The Pinhas Sapir Center
for Development
Tel Aviv University

[redacted]

[redacted]

TEL AVIV UNIVERSITY
The Gordon Faculty of
Social Sciences
Department of Public
Policy

TWELFTH CEPR/JIE CONFERENCE ON APPLIED INDUSTRIAL ORGANIZATION

Hosted by

The Eitan Berglas School of Economics, Recanati School of Business, Department of Public Policy at Tel Aviv University

Supported by

The Pinhas Sapir Center for Development, Tel Aviv University
Recanati School of Business, Tel Aviv University
Department of Public Policy, Tel Aviv University
Journal of Industrial Economics (JIE)
CEPR

Tel Aviv; 24-27 May 2011

Identification and Estimation of Dynamic Games when Players' Beliefs are not in Equilibrium

*Victor Aguirregabiria (University of Toronto)
Arvind Magesan (University of Calgary)

Identification and Estimation of Dynamic Games when Players' Beliefs Are Not in Equilibrium

Victor Aguirregabiria* University of Toronto

Arvind Magesan* University of Calgary

First version: September 2009. This version: May 15, 2011

Abstract

This paper deals with the identification and estimation of dynamic games when players' beliefs about other players' actions may not be in equilibrium, i.e., they do not represent the actual behavior of other players. This type of model applies naturally to competition in oligopoly industries when firms face significant strategic uncertainty such that they have imperfect knowledge about the strategies of their competitors. Our approach can be used to estimate the evolution of players' beliefs when the researcher does not have data on elicited beliefs. First, we show that a standard exclusion restriction, that is typically used to identify payoffs in empirical games, provides testable nonparametric restrictions of the null hypothesis of equilibrium beliefs. Second, we prove that an additional assumption, that we call *no strategic uncertainty at two 'extreme' points*, is sufficient for nonparametric point-identification of payoff functions and players' beliefs. Third, we propose a simple two-step estimation method and a sequential generalization of the method that improves its asymptotic and finite sample properties. Finally, we illustrate our model and methods with an empirical application of a dynamic game of store location by retail chains.

Keywords: Dynamic games; Rational behavior; Rationalizability; Identification; Estimation; Market entry-exit.

Corresponding Author: Victor Aguirregabiria. Department of Economics. University of Toronto. 150 St. George Street. Toronto, Ontatio. `victor.aguirregabiria@utoronto.ca`

*We would like to thank comments from Chris Auld, Dan Bernhardt, Sridhar Moorthy, Katsumi Shimotsu, Jean-Francois Wen and seminar participants at Calgary, Carlos III-Madrid, New York University, UI Urbana-Champaign, at the Canadian Econometric Study Group in Ottawa, and the Econometric Society World Congress in Shanghai. The first author thanks the Social Science and Humanities Research Council of Canada (SSHRCC) for financial support.

Introduction

The principle of *revealed preference* (Samuelson, 1938) is a cornerstone in the structural empirical analysis of decision models, either static or dynamic, single-agent decision problems or games. Under the principle of *revealed preference*, agents maximize expected payoffs and their actions reveal information on the structure of payoff functions. This simple but powerful concept has allowed econometricians to use data on agents' decisions to identify important structural parameters for which there is very limited information from other sources. Some examples of the different types of parameters and functions that have been estimated using the principle of revealed preference are agents' degree of risk aversion, intertemporal rates of substitution, market entry costs, adjustment costs and switching costs, consumer willingness to pay, preference for a political party, or the cost of a merger. In the context of empirical games, a player's expected payoff depends on his beliefs about the behavior of other players. Most empirical applications of games have combined the principle of revealed preference with the assumption that players' beliefs are in equilibrium. There are multiple reasons why the assumption of equilibrium beliefs is useful in the estimation of games. Equilibrium restrictions have identification power even in models with multiple equilibria (Tamer, 2003, and Aradillas-Lopez and Tamer, 2008).[1] Imposing these restrictions contributes to improve asymptotic and finite sample properties of estimators (Kasahara and Shimotsu, 2008). Furthermore, structural models where agents' beliefs are endogenously determined in equilibrium are attractive for the evaluation of counterfactual policy experiments because they take into account how agents' beliefs may change in the counterfactual scenario.

Despite these attractive implications of the assumption of equilibrium beliefs, there are empirical applications of games where the assumption is not realistic and it is of interest to relax it. The following are three examples.[2]

Example 1. Competition in oligopoly industries is often characterized by strategic uncertainty (Besanko et al., 2010). Firm managers are very secretive about their own strategies and face

[1] In games with multiple equilibria, the assumption of equilibrium beliefs is key to have point identification of structural parameters and beliefs, and for the implementation of relatively simple methods of estimation. That is the case in games of incomplete information under the assumption that the observations in the data come from the same equilibrium (Aguirregabiria and Mira, 2007). Under this assumption, it is possible to estimate players' beliefs consistently using a nonparametric estimator of the distribution of players' actions. This nonparametric estimator of beliefs can be used to construct players' expected payoffs and to obtain an estimator of structural parameters that optimizes a sample criterion function based on players' best responses to the estimated beliefs from the data. This simple two-step approach for identification and estimation cannot be applied when players beliefs are not in equilibrium.

[2] See also Morris and Song (2002) for different examples of models with strategic uncertainty and the related experimental evidence.

significant uncertainty about the strategies of their competitors. In fact, it is often the case that firms have incentives to misrepresent their own strategies. For example, a firm would want its rival to believe that it is planning an expansion in a particular location to deter the rival from entering into the location, when in fact there is no such plan. In this context, it can be difficult for firms to construct correct beliefs about the behavior of competitors.

Example 2. Our second example deals with the empirical evaluation of the effect of a policy change in a strategic environment. Suppose that we have panel data of firms competing in an industry, and that an important policy change occurred in the middle of our sample period, e.g., the introduction of a new government subsidy that tries to encourage firms' adoption of a new technology. It seems reasonable to think that it will take time for firms to learn about the new strategies of competitors after the policy change, and for a while firms' beliefs will be out of equilibrium. Imposing the restriction of equilibrium beliefs may bias our estimates of the effects of the new policy.

Example 3. It is well established in the experimental economics literature that there is significant heterogeneity in agents' elicited beliefs, and that this heterogeneity is often one of the most important factors in explaining heterogeneity in observed behavior in laboratory experiments (Camerer, 2003). This evidence is at odds with the standard approach in empirical applications of games using non-experimental data. Most of these applications impose the assumption of equilibrium beliefs, and this implies that beliefs play no role in explaining heterogeneous behavior.

These examples illustrate the importance of relaxing the assumption of equilibrium beliefs in some empirical applications of games. However, a well-known problem in economics is that there is not an alternative assumption that is more acceptable and reasonable than the assumption of equilibrium beliefs. In this context, our paper is concerned with the identification of players' payoff functions and beliefs when the researcher does not have data on elicited beliefs and does not want to make any arbitrary assumption on beliefs.[3]

In this paper we study nonparametric identification, estimation, and inference in dynamic discrete games of incomplete information when we assume that players are rational, in the sense that each player maximizes expected payoff given some arbitrary beliefs, but we relax the assumption of equilibrium beliefs and impose minimum restrictions on beliefs. In the class of econometric models that we consider, players' beliefs are probability distributions over the set of other players' actions. These distributions are nonparametrically specified and they are treated as incidental parameters that, together with the structural parameters of the game, determine the stochastic process followed

[3] For instance, data on elicited beliefs of firm managers is very rare and typically of low quality.

by players' actions. When players beliefs are not in equilibrium they are different from the actual distribution of players' actions in the population. Therefore, without other restrictions, beliefs cannot be identified and estimated by simply using a nonparametric estimator of the distribution of players' actions. First, we show that a standard exclusion restriction, that is typically used to identify payoffs in empirical games, provides testable nonparametric restrictions of the null hypothesis of equilibrium beliefs. Second, we present new results on the nonparametric point-identification of payoff functions and beliefs. We show that, together with the exclusion restriction, a large-support condition on one of the explanatory variables is sufficient for point-identification. While the type of exclusion restriction that we need is quite plausible in dynamic games of oligopoly competition, the large support condition is not satisfied in most applications. We show that this condition can be replaced with a more general restriction that we call *no strategic uncertainty at two 'extreme' points*. Third, we propose a simple two-step estimation method of structural parameters and beliefs, and a sequential extension of this method that provides estimators with better statistical properties. We also present a procedure for testing the null hypothesis of equilibrium beliefs. Finally, we illustrate our model and methods with an empirical application of a dynamic game of store location by retail chains.

This paper builds on the recent literature on estimation of dynamic games of incomplete information (see Aguirregabiria and Mira, 2007, Bajari, Benkard and Levin, 2007, Pakes, Ostrovsky and Berry, 2007, and Pesendorfer and Schmidt-Dengler, 2008). All the papers in this literature have assumed that the data come from a Markov Perfect Equilibrium. We relax that assumption. Our paper also builds and extends the work of Aradillas-Lopez and Tamer (2008) who study the identification power of the assumption of equilibrium beliefs in static games. We extend their work in the following points: (i) they consider static games while we study dynamic games;[4] (ii) they concentrate on identification while we also propose and implement new tests and estimators; and (iii) they deal with models with a parametric specification of the payoff function while we consider nonparametric payoffs.

To illustrate our model and methods, we consider an empirical application of a dynamic game of store location between McDonalds and Burger King. Most empirical studies on bounded rationality have concentrated on individual behavior, and there is very little empirical work on bounded

[4] The implications of dropping the assumption of equilibrium beliefs, and the associated identification issues, are different between static and dynamic games. The characterization and derivation of bounds on choice probabilities is significantly more complicated in dynamic than in static games, and some results in Aradillas-Lopez and Tamer do not extend to dynamic games.

rationality of firms.[5] The estimation of reduced form models using panel data of McDonalds' and Burger King's store location decisions show that the probability that these firms open a new store in a local market does not respond to the number of stores of the competing firm, or even responds positively (Toivanen and Waterson, 2005). This evidence is robust to controlling for unobserved market heterogeneity, and it cannot be explained by a standard static model of store location. We propose and estimate a structural dynamic game of entry in local markets that incorporates three alternative explanations for this puzzle: positive spillover effects; firms' forward looking behavior; and biased beliefs about the behavior of the competitor.

The rest of the paper includes the following sections. Section 2 presents the model and basic assumptions. In section 3, we present our identification results. Section 4 describes estimation methods and testing procedures. The empirical application is described in section 5. We summarize and conclude in section 6.

2 Model

This section presents a dynamic game of incomplete information where two players make binary choices over T periods.[6] The time horizon T can be either finite or infinite. We use the indexes $i \in \{1, 2\}$ and $j \in \{1, 2\}$ to represent a player and his opponent, respectively. Time is discrete and indexed by $t \in \{1, 2, ..., T\}$. Every period t, players choose simultaneously and non-cooperatively between alternatives 0 and 1. Let $Y_{it} \in \{0, 1\}$ represent the choice of player i at period t. Each player makes this decision to maximize his expected and discounted payoff, $E_t(\sum_{s=0}^{T} \beta_i^s \Pi_{i,t+s})$, where $\beta_i \in (0, 1)$ is player i's discount factor, and Π_{it} is his payoff at period t. The one-period payoff function has the following structure:

$$\Pi_{it} = \begin{cases} \pi_{it}(Y_{jt}, \mathbf{X}_t) - \varepsilon_{it} & \text{if} \quad Y_{it} = 1 \\ \\ 0 & \text{if} \quad Y_{it} = 0 \end{cases} \tag{1}$$

Y_{jt} represents the current action of the other player; $\mathbf{X}_t$ is a vector of state variables which are common knowledge for both players; variable ε_{it} is private information of firm i at period t; and $\pi_{it}(.)$ is a real valued function. Note that we normalize to zero the one-period payoff of alternative 0. In the context of nonparametric identification of games, Bajari and Hong (2005) and Bajari et

[5] An exception is the recent paper by Goldfarb and Xiao (2011) that studies entry decisions in the US local telephone industry and finds significant heterogeneity in firms' beliefs about other firms' strategic behavior.

[6] The results in the paper can be generalized to models with more than two players or choice alternatives. A key result in our paper is the characterization of rational beliefs in section 3. That characterization is used in the identification results and in the estimation methods in sections 4 and 5. It is possible to extend that representation to multinomial choice models and games with more than two players.

al. (2009) use also this normalization.[7] However, the results in this paper do not depend on the specific choice of normalization of the payoff function. For instance, in the empirical application in section 5 we make a different normalization assumption.

The vector of common knowledge state variables $\mathbf{X}_t$ has three different components: $\mathbf{X}_t = (\mathbf{W}_t, S_{it}, S_{jt})$. $\mathbf{W}_t \in \mathcal{W}$ is a vector of state variables that evolve exogenously according to a Markov process with transition probability function $f_{Wt}(\mathbf{W}_{t+1}|\mathbf{W}_t)$. The variables $\mathbf{S}_t \equiv (S_{it}, S_{jt}) \in \mathcal{S} \times \mathcal{S}$ are endogenous state variables, and they evolve over time according to a transition probability function $f_{St}(\mathbf{S}_{t+1}|Y_{it}, Y_{jt}, \mathbf{X}_t)$. We refer to S_{it} as the stock or the "capacity" of player i. The private information shocks ε_{1t} and ε_{2t} are independent of $\mathbf{W}_t$, independent of each other, and independently and identically distributed over time. Their distribution functions, Λ_1 and Λ_2, are absolutely continuous and strictly increasing with respect to the Lebesgue measure on $\mathbb{R}$.

EXAMPLE: Dynamic game of market entry and exit. Consider two firms competing in a market. Each firm sells a differentiated product. Every period, firms decide whether or not to be active in the market. Then, incumbent firms compete in prices. Let $Y_{it} \in \{0, 1\}$ represent the decision of firm i to be active in the market at period t. Let S_{it} be a state variable that represents the number of consecutive periods that the firm has been actively operating in the market, i.e., a firm's *experience*. That is, $S_{it} = 0$ means that the firm was not active at $t-1$, and $S_{it} = k > 0$ means that the firm entered the market at $t-k-1$ and has remained active every period until $t-1$. The transition rule of firm experience is deterministic, $S_{it+1} = Y_{it}(S_{it} + Y_{it})$. Let $\mathbf{W}_t$ be a vector of exogenous market characteristics, e.g., market size. The profit function of firm i is: $\Pi_{it} = \pi_i(Y_{jt}, S_{it}, \mathbf{W}_t) - \varepsilon_{it}$ if $Y_{it} = 1$, and $\Pi_{it} = 0$ if $Y_{it} = 0$. In this example, we have that $\pi_i(0, k, \mathbf{W}_t)$ is the profit of firm i when it is a monopolist that has been active in the market for k periods. Similarly, $\pi_i(1, k, \mathbf{W}_t)$ is the profit of firm i when it is a duopolist with k periods of experience. The variable ε_{it} is a private information shock in the fixed operating cost, and it is i.i.d. normally distributed. We study identification and estimation of the nonparametric functions π_i. We also propose an estimation method of player's beliefs and payoffs when the payoff function is parametrically specified. For instance, the profit function could be:

$$\pi_i(Y_{jt}, S_{it}, W_t) \quad = \quad W_t\left((1 - Y_{jt})\,\theta_i^M + Y_{jt}\,\theta_i^D\right) - \theta_{i0}^{FC} - \theta_{i1}^{FC}\exp\{-S_{it}\} - 1\{S_{it} = 0\}\,\theta_i^{EC} \qquad (2)$$

where $1\{.\}$ is the binary indicator function. W_t represents market size, and θ_i^M, θ_i^D, θ_{i0}^{FC}, θ_{i1}^{FC}, and

[7] It is relevant to point out that, in contrast to static decision problems, this normalization is not innocuous in dynamic decision models and it has implications for the predictions of some counterfactual experiments such as those that involve an hypothetical change in the transition probabilities of the state variables. See Aguirregabiria (2010) for an study of this issue in the context of single-agent dynamic decision models.

θ_i^{EC} are structural parameters. θ_i^M and θ_i^D represent the per capita variable profit of firm i when the firm is a monopolist and when it is a duopolist, respectively. θ_i^{EC} is a parameter that represents market entry cost. And θ_{i0}^{FC} and θ_{i1}^{FC} are parameters that represent the fixed operating costs and how they depend on firm's experience. ■

Most previous literature on estimation of dynamic discrete games assumes that the data comes from a Markov Perfect Equilibrium (MPE). This equilibrium concept incorporates three main assumptions.

ASSUMPTION 1 (Payoff relevant state variables): Players' strategy functions depend only on payoff relevant state variables: $\mathbf{X}_t$ and ε_{it}.

ASSUMPTION 2 (Rational beliefs on own future behavior): Players are forward looking, maximize expected intertemporal payoffs, and have rational expectations on their own behavior in the future.

ASSUMPTION 'EQUIL': (Rational or equilibrium beliefs on other players' actions): Strategy functions are common knowledge, and players' have rational expectations on the current and future behavior of other players. That is, players beliefs about other players' behavior are consistent with the actual behavior of other players.

First, let us examine the implications of imposing only Assumption 1. The payoff-relevant information set of player i is $\{\mathbf{X}_t, \varepsilon_{it}\}$. The space of $\mathbf{X}_t$ is $\mathcal{X} \equiv \mathcal{W} \times \mathcal{S}^2$. At period t, players observe $\mathbf{X}_t$ and choose their respective actions. Let $\sigma_{it}(\mathbf{X}_t, \varepsilon_{it})$ be a strategy function for player i at period t. This is a function from the support of $(\mathbf{X}_t, \varepsilon_{it})$ into the binary set $\{0, 1\}$, i.e., $\sigma_{it} : \mathcal{X} \times \mathbb{R} \rightarrow \{0, 1\}$. Given any strategy function σ_{it}, we can define a choice probability function $P_{it}(\mathbf{X}_t)$ that represents the probability of $Y_{it} = 1$ conditional on $\mathbf{X}_t$ given that player i follows strategy σ_{it}. That is,

$$P_{it}(\mathbf{X}_t) \equiv \int 1\left\{\sigma_{it}\left(\mathbf{X}_t, \varepsilon_{it}\right) = 1\right\}\ d\Lambda_i\left(\varepsilon_{it}\right) \tag{3}$$

It is convenient to represent players' behavior using these *Conditional Choice Probability* (CCP) functions. When the variables in $\mathbf{X}_t$ have a discrete support, we can represent the CCP function $P_{it}(.)$ using a finite dimension vector $\mathbf{P}_{it} \equiv \{P_i(\mathbf{X}_t) : \mathbf{X}_t \in \mathcal{X}\}$. Throughout the paper we use either the function $P_{it}(.)$ or the vector $\mathbf{P}_{it}$ to represent the *actual behavior* of player i at period t.

Without imposing Assumption 'Equil' ('Equilibrium Beliefs'), a player's beliefs about the behavior of other players do not necessarily represent the actual behavior of the other player. Therefore, we need functions other than $\sigma_{jt}(.)$ and $P_{jt}(.)$ to represent players i' beliefs about the strategy of player j. Let $b_{jt}^{(t')}\left(\mathbf{X}_t, \varepsilon_{jt}\right)$ be player i's belief at period t' about the strategy function of player j

at period t. In principle, this function may vary with t' due to players' learning and forgetting, or to other factors that make players to change their beliefs over time. Allowing for this very flexible form of time-varying beliefs introduces serious complexities not only in terms of identification but also in the structure of the model. For instance, in this general framework, belief functions become part of the set of payoff relevant state variables, and we should introduce players' beliefs about the evolution of their own future beliefs. Therefore, we assume that a player's belief at period t' about the strategy of his opponents at period t may depend on the time period of the opponents' behavior (t) but not on the time period when the prediction is made (t'), i.e., $b_{jt}^{(t')}(\mathbf{X}_t, \varepsilon_{jt}) = b_{jt}(\mathbf{X}_t, \varepsilon_{jt})$ for every t'.[8]

Let $B_{jt}(\mathbf{X}_t)$ be the choice probability associated with $b_{jt}(\mathbf{X}_t, \varepsilon_{jt})$, i.e., $B_{jt}(\mathbf{X}_t) \equiv \int 1\{b_{jt}(\mathbf{X}_t, \varepsilon_{jt}) = 1\}\ d\Lambda_j(\varepsilon_{jt})$. When $\mathcal{X}$ is a discrete and finite space, we can represent function $B_{jt}(.)$ using a finite-dimensional vector $\mathbf{B}_{jt} \equiv \{B_{jt}(\mathbf{X}_t) : \mathbf{X}_t \in \mathcal{X}\}$, and the vector $\widetilde{\mathbf{B}}_j$ represents $\{\mathbf{B}_{jt} : t = 1, 2, ..., T\}$. Using this notation, Assumption 'Equil' can be represented as $B_{jt}(\mathbf{X}_t) = P_{jt}(\mathbf{X}_t)$ for every t and $\mathbf{X}_t \in \mathcal{X}$.

Suppose that we impose Assumptions 1 and 2, but not Assumption 'Equil'. We say that a strategy function $\sigma_{it}(.)$ (and the associated CCP function P_{it}) is *rational* if for every possible value of $(\mathbf{X}_t, \varepsilon_{it}) \in \mathcal{X} \times \mathbb{R}$ the action $\sigma_{it}(\mathbf{X}_t, \varepsilon_{it})$ maximizes player i's expected and discounted value given his beliefs on the opponent's strategy. Given beliefs $\widetilde{\mathbf{B}}_j$, player i's best response at period t is the optimal solution of a single-agent dynamic programming (DP) problem. This DP problem can be described in terms of: (i) a discount factor, β_i; (ii) a sequence of expected one-period payoff functions, $\{Y_{it}\ (\pi_{it}^{\mathbf{B}}(\mathbf{X}_t) - \varepsilon_{it}) : t = 1, 2, ..., T\}$, where

$$\pi_{it}^{\mathbf{B}}(\mathbf{X}_t) = (1 - B_{jt}(\mathbf{X}_t))\ \pi_{it}(0, \mathbf{X}_t) + B_{jt}(\mathbf{X}_t)\ \pi_{it}(1, \mathbf{X}_t); \tag{4}$$

and (iii) a sequence of transition probability functions $\{f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|Y_{it}, \mathbf{X}_t) : t = 1, 2, ..., T\}$, where

$$\begin{aligned} f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|Y_{it}, \mathbf{X}_t) \quad = \quad & f_{Wt}(\mathbf{W}_{t+1}|\mathbf{W}_t) \\ & [\ (1 - B_{jt}(\mathbf{X}_t))\ f_{St}(\mathbf{S}_{t+1}|Y_{it}, 0, \mathbf{X}_t) + B_{jt}(\mathbf{X}_t)\ f_{St}(\mathbf{S}_{t+1}|Y_{it}, 1, \mathbf{X}_t)\ ] \end{aligned} \tag{5}$$

Let $V_{it}^{\mathbf{B}}(\mathbf{X}_t, \varepsilon_{it})$ be the value function for player i's DP problem given beliefs $\widetilde{\mathbf{B}}_j$. By Bellman's principle, the sequence of value functions $\{V_{it}^{\mathbf{B}} : t = 1, 2, ..., T\}$ can be obtained recursively using backwards induction in the following *Bellman equation*:

$$V_{it}^{\mathbf{B}}(\mathbf{X}_t, \varepsilon_{it}) \quad = \quad \max_{Y_{it} \in \{0,1\}} Y_{it}(\pi_{it}^{\mathbf{B}}(\mathbf{X}_t) - \varepsilon_{it}) + \beta_i \int V_{it+1}^{\mathbf{B}}(\mathbf{X}_{t+1}, \varepsilon_{it+1})\ d\Lambda_i(\varepsilon_{it+1})\ df_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|Y_{it}, \mathbf{X}_t) \tag{6}$$

[8]Note that our model can incorporate Bayesian learning about some exogenous variables in the model, e.g., about demand, or costs. What our assumption rules out is learning about other players' strategies.

The best response function of player i at period t given beliefs $\widetilde{\mathbf{B}}_j$ is the optimal decision rule of this DP problem. Define the *threshold value function*:

$$v_{it}^{\mathbf{B}}(\mathbf{X}_t) \equiv \pi_{it}^{\mathbf{B}}(\mathbf{X}_t) + \beta_i \sum_{\mathbf{X}_t} \left[f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|1, \mathbf{X}_t) - f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|0, \mathbf{X}_t) \right] \bar{V}_{it+1}^{\mathbf{B}}(\mathbf{X}_{t+1}) \tag{7}$$

where $\bar{V}_{it}^{\mathbf{B}}(\mathbf{X}_t)$ is the *integrated value function* $\int V_{it}^{\mathbf{B}}(\mathbf{X}_t, \varepsilon_{it}) d\Lambda_i(\varepsilon_{it})$. We denote $v_{it}^{\mathbf{B}}$ a *threshold value function* because it represents the threshold value of ε_{it} that makes player i indifferent between the choice of alternatives 0 and 1. The optimal response function is:

$$\{Y_{it} = 1\} \ \text{ iff } \ \left\{\varepsilon_{it} \leq v_{it}^{\mathbf{B}}(\mathbf{X}_t)\right\} \tag{8}$$

The *best response probability function* is a probabilistic representation of the best response function. More precisely, it is the best response function indegrated over the distribution of ε_{it}. In this model, the *best response probability function* is $\Lambda_i(v_{it}^{\mathbf{B}}(\mathbf{X}_t))$, i.e., $\Pr(Y_{it} = 1|\mathbf{X}_t) = \int 1\left\{\varepsilon_{it} \leq v_{it}^{\mathbf{B}}(\mathbf{X}_t)\right\} d\Lambda_i(\varepsilon_{it}) = \Lambda_i(v_{it}^{\mathbf{B}}(\mathbf{X}_t))$. Therefore, under Assumptions 1 and 2 the actual behavior of player i, represented by the CCP function $P_{it}(.)$, satisfies the following condition:

$$P_{it}(\mathbf{X}_t) = \Lambda_i\left(v_{it}^{\mathbf{B}}(\mathbf{X}_t)\right) \tag{9}$$

This equation summarizes all the restrictions that Assumptions 1 and 2 impose on players' choice probabilities. The right hand side of equation (9) is the best response function of a rational player. We use $\Lambda_i(\mathbf{v}_{it}^{\mathbf{B}})$ to represent the vector-value function $\{\Lambda_i\left(v_{it}^{\mathbf{B}}(\mathbf{X}_t)\right) : \mathbf{X}_t \in \mathcal{X}\}$.

The concept of Markov Perfect Equilibrium (MPE) is completed with Assumption 'Equil' ('Equilibrium Beliefs'). Under this assumption, players' beliefs are in equilibrium, i.e., $B_{jt}(\mathbf{X}_t) = P_{jt}(\mathbf{X}_t)$ for every period t and state $\mathbf{X}_t$. A MPE can be described as a sequence of CCP vectors, $\{\mathbf{P}_{it}, \mathbf{P}_{jt} : t = 1, 2, ..., T\}$ such that for every player i and time period t, we have that

$$\mathbf{P}_{it} = \Lambda_i\left(\mathbf{v}_{it}^{\mathbf{P}_j}\right) \tag{10}$$

For the rest of the paper, we maintain Assumptions 1 and 2 but we relax the assumption of 'Equilibrium Beliefs'. Our approach is agnostic about the formation of players' beliefs. Assumption 2 establishes that every player is rational, in the sense that his strategy maximizes his expected and discounted payoff given his beliefs on other players' behavior.

3 Identification

Suppose that the researcher has a random sample with realizations of the game over multiple locations and time periods. Using the terminology in empirical applications of games in Industrial

Organization, we employ the term *local market* to refer to a location. We use the letter m to index local markets. The researcher observes a random sample of M local markets with information on $\{Y_{imt},\ S_{imt},\ \mathbf{W}_{mt}\}$ for every player $i \in \{1, 2\}$ and every period $t \in \{1, 2, ..., T\}$. For the moment, we assume that the dynamic game has a finite horizon T. The number of local markets, M, is large and for the identification results in this section we assume that M is infinite. The unobservable variables $\{\varepsilon_{imt}\}$ are assumed to be independently and identically distributed across markets and over time.[9]

We want to use this sample to estimate the model structural 'parameters' or functions: i.e., payoffs $\{\pi_{it},\ \beta_i\}$; transition probabilities $\{f_{St},\ f_{Wt}\}$; distribution of unobservables $\{\Lambda_i : i = 1,\ 2\}$; and beliefs $\{B_{it}\}$. For primitives other than players' beliefs, we make some assumptions that are standard in previous research on identification of static games and of dynamic structural models with rational or equilibrium beliefs.[10] We assume that the distribution of the unobservables, Λ_i, is known to the researcher up to a scale parameter. We study identification of the payoff functions π_{it} up to scale, but for notational convenience we omit the scale parameter.[11] Following the standard approach in dynamic decision models, we assume that the discount factors, β_i, is known to the researcher. Finally, note that the transition probability functions f_{St} and f_{Wt} are nonparametrically identified.[12] Therefore, we concentrate on the identification of the payoff functions π_{it} and belief function B_{it} and assume that $\{f_{St}, f_{Wt}, \Lambda_i, \beta_i : i = 1, 2\}$ are known.

Let $P^0_{imt}(\mathbf{X})$ be the true conditional probability function $\Pr(Y_{imt} = 1|\mathbf{X}_{mt} = \mathbf{X})$ that represents the actual behavior of player i in market m at period t. Let $B^0_{jmt}(\mathbf{X})$ be the probability function with player i's 'true' beliefs in market m at period t. And let $\boldsymbol{\pi}^0 \equiv \{\pi^0_{it} : t = 1, 2;\ t = 1, 2, ..., T\}$ be the true payoff functions in the population. Assumption 3 summarizes our conditions on the Data Generating Process.

ASSUMPTION 3: (A) For every player i, P^0_{imt} is the best response of player i given his beliefs B^0_{jm} and the payoff functions $\boldsymbol{\pi}^0$, i.e., $\mathbf{P}^0_{imt} = \Lambda_i(\mathbf{v}^{\mathbf{B}_m}_{imt}(\pi^0))$. (B) Players' beliefs may vary over time without restrcitions, and may vary over markets with the observable variables $\mathbf{X}$, but a player

[9] Our identification results can be extended to incorporate time-invariant, local market-specific unobservables for the econometrician which are common knowledge to players and that have a distribution with finite support, i.e., finite mixture model. Kasahara and Shimotsu (2009) have derived sufficient conditions for nonparametric identification of the CCP functions and of the finite mixture distribution in Markov dynamic discrete choice models. Under these conditions, our results on identification of beliefs and payoff functions also hold.

[10] See Bajari and Hong (2005), or Bajari et al (2010), among others.

[11] Based on results by Matzkin (1992 and 1994), Aguirregabiria (2010) shows the nonparametric identification of the distribution of the unobservables Λ in a single-agent dynamic structural model where the researcher observes a component of the payoff function, e.g., the firm's revenue, or a component of its fixed cost.

[12] Note that $f_{St}(S'|Y, \mathbf{X}) = \Pr(S_{mt} \quad = S'|Y_{mt} = Y, \mathbf{X}_{mt} = \mathbf{X})$ and $f_{Wt}(\mathbf{W}^{'}|\mathbf{W}) = \Pr(\mathbf{W}_{mt} \quad = \mathbf{W}'|\mathbf{W}_{mt} = \mathbf{W})$. We can estimate consistently these conditional distributions using, for intance, kernel estimators.

has the same beliefs in markets with the same observable characteristics $\mathbf{X}$ *at the same time period, i.e., for every market* m *with* $\mathbf{X}_{mt} = \mathbf{X}$, $B^0_{jmt}(\mathbf{X}) = B^0_{jt}(\mathbf{X})$.

Assumption 3(A) simply establishes that players are rational in the sense that their actual behavior is the best response given their beliefs. Assumption 3(B) establishes that a player should have the same beliefs in two markets with the same state variables and at the same period of time. Note that beliefs can vary across markets according to the state variables in $\mathbf{X}_{mt}$.[13] This assumption plays a similar role for identification as the assumption of *"no multiple equilibria in the data"* in dynamic games under the condition that beliefs are in equilibrium.[14]

In dynamic games where beliefs are in equilibrium, the "no multiple equilibria in the data" assumption effectively allows the researcher to identify player beliefs. Under this assumption, conditional choice probabilities are identified, and if beliefs are in equilibrium, these beliefs are equal to the conditional choice probabilities. When beliefs are not in equilibrium we can not identify beliefs in this way. However, assumption 3B still implies that CCPs are identified from the data. This assumption implies that for any player i, any period t, and any value of $\mathbf{X} \in \mathcal{X}$, we have that $P^0_{imt}(\mathbf{X}) = P^0_{it}(\mathbf{X}) = E(Y_{imt}|\mathbf{X}_{mt} = \mathbf{X})$, and this conditional expectation can be estimated consistently using data on Y_{imt} and $\mathbf{X}_{mt}$. This in turn, as we will show, is important for the identification of beliefs themselves.

For notational simplicity, we omit the market subindex m for the rest of this section. The model restrictions are summarized in the best response conditions $P^0_{it}(\mathbf{X}_t) = \Lambda_i\left(v^{\mathbf{B}}_{it}\ (\mathbf{X}_t, \pi^0)\right)$. Given these conditions, we want to identify payoffs π^0 and beliefs $\mathbf{B}^0_j$. It is simple to verify that, without further restriction, the order condition for identification is not satisfied. Suppose that the state space $\mathcal{X}$ is finite, and $|\mathcal{X}|$ represents its dimension or number of elements in the set. For a given time period and player, the number of restrictions in the model is $|\mathcal{X}|$ and the total number of parameters is $3|\mathcal{X}|$, so the order condition of identification does not hold.[15] It is important to note that this order condition does not hold even if we assume that players' beliefs are in equilibrium and equal to CCPs. Under that restriction, the model has $2|\mathcal{X}|$ restrictions and $3|\mathcal{X}|$ unknowns.

[13] Our sampling design with large M and small T is the standard case in applications of empirical games in Industrial Organization. Alternatively, if the sampling design is such that the number of periods T is large and the number of markets M is small, then we can allow beliefs to vary over markets but being constant over time.

[14] See Assumptions 5(A) and 5'(A) in Aguirregabiria and Mira (2007) and the discussion of these assumptions in pages 14-15 and 25-26 of that paper. Other recent papers on the estimation of dynamic games, such as Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007), and Pesendorfer ans Schmidt-Dengler (2008) consider similar assumptions. Aguirregabiria and Mira (2009) relax that assumption to allow for multiple equilibria in the data.

[15] The number of parameters in the payoff function is $2|\mathcal{X}|$ (i.e., $|\mathcal{X}|$ parameters for each possible value of the opponent's action, Y_j), and the number of parameters in beliefs is $|\mathcal{X}|$.

A common restriction that has been used to obtain nonparametric identification of payoff functions in games with equilibrium beliefs is a particular kind of exclusion restriction (see Bajari and Hong, 2005, and Bajari et al., 2010). We follow this approach.

ASSUMPTION 4 (Exclusion Restriction): The one-period payoff function of player i depends on the actions of both players, Y_{it} and Y_{jt}, the common state variables $\mathbf{W}_t$, and the own stock variable, S_{it}, but it does not depend on the stock variable of the other player, S_{jt}.

$$\pi_{it}(Y_{jt}, \mathbf{X}_t) = \pi_{it}(Y_{jt}, \mathbf{W}_t, S_{it})$$

This type of exclusion restriction appears naturally in some dynamic games. For instance, this assumption holds in the dynamic game of market entry and exit of our Example in section 2 above. A firm's profits depend on whether the competitor is currently in the market or not but, given that, it does not depend on whether the competitor was in the market or not at previous period. Similarly, in a dynamic game of Cournot competition with capacity constraints (e.g., Besanko and Doraszelski, 2004) a firm's profit depend on the current capacity of competing firms but not on the capacity of these firms at previous periods. A similar structure appears in the dynamic game of retail store competition that we present and estimate in section 5.

The exclusion restriction in Assumption 4 is sufficient to identify the payoff function in a dynamic game where beliefs are assumed in equilibrium. However, it is not enough for the identification of our model. The order condition of identification is not satisfied. The number of restrictions is still the same, $|\mathcal{X}| = |\mathcal{W}|\ |\mathcal{S}|^2$, the number of parameters in the payoff function is now $2|\mathcal{W}|\ |\mathcal{S}|$, and the number of parameters in beliefs is still $|\mathcal{X}| = |\mathcal{W}|\ |\mathcal{S}|^2$.

Before we introduce additional restrictions that provide identification of payoffs and beliefs in our model, we want to show that based on the exclusion restriction in Assumption 4, and *without further restrictions*, we can test the null hypothesis of equilibrium beliefs.

PROPOSITION 1: Under Assumptions 1 to 4, the null hypothesis of equilibrium beliefs is testable.

Proof: Consider the model at the last period T. For notational simplicity we omit the time subindex. Under Assumptions 1 to 3, the CCP functions P_i^0 and P_j^0 are identified everywhere in the support of $\mathbf{X}$. At period T, the threshold function $v_i^{\mathbf{B}}(\mathbf{X}, \pi)$ is equal to $\pi_i^{\mathbf{B}}(\mathbf{X}) = (1 - B_j(\mathbf{X}))\ \pi_i(0, \mathbf{X}) + B_j(\mathbf{X})\ \pi_i(1, \mathbf{X})$. Therefore, the best response condition is $P_i^0(\mathbf{X}) = \Lambda_i((1 - B_j^0(\mathbf{X}))\pi_i^0(0, \mathbf{X}) + B_j^0(\mathbf{X})\pi_i^0(1, \mathbf{X}))$. Define the function $q_i^0(\mathbf{X}) \equiv \Lambda_i^{-1}(P_i^0(\mathbf{X}))$. Given that the distribution function Λ_i is invertible and it is known (up to scale) to the researcher, the function $q_i^0(.)$ is identified

everywhere in the support of $\mathbf{X}$. By definition of q_i^0 and by the best response condition, we have that:

$$q_i^0(\mathbf{X}) = \pi_i^0(0,\mathbf{X}) + \left[\pi_i^0(1,\mathbf{X}) - \pi_i^0(0,\mathbf{X})\right]\ B_j^0(\mathbf{X}) \tag{11}$$

Consider equation (11) evaluated at four different values of the vector $\mathbf{X}$, say $\mathbf{X}^a$, $\mathbf{X}^b$, $\mathbf{X}^c$, and $\mathbf{X}^d$. These four vectors have been constructed such that they have exactly the same value of the component $(S_i, \mathbf{W})$, but they can have different values for the element S_j, say S_j^a, S_j^b, S_j^c,and S_j^d. The exclusion restriction in Assumption 4 implies that the payoff function π_i^0 does not depend on variable S_j. Therefore, we have that $\pi_i^0(0,\mathbf{X}^a) = \pi_i^0(0,\mathbf{X}^b) = \pi_i^0(0,\mathbf{X}^c) = \pi_i^0(0,\mathbf{X}^d)$. If we subtract equation b from equation a, we obtain the expression $q_i^0(\mathbf{X}^a) - q_i^0(\mathbf{X}^b) = \left[\pi_i^0(1,S_i,\mathbf{W}) - \pi_i^0(0,S_i,\mathbf{W})\right]$ $[B_j^0(\mathbf{X}^a) - B_j^0(\mathbf{X}^b)]$, and similarly if we substract equation d from c, we get $q_i^0(\mathbf{X}^c) - q_i^0(\mathbf{X}^d) = \left[\pi_i^0(1,S_i,\mathbf{W}) - \pi_i^0(0,S_i,\mathbf{W})\right]$ $[B_j^0(\mathbf{X}^c) - B_j^0(\mathbf{X}^d)]$. Finally, if $S_j^a \neq S_j^b$ or $S_j^c \neq S_j^d$, we can obtain the ratio between these two pairwise-difference equations, and this ratio implies that:

$$\frac{q_i^0(\mathbf{X}^a) - q_i^0(\mathbf{X}^b)}{q_i^0(\mathbf{X}^c) - q_i^0(\mathbf{X}^d)} = \frac{B_j^0(\mathbf{X}^a) - B_j^0(\mathbf{X}^b)}{B_j^0(\mathbf{X}^c) - B_j^0(\mathbf{X}^d)} \tag{12}$$

This expression shows that, under assumptions 1 to 4, there is a function of beliefs that is identified, without having to impose the assumption of equilibrium beliefs. This result provides a nonparametric test for the null hypothesis of equilibrium beliefs. Define the function:

$$\delta_i(\mathbf{X}^a,\mathbf{X}^b,\mathbf{X}^c,\mathbf{X}^d) \equiv \left\{\frac{q_i^0(\mathbf{X}^a) - q_i^0(\mathbf{X}^b)}{q_i^0(\mathbf{X}^c) - q_i^0(\mathbf{X}^d)}\right\} - \left\{\frac{P_j^0(\mathbf{X}^a) - P_j^0(\mathbf{X}^b)}{P_j^0(\mathbf{X}^c) - P_j^0(\mathbf{X}^d)}\right\} \tag{13}$$

It is clear that we can nonparametrically identify δ_i. If beliefs are in equilibrium, δ_i should be equal to zero for any value of $(\mathbf{X}^a,\mathbf{X}^b,\mathbf{X}^c,\mathbf{X}^d)$. Let $\hat{\delta}_i$ be a root-M consistent and asymptotically normal nonparametric estimator of δ_i. The hypothesis of equilibrium beliefs implies $\delta_i = 0$, and we can test equilibrium beliefs using a simple LM test of the null hypothesis of $\delta_i = 0$. We describe this test in more detail in section 4. ■

We now present two restrictions such that either of them, together with the exclusion restriction, is sufficient to identify payoffs and beliefs in the model.

ASSUMPTION 5a (Monotonic beliefs and large support): (i) The beliefs function $B_{jt}^0(\mathbf{X}_t)$ is strictly monotonically decreasing (or increasing) in the stock variable S_{jt}, i.e., $\partial B_{jt}^0(\mathbf{X}_t)/\partial S_{jt} < 0$. *(ii) Conditional on $(\mathbf{W}_t, S_{it})$ the probability distribution of S_{jt} has unbounded support over the whole real line.*

ASSUMPTION 5b (No strategic uncertainty at two 'extreme' points): There are two values in the support of the distribution of S_{jt}, say S_j^{low} and S_j^{high}, such that for every value of (S_i, W) we have

that beliefs are in equilibrium, i.e., $B^0_{jt}(S^{low}_j, S_i, W) = P^0_{jt}(S^{low}_j, S_i, W)$*, and* $B^0_{jt}(S^{high}_j, S_i, W) = P^0_{jt}(S^{high}_j, S_i, W)$*.*

Assumptions 5a and 5b are two alternative conditions that, together with the exclusion restriction, can provide nonparametric identification of payoff functions and beliefs. These two assumptions are related and the intuition behind them is similar. Assumption 5a implies that when the opponent's stock variable is large enough (or small enough) strategic uncertainty disappears because the opponent's choice becomes certain (i.e., the choice probability gets arbitrarily close to zero or one). But the conditions for the absence of strategic uncertainty may be more general, as in Assumption 5b.

The following Proposition summarizes our main identification result.

PROPOSITION 2: Under Assumptions 1-4 and either 5a or 5b the payoff functions $\{\pi^0_{it}$ *for any* $i, t\}$ *and the beliefs functions* $\{B^0_{it}$ *for any* $i, t\}$ *are nonparametrically identified.*

Proof. The proof of this Proposition follows a backwards induction approach. First, we show identification of beliefs and payoff functions at period T. Then, we apply backwards induction to prove identification of these functions at any period $t < T$. Throughout this proof, we assume that the time horizon T is finite and that the researcher has panel data that covers all the periods until T. At the end of this section we discuss identification when we relax these conditions.

Let us examine the identification of the payoff function π^0_i and the beliefs function B^0_j from the best response equation (11) at the last period T. First, we present the proof using Assumption 5a (*monotonic beliefs and large support condition*). By the monotonicity of the beliefs function B^0_j, as $S_j \to +\infty$ we have that $B^0_j(\mathbf{X}) \to 0$. Therefore, for arbitrarily large values of S_j we have that $\pi^0_i(0, S_i, \mathbf{W}) = q^0_i(\mathbf{X})$ and variables $(S_i, \mathbf{W})$ have variability over their whole support. Therefore, $\pi^0_i(0, .)$ is identified. Similarly, as $S_j \to -\infty$ we have that $B^0_j(\mathbf{X}) \to 1$, and $\pi^0_i(1, S_i, \mathbf{W}) = q^0_i(\mathbf{X})$ such that $\pi^0_i(1, .)$ is identified over the whole support of $(S_i, \mathbf{W})$. Given that π^0_i is identified, then we can identify the beliefs function by solving beliefs in equation (11) to obtain $B^0_j(\mathbf{X}) = [q^0_i(\mathbf{X}) - \pi^0_i(0, \mathbf{X})]/[\pi^0_i(1, \mathbf{X}) - \pi^0_i(0, \mathbf{X})]$ for any value of $\mathbf{X}$ with $\pi^0_i(1, \mathbf{X}) - \pi^0_i(0, \mathbf{X}) \neq 0$.

Now, we prove identification at period T when we replace Assumption 5a by Assumption 5b. First, note that the order condition of identification is satisfied: the number of restrictions is $|\mathcal{X}| = |\mathcal{W}|\ |\mathcal{S}|^2$, the number of unknown parameters in the payoff function is $2|\mathcal{W}|\ |\mathcal{S}|$, and the number of unknown parameters in beliefs is $|\mathcal{W}||\mathcal{S}|^2 - 2|\mathcal{W}|\ |\mathcal{S}|$. Therefore, the number of unknown parameters is equal to the number of restrictions. To show that the rank condition of identification is also satisfied, we now derive closed-form expressions of the unknown parameters in terms of

the identified functions q_i^0 and P_j^0. Substituting the restrictions imposed by Assumption 5b into equation (11), we have that for any value of $(S_i, \mathbf{W})$:

$$\begin{aligned} q_i^0(S_j^{low}, S_i, \mathbf{W}) &= \pi_i^0(0, S_i, \mathbf{W}) + \left[\pi_i^0(1, S_i, \mathbf{W}) - \pi_i^0(0, S_i, \mathbf{W})\right] \; P_j^0(S_j^{low}, S_i, \mathbf{W}) \\ q_i^0(S_j^{high}, S_i, \mathbf{W}) &= \pi_i^0(0, S_i, \mathbf{W}) + \left[\pi_i^0(1, S_i, \mathbf{W}) - \pi_i^0(0, S_i, \mathbf{W})\right] \; P_j^0(S_j^{high}, S_i, \mathbf{W}) \end{aligned} \tag{14}$$

Combining these two equations, we obtain $[\pi_i^0(1, S_i, \mathbf{W}) - \pi_i^0(0, S_i, \mathbf{W})] = [q_i^0(S_j^{high}, S_i, \mathbf{W}) - q_i^0(S_j^{low}, S_i, \mathbf{W})] \,/\, [P_j^0(S_j^{high}, S_i, \mathbf{W}) - P_j^0(S_j^{low}, S_i, \mathbf{W})]$. And substituting this expression into the equations in (14), we identify the payoff function as:

$$\begin{aligned} \pi_i^0(0, S_i, \mathbf{W}) &= q_i^0(\mathbf{X}^{low}) - P_j^0(\mathbf{X}^{low}) \left[\frac{q_i^0(\mathbf{X}^{high}) - q_i^0(\mathbf{X}^{low})}{P_j^0(\mathbf{X}^{high}) - P_j^0(\mathbf{X}^{low})}\right] \\ \pi_i^0(1, S_i, \mathbf{W}) &= q_i^0(\mathbf{X}^{low}) + \left[1 - P_j^0(\mathbf{X}^{low})\right] \left[\frac{q_i^0(\mathbf{X}^{high}) - q_i^0(\mathbf{X}^{low})}{P_j^0(\mathbf{X}^{high}) - P_j^0(\mathbf{X}^{low})}\right] \end{aligned} \tag{15}$$

where $\mathbf{X}^{low} \equiv (S_j^{low}, S_i, \mathbf{W})$ and $\mathbf{X}^{high} \equiv (S_j^{high}, S_i, \mathbf{W})$. Again, given the identification of the payoff function, we can obtain beliefs as $B_j^0(\mathbf{X}) = [q_i^0(\mathbf{X}) - \pi_i^0(0, \mathbf{X})] \,/\, [\pi_i^0(1, \mathbf{X}) - \pi_i^0(0, \mathbf{X})]$.

The equations in (15) deserve further comments. First, these equations contain as a particular case the model with the large support condition in Assumption 5a. That is, the conditions of Assumption 5a can be seen as a particular case of the conditions in Assumption 5b. In particular, suppose that S_j^{high} is so large that $P_j^0(S_i, S_j^{high}, \mathbf{W}) = 0$, and S_j^{low} is so small that $P_j^0(S_i, S_j^{low}, \mathbf{W}) = 1$. Then, the equations in (15) become $\pi_i^0(0, \mathbf{X}) = q_i^0(S_j^{low}, S_i, \mathbf{W})$ and $\pi_i^0(1, \mathbf{X}) = q_i^0(S_j^{high}, S_i, \mathbf{W})$, which are also the expressions that identify the payoff function under the large support conditions in Assumption 5a.

Second, the equations in (15) relate also to the model with equilibrium beliefs. When we impose the assumption of equilibrium beliefs, the equations in (15) are satisfied not only for the pair of values (S_j^{low}, S_j^{high}) but for every pair of values of S_j, say S_j^a and S_j^b, such that $P_j^0(S_j^a, S_i, \mathbf{W}) - P_j^0(S_j^b, S_i, \mathbf{W}) \neq 0$. Therefore, the payoff function is over-identified, and we can test these over identifying restrictions, as we have mentioned above.

Given beliefs and payoffs at period T, we can use backwards induction to prove the identification of the payoff function and beliefs at any period $t < T$. Suppose that the functions π_{it+s}^0 and B_{jt+s}^0 are known for every $s \geq 1$. Given this information, we want to identify functions π_{it}^0 and B_{jt}^0. Player i's best response implies that $P_{it}^0(\mathbf{X}_t) = \Lambda_i(v_{it}^{\mathbf{B}}\ (\mathbf{X}_t))$, where $v_{it}^{\mathbf{B}}\ (\mathbf{X}_t) = (1 - B_{jt}^0(\mathbf{X}_t))\ \pi_{it}^0(0, \mathbf{X}_t) + B_{jt}^0(\mathbf{X}_t)\ \pi_{it}^0(1, \mathbf{X}_t) +\ \beta_i \sum_{\mathbf{X}_t}\ [f_{it}^{\mathbf{B}}\ (\mathbf{X}_{t+1}|1, \mathbf{X}_t) -\ f_{it}^{\mathbf{B}}\ (\mathbf{X}_{t+1}|0, \mathbf{X}_t)]\ \bar{V}_{it+1}^{\mathbf{B}}(\mathbf{X}_{t+1})$, and $\bar{V}_{it+1}^{\mathbf{B}}$ is the integrated values function $\int V_{it+1}^{\mathbf{B}}(\mathbf{X}_{t+1}, \varepsilon_{it+1}) d\Lambda_i(\varepsilon_{it+1})$. It is straightforward to show that given

$\{P^0_{it+s}, \pi^0_{it+s}, B^0_{jt+s} : s = 1, 2, ..., T - t\}$ the integrated value function $\bar{V}^{\mathbf{B}}_{it+1}$ is known. Then, define the function:

$$q^0_{it}(\mathbf{X}) \equiv \Lambda_i^{-1}(P^0_{it}(\mathbf{X})) - \beta_i \sum_{\mathbf{X}_t} [f^{\mathbf{B}}_{it}\ (\mathbf{X}_{t+1}|1, \mathbf{X}_t) - f^{\mathbf{B}}_{it}\ (\mathbf{X}_{t+1}|0, \mathbf{X}_t)]\ \bar{V}^{\mathbf{B}}_{it+1}(\mathbf{X}_{t+1}) \tag{16}$$

This function is identified everywhere in the support of $\mathbf{X}$. Furthermore, by the best response condition and the definition of q^0_{it}, we have that:

$$q^0_{it}(\mathbf{X}) = \pi^0_{it}(0, \mathbf{X}) + \left[\pi^0_{it}(1, \mathbf{X}) - \pi^0_{it}(0, \mathbf{X})\right]\ B^0_{jt}(\mathbf{X}) \tag{17}$$

This equation has the same structure as the best response condition at last period T. Therefore, we can use exactly the same arguments as above to show the identification of functions π^0_{it} and B^0_{jt}. ■

So far, we have assumed that the time horizon T is finite and that the researcher has panel data that covers all the periods until T. Can the identification result in Proposition 2 be extended to a model with infinite time horizon? For the sake of saving notation, we continue using T to represent the number of periods in the sample, but now T does not represent the time horizon of the decision problem because that horizon infinite. In order to get identification of this model we need to impose some restrictions on preferences, transition probabilities, and players' beliefs at periods $t > T$ that are out of our sample. We could consider different types of assumptions. Here we describe the approach we have followed in our application in section 5. In the case of an infinite horizon, it seems natural to assume that preferences, transition probabilities, and beliefs do not vary over time after period T: for any period $t \geq T$, $\pi^0_{it} = \pi^0_{iT}$, $f_{Wt} = f_{WT}$, $f_{St} = f_{ST}$, and $B^0_{it} = B^0_{iT}$. This implies that from period T the game has a stationary Markov structure. Therefore, best response probabilities are also time-invariant: for any period $t \geq T$, $P^0_{it} = P^0_{iT}$. For given beliefs, the CCP function (or vector), $P^0_{iT}(.)$ is the unique solution of a contraction mapping (see Aguirregabiria and Mira, 2002 and 2007). We assume that from period T and forever in the future, agents have rational beliefs on the behavior of the opponent: for any period $t \geq T$, $B^0_{it} = P^0_{iT}$. Aguirregabiria (2005) and Tan (2008) have proved nonparametric identification of preferences in this case.

4 Estimation and Inference

We begin this section by presenting a simple nonparametric test of the null hypothesis of equilibrium beliefs. Sections 4.2 and 4.3 deal with estimation. Our proof of identification above suggests a method for the estimation of the just-identified nonparametric model. Section 4.2 provides a

detailed description of that estimation method. In most empirical applications, the specification of the payoff function involves parametric restrictions. Therefore, we extend the estimation method to deal with parametric models. Finally, for parametric models, we describe how the two-step method can be extended recursively to generate a sequence of estimators with better statistical properties.

4. Test of Equilibrium Beliefs

In principle, we could use a standard Lagrange Multiplier (LM) or Score test of the null hypothesis of equilibrium beliefs. That test is based on the constrained maximum likelihood estimation (MLE) of structural parameters and beliefs. Let $\boldsymbol{\theta}$ be the vector of structural parameters in the payoff function. Define the log-likelihood function:

$$l(\boldsymbol{\theta},\mathbf{P}) \equiv \sum_{m=1}^{M}\sum_{t=1}^{T}\sum_{i=1}^{2} Y_{imt}\ \log \Lambda_i(v_{it}^{\mathbf{P}}(\mathbf{X}_{mt},\boldsymbol{\theta})) + (1-Y_{imt})\ \log(1-\Lambda_i(v_{it}^{\mathbf{P}}(\mathbf{X}_{mt},\boldsymbol{\theta}))) \tag{18}$$

The constrained MLE is defined as a vector $(\hat{\boldsymbol{\theta}}_{MLE}, \hat{\mathbf{P}}_{MLE})$ such that:

$$\begin{array}{rcl} (\hat{\boldsymbol{\theta}}_{MLE}, \hat{\mathbf{P}}_{MLE}) & = & \arg\max\limits_{(\theta,\mathbf{P})}\ l(\boldsymbol{\theta},\mathbf{P}) \\ & & \text{subject to:}\quad \mathbf{P} = \Lambda(v^{\mathbf{P}}(\boldsymbol{\theta})) \end{array} \tag{19}$$

We want to test the null hypothesis $\mathbf{P} = \Lambda(v^{\mathbf{P}}(\boldsymbol{\theta}))$, that consists of $2|\mathcal{X}|$ constrains on $(\boldsymbol{\theta},\mathbf{P})$. The standard LM statistic for testing this null hypothesis is:

$$LM = \frac{\partial l(\hat{\boldsymbol{\theta}}_{MLE}, \hat{\mathbf{P}}_{MLE})}{\partial(\boldsymbol{\theta},\mathbf{P})'} \left[\frac{\partial^2 l(\hat{\boldsymbol{\theta}}_{MLE}, \hat{\mathbf{P}}_{MLE})}{\partial(\boldsymbol{\theta},\mathbf{P})\partial(\boldsymbol{\theta},\mathbf{P})'}\right]^{-1} \frac{\partial l(\hat{\boldsymbol{\theta}}_{MLE}, \hat{\mathbf{P}}_{MLE})}{\partial(\boldsymbol{\theta},\mathbf{P})} \tag{20}$$

Under the null hypothesis, this statistic is asymptotically distributed as a chi-square with $2|\mathcal{X}|$ degrees of freedom.

This LM test has at least two important limitations. A first limitation is its implementation. Maximum likelihood estimation of dynamic games is computationally very demanding both because the high dimension of the state space and because of the existence of multiple equilibria. Second, this is a general specification test. The null hypothesis is not only that beliefs are in equilibrium but also that the parametric specification of preferences and the distribution of unobservables is correct. We would like to have a procedure that specifically tests for the equilibrium beliefs and not for other specification assumptions of the model.

The test that we propose is the following. Let $\mathbf{X}^a$, $\mathbf{X}^b$, $\mathbf{X}^c$, and $\mathbf{X}^d$ be four different values of the vector $\mathbf{X}$ that have the same value of the component $(S_i,\mathbf{W})$ and different values for the element S_j, say S_j^a, S_j^b, S_j^c,and S_j^d, such that $S_j^a \neq S_j^b$ and $S_j^c \neq S_j^d$. Define the function δ_i^0:

$$\delta_i^0(\mathbf{X}^a,\mathbf{X}^b,\mathbf{X}^c,\mathbf{X}^d) \equiv \left\{\frac{q_i^0(\mathbf{X}^a)-q_i^0(\mathbf{X}^b)}{q_i^0(\mathbf{X}^c)-q_i^0(\mathbf{X}^d)}\right\} - \left\{\frac{P_j^0(\mathbf{X}^a)-P_j^0(\mathbf{X}^b)}{P_j^0(\mathbf{X}^c)-P_j^0(\mathbf{X}^d)}\right\} \tag{21}$$

As shown in section 3, under Assumptions 1-4, if player i has rational beliefs then $\delta_i^0(\mathbf{X}^a, \mathbf{X}^b, \mathbf{X}^c, \mathbf{X}^d) = 0$ for every value of $(\mathbf{X}^a, \mathbf{X}^b, \mathbf{X}^c, \mathbf{X}^d)$. Therefore, testing $\delta_i^0(\mathbf{X}^a, \mathbf{X}^b, \mathbf{X}^c, \mathbf{X}^d) = 0$ implies testing the null hypothesis of rational beliefs (and Assumptions 1 to 4).

Let H be the number of all possible combinations of four different values of S_j with $S_j^a \neq S_j^b$ and $S_j^c \neq S_j^d$. We index these values by h. Let $\tilde{\mathbf{X}}_{mt}^{(h)}$ be the quadruplet h when the values of S_i and $\mathbf{W}$ are the ones in observation (m,t): i.e., $\tilde{\mathbf{X}}_{mt}^{(h)} = (\mathbf{X}_{mt}^{(h)a}, \mathbf{X}_{mt}^{(h)b}, \mathbf{X}_{mt}^{(h)c}, \mathbf{X}_{mt}^{(h)d}) = ([S_j^{(h)a}, S_{imt}, \mathbf{W}_{mt}], [S_j^{(h)b}, S_{imt}, \mathbf{W}_{mt}], [S_j^{(h)c}, S_{imt}, \mathbf{W}_{mt}], [S_j^{(h)d}, S_{imt}, \mathbf{W}_{mt}])$. Under the hypothesis of equilibrium beliefs, we have that $E(\delta_i^0(\tilde{\mathbf{X}}_{mt}^{(h)})) = 0$ for every h, where the expectation $E(.)$ is taken over the distribution of $(S_{imt}, \mathbf{W}_{mt})$. This is exactly the null hypothesis that we test:

$$H_0 : E\left(\delta_i^0\left(\tilde{\mathbf{X}}_{mt}^{(h)}\right)\right) = 0 \quad \text{for every quadruple } h \tag{22}$$

Let $\hat{\delta}_i^0(.)$ be the estimator of $\delta_i^0(.)$ that we obtain when we replace P_i^0 and P_j^0 by nonparametric estimates of these CCP functions. Define the statistic $\hat{d}_i^{(h)}$ as the sample mean of $\hat{\delta}_i^0\left(\tilde{\mathbf{X}}_{mt}^{(h)}\right)$, i.e., $\hat{d}_i^{(h)} = (MT)^{-1}\sum_{m=1}^{M}\sum_{t=1}^{T}\hat{\delta}_i^0\left(\tilde{\mathbf{X}}_{mt}^{(h)}\right)$. Then, define the statistic:

$$\hat{D} = \sum_{h=1}^{H}\left(\frac{\hat{d}_i^{(h)}}{se(\hat{d}_i^{(h)})}\right)^2 \tag{23}$$

where $se(\hat{d}_i^{(h)})$ is the standard error of $\hat{d}_i^{(h)}$, that we can obtain using nonparametric bootstrap. Under the null hypothesis, $\hat{S}$ is asymptotically distributed as a Chi-square with H degrees of freedom.

4.2 Estimation with nonparametric payoff function

Step 1: Nonparametric estimation of CCPs, $\hat{P}_{it}^0$, for every player, time period, and state $\mathbf{X}$, and (if needed) of the transition probabilities f_{St} and f_{Wt}.

Step 2: Estimation of preferences and beliefs. At the last period T, we construct $\hat{q}_{iT}^0(\mathbf{X}) = \Lambda_i^{-1}(\hat{P}_{iT}^0(\mathbf{X}))$ for any value of $\mathbf{X}$ in the support $\mathcal{X}$. The estimated payoff function at point $(S_i, \mathbf{W})$ is:

$$\begin{aligned}
\hat{\pi}_{iT}^0(0, S_i, \mathbf{W}) &= \hat{q}_{iT}^0(\mathbf{X}^{low}) - \hat{P}_{jT}^0(\mathbf{X}^{high})\left[\frac{\hat{q}_{iT}^0(\mathbf{X}^{high}) - \hat{q}_{iT}^0(\mathbf{X}^{low})}{\hat{P}_{jT}^0(\mathbf{X}^{high}) - \hat{P}_{jT}^0(\mathbf{X}^{low})}\right] \\
\hat{\pi}_{iT}^0(1, S_i, \mathbf{W}) &= \hat{q}_{iT}^0(\mathbf{X}^{low}) + \left[1 - \hat{P}_{jT}^0(\mathbf{X}^{high})\right]\left[\frac{\hat{q}_{iT}^0(\mathbf{X}^{high}) - \hat{q}_{iT}^0(\mathbf{X}^{low})}{\hat{P}_{jT}^0(\mathbf{X}^{high}) - \hat{P}_{jT}^0(\mathbf{X}^{low})}\right]
\end{aligned} \tag{24}$$

where $\mathbf{X}^{low} \equiv (S_j^{low}, S_i, \mathbf{W})$ and $\mathbf{X}^{high} \equiv (S_j^{high}, S_i, \mathbf{W})$. The estimated beliefs function is:

$$\hat{B}_{jT}^0(\mathbf{X}) = \frac{\hat{q}_{iT}^0(\mathbf{X}) - \hat{\pi}_{iT}^0(0, \mathbf{X})}{\hat{\pi}_{iT}^0(1, \mathbf{X}) - \hat{\pi}_{iT}^0(0, \mathbf{X})} \tag{25}$$

We also construct an estimate of the integrated value function at period T, that we will use later for the estimation of payoffs and beliefs at periods earlier than T. If the distribution Λ_i is logistic, then:

$$\hat{V}_{iT}^{\mathbf{B}}(\mathbf{X}) = -\ln\left(1 - \hat{P}_{iT}^{0}(\mathbf{X})\right) \tag{26}$$

At any period $t < T$, we first construct the values $\hat{q}_{it}^{0}(\mathbf{X})$ for any value of $\mathbf{X}$ in the support $\mathcal{X}$. To construct $\hat{q}_{it}^{0}(\mathbf{X})$ we need the integrated value function at $t+1$ that we have obtained before:

$$\hat{q}_{it}^{0}(\mathbf{X}) = \Lambda_i^{-1}(\hat{P}_{it}^{0}(\mathbf{X})) - \beta_i \sum_{\mathbf{X}_t} \left[f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|1,\mathbf{X}) - f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|0,\mathbf{X})\right] \hat{V}_{it+1}^{\mathbf{B}}(\mathbf{X}_{t+1}) \tag{27}$$

Given $\hat{q}_{it}^{0}$, we can estimate payoffs and beliefs using the same type of expression as for period T:

$$\begin{aligned}
\hat{\pi}_{it}^{0}(0,S_i,\mathbