5}$$

Given this pseudo likelihood function, the NPL estimator is defined as in Section 3.4. In order 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 $\theta^0 = \text{argmax}_\theta Q_\infty(\theta, P^0, f)$ where

$$Q_\infty(P, \theta, f) \equiv E \left(\ln \left(\sum_{l=1}^{L} \varphi_l \left(\prod_{t=1}^{T} \prod_{i=1}^{N} \Psi_i(a_{imt} | x_{mt}, \omega^l; P_l, \theta) \right) p^*(x_{m1}; f, P_l) \right) \right)$$

Notice that to start the NPL algorithm we need guesses of all conditional choice probability vectors $\{P_l : l = 1, \ldots, L\}$, all of which will be updated at each iteration: $\hat{P}_{l,K} = \Psi(\hat{\theta}_K, \hat{P}_{l,K-1})$. At each NPL iteration we also need to re-calculate the steady state distributions $p^*(.; f, P_l)$. However,

these steady state probabilities are fixed within an NPL iteration, which facilitates very much the estimation of this model with permanent unobserved heterogeneity.

REFERENCES:

Aguirregabiria, V., and P. Mira (2002): "Swapping the nested fixed point algorithm: A class of estimators for Markov decision models," *Econometrica,* 70(4), 1519-43.

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

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

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

Bresnahan, T., and P. Reiss (1991b): "Entry and competition in concentrated markets", *Journal of Political Economy,* 99, 977-1009.

Pakes, A., M. Ostrowsky, and S. Berry (2003): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," manuscript, Harvard University.

Bajari, P., L. Benkard and J. Levin (2003): "Estimating dynamic models of imperfect competition," manuscript, Stanford University.

Bajari, P., H. Hong, and S. Ryan (2004): "Identification and estimation of normal form games," manuscript, Duke University.

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

De Juan, R. (2001): "Entry and Exit in Independent markets: An application to the Spanish retail banking market," Fundacion Empresa Publica, Working Paper 0105.

Engers, M., and S. Stern (2002): "Long-Term Care and Family Bargaining," *International Economic Review*, 43, 73-114

Einav, L. (2003): "Not All Rivals Look Alike: An Empirical Model for Discrete Games with Asymmetric Rivals," working paper, GSB, Stanford University.

Ellickson, P. (2003): "Competition in the Supermarket Industry: Sunk Costs and Market Structure," manuscript, Department of Economics, Duke University.

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

Gong, G. and F. Samaniego (1981): "Pseudo Maximum Likelihood Estimation: Theory and Applications," *The Annals of Statistics*, 9, 861-869.

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

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

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

Heckman, J. (1981): "The incidental parameters problem and the problem of initial conditions in estimating a discrete time - discrete data stochastic process," in C. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data with Econometric Applications.* MIT Press.

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

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

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

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

McFadden, D., and W. Newey (1994): "Large Sample Estimation and Hypothesis Testing," in R. E. Engle and McFadden (eds.) *Handbook of Econometrics Volume 4*, North-Holland.

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

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

Moro, A. (2003): "The effect of statistical discrimination on black-white wage inequality: Estimating a model with multiple equilibria," *International Economic Review.*

Moro, A., A. Bisin, and G. Topa (2003): "The empirical implications of models with multiple equilibria," manuscript, University of Minnesota.

Netz, J. and B. Taylor (2002): "Maximum or minimum differentiation? Location patterns or retail outlets," *Review of Economics and Statistics*, 84(10, 162-175.

Pakes, A. and P. McGuire (1994): "Computing Markov-Perfect Nash equilibria: Numerical implications of a dynamic differentiated product model, *Rand Journal of Economics*, 555-589.

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

Pesendorfer, M. and Schmidt-Dengler (2003): "Identification and estimation of dynamic games," London School of Economics Working Paper.

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

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

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

Tamer, E. (2004): "Empirical strategies for estimating discrete games with multiple equilibria," manuscript. Princeton University.

Tauchen, G. (1986): "Finite state Markov chain approximantion to univariate and vector autoregressions," *Economics Letters*, 20, 177-181.

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

Zhang-Foutz, N., and V. Kadiyali (2003): "Competitive Dynamics in Optimal Release Timing of Motion Pictures," manuscript, The S. C. Johnson Graduate School of Management, Cornell University.

Table 1

Table 1
Monte Carlo Experiments
Parameters[1] and Description of the Markov Perfect Equilibrium

Descriptive Statistics	Exp. 1 $\alpha_2 = 1.0$ $\delta = 0.0$	Exp. 2 $\alpha_2 = 1.0$ $\delta = 1.0$	Exp. 3 $\alpha_2 = 1.0$ $\delta = 2.0$	Exp. 4 $\alpha_2 = 0.0$ $\delta = 1.0$	Exp. 5 $\alpha_2 = 2.0$ $\delta = 1.0$	Exp. 6 $\alpha_2 = 4.0$ $\delta = 1.0$
(1) Number active firms: Average	3.676	2.760	1.979	2.729	2.790	2.801
(2) Number of firms: Std. Dev.	1.551	1.661	1.426	1.515	1.777	1.905
(3) AR(1) for number active firms (autoregressive parameter)	0.744	0.709	0.571	0.529	0.818	0.924
(4) Number of Entrants (or Exits) per period	0.520	0.702	0.748	0.991	0.463	0.206
(5) Excess Turnover[2] (in # of firms per period)	0.326	0.470	0.516	0.868	0.211	0.029
(6) Correlation between entries and exits	-0.015	-0.169	-0.220	-0.225	-0.140	-0.110
(7) Prob. 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

Note 1: For all these experiments, the values of the rest of the parameters are: $N = 5$, $\alpha_{01} = -1.9$, $\alpha_{02} = -1.8$, $\alpha_{03} = -1.7$, $\alpha_{04} = -1.6$, $\alpha_{05} = -1.5$, $\alpha_1 = 1.0$, $\sigma_\varepsilon = 1$, and $\beta = 0.95$.

Note 2: Excess turnover is defined as (#Entrants + #Exits)-abs(#Entrants - #Exits).

	Table 2		
	Monte Carlo Experiments		
	Median Number of Iterations of the NPL Algorithm		
	Initial Probabilities		
	Frequencies	Logits	Random
Experiment 1	8	4	6
Experiment 2	11	7	9
Experiment 3	27	19	23
Experiment 4	16	8	11
Experiment 5	12	7	9
Experiment 6	13	9	10

Estimator		*Parmeters*			
		α_{01}	α_1	α_2	δ
Experiment 1	True values	-1.900	1.000	1.000	0.000
	1-stage (True P^0)	-1.915 (0.273)	1.007 (0.152)	1.002 (0.139)	0.002 (0.422)
	1-stage (NP freq. $\hat{P}^0$)	-0.458 (0.289)	0.374 (0.141)	1.135 (0.190)	0.200 (0.364)
	1-stage (Logit $\hat{P}^0$)	-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)
Experiment 2	True values	-1.900	1.000	1.000	1.000
	1-stage (True P^0)	-1.894 (0.212)	1.002 (0.186)	1.007 (0.118)	1.007 (0.583)
	1-stage (NP freq. $\hat{P}^0$)	-0.919 (0.208)	0.351 (0.119)	0.886 (0.123)	0.095 (0.337)
	1-stage (Logit $\hat{P}^0$)	-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)
Experiment 3	True values	-1.900	1.000	1.000	2.000
	1-stage (True P^0)	-1.910 (0.183)	1.006 (0.209)	1.000 (0.112)	2.008 (0.783)
	1-stage (NP freq. $\hat{P}^0$)	-1.126 (0.189)	0.286 (0.094)	0.792 (0.107)	0.027 (0.311)
	1-stage (Logit $\hat{P}^0$)	-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)
Experiment 4	True values	-1.900	1.000	0.000	1.000
	1-stage (True P^0)	-1.890 (0.516)	1.020 (0.329)	0.001 (0.119)	1.063 (1.345)
	1-stage (NP freq. $\hat{P}^0$)	-0.910 (0.243)	0.337 (0.104)	0.239 (0.113)	0.127 (0.354)
	1-stage (Logit $\hat{P}^0$)	-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)
Experiment 5	True values	-1.900	1.000	2.000	1.000
	1-stage (True P^0)	-1.912 (0.178)	1.007 (0.142)	2.008 (0.132)	1.006 (0.359)
	1-stage (NP freq. $\hat{P}^0$)	-0.840 (0.218)	1.379 (0.130)	1.591 (0.143)	0.181 (0.302)
	1-stage (Logit $\hat{P}^0$)	-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)
Experiment 6	True values	-1.900	1.000	4.000	1.000
	1-stage (True P^0)	-1.899 (0.206)	1.003 (0.132)	4.050 (0.203)	1.006 (0.238)
	1-stage (NP freq. $\hat{P}^0$)	-0.558 (0.228)	0.332 (0.128)	2.745 (0.211)	0.206 (0.238)
	1-stage (Logit $\hat{P}^0$)	-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)

Table 4
Square-root Mean Square Error
Relative to the 1-stage PML with true P^0

	Estimator	Parmeters			
		α_{01}	α_1	α_2	δ
Experiment 1	1-stage (NP freq. $\hat{P}^0$)	5.380	4.222	1.676	0.983
	1-stage (Logit $\hat{P}^0$)	1.027	1.006	1.002	1.022
	NPL	1.019	1.040	0.996	1.044
Experiment 2	1-stage (NP freq. $\hat{P}^0$)	4.736	3.553	1.415	1.655
	1-stage (Logit $\hat{P}^0$)	1.070	1.066	1.029	1.034
	NPL	1.098	1.188	1.020	1.171
Experiment 3	1-stage (NP freq. $\hat{P}^0$)	4.347	3.440	2.095	2.549
	1-stage (Logit $\hat{P}^0$)	1.357	1.462	1.301	1.419
	NPL	1.268	0.935	1.038	0.892
Experiment 4	1-stage (NP freq. $\hat{P}^0$)	1.977	2.035	2.228	0.699
	1-stage (Logit $\hat{P}^0$)	0.906	0.848	1.000	0.850
	NPL	0.935	0.884	0.969	0.881
Experiment 5	1-stage (NP freq. $\hat{P}^0$)	6.054	4.459	3.279	2.429
	1-stage (Logit $\hat{P}^0$)	1.146	1.176	1.043	1.130
	NPL	1.143	1.250	1.037	1.210
Experiment 6	1-stage (NP freq. $\hat{P}^0$)	6.591	5.589	6.072	3.487
	1-stage (Logit $\hat{P}^0$)	1.162	1.209	1.020	1.166
	NPL	1.158	1.248	1.010	1.197

Table 5
Descriptive Statistics
189 markets. Years 1994-1999

	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 %
Markets with 1 firm	1.3 %	15.3 %	15.8 %	10.8 %	9.6 %
Markets with 2 firms	1.2 %	7.8 %	8.0 %	6.7 %	5.0 %
Markets with 3 firms	0.5 %	5.2 %	6.9 %	3.8 %	3.4 %
Markets with 4 firms	1.2 %	4.0 %	3.6 %	2.7 %	2.0 %
Markets with 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)[1]	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)[2]	-0.019 (0.034)	0.153 (0.082)	-0.066 (0.050)	0.223 (0.081)	0.097 (0.111)
Entry rate (%)[3]	9.8	14.6	19.7	12.8	21.3
Exit rate (%)[4]	9.9	7.4	13.5	10.4	14.5
Survival rate (Hazard rate): 1 year (%)[5]	86.2 (13.8)	89.5 (10.5)	84.0 (16.0)	86.8 (13.2)	79.7 (20.3)
Survival rate (Hazard rate): 2 years (%)	69.5 (19.5)	88.5 (1.1)	70.0 (16.6)	71.1 (18.2)	58.1 (27.2)
Survival rate (Hazard rate): 3 years (%)	60.1 (14.9)	84.6 (4.3)	60.0 (14.3)	52.6 (25.1)	44.6 (23.3)

Note 1: Market size = Population. Regression included time dummies. Standard errors in parentheses.

Note 2: Firm size = Revenue per firm. Regression included time dummies. Standard errors in parentheses.

Note 3: Entry rate = Entrants / Incumbents

Note 4: Exit rate = Exits / Incumbents

Note 5: Survival and hazard rates are calculated using the subsample of new entrants in years 1995 and 1996.

<div align="center">

Table 6
Distribution of the estimated number of potential entrants

</div>

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

Table 7
NPL estimation of Entry-Exit model

Parameters	Restaurants	Gas stations	Bookstores	Shoe shops	Fish shops
Constant: $\dfrac{\alpha_0}{\sigma}$	-9.519	-12.769	-15.997	-14.497	-6.270
	(0.478)	(1.251)	(0.141)	(1.206)	(1.233)
ln(Pop): $\dfrac{\alpha_1}{\sigma}$	1.743	1.929	2.029	2.030	0.914
	(0.045)	(0.127)	(0.076)	(0.121)	(0.125)
Entry cost: $\dfrac{\alpha_2}{\sigma}$	5.756	10.441	5.620	5.839	4.586
	(0.030)	(0.150)	(0.081)	(0.145)	(0.121)
ln(# firms): $\dfrac{\delta}{\sigma}$	1.643	2.818	1.606	2.724	1.395
	(0.176)	(0.325)	(0.201)	(0.316)	(0.234)
Std. dev of ω_m	1.322	2.029	1.335	2.060	1.880
Number of observations	945	945	945	945	945
R-square: entries	0.298	0.196	0.442	0.386	0.363
R-square: exits	0.414	0.218	0.234	0.221	0.298
Standarized parameters:					
(a) $exp\left(-\alpha_0 \ / \ \alpha_1\right)$	235	750	2658	1267	951
(b) $exp\left(\alpha_2 \ / \ \alpha_1\right)$	27	224	16	18	151
(c) $\delta \ln(2) \ / \ \alpha_1$	65.3 %	101.3 %	54.9 %	93.1 %	105.7 %

Note (a): The value $exp\left(-\alpha_0/\alpha_1\right)$ represents the minimum market population size such that variable profits of a monopolist can cover fixed operating costs.

Note (b): The value $exp\left(\alpha_2/\alpha_1\right)$ represents the minimum market population size such that variable profits of a monopolist can cover entry sunk costs.

Note (c): The value $\delta \ln(2)/\alpha_1$ represents the percentage increase in market size such that profits of a duopolist in the larger market are equal to profits of a monopolist in the smaller market.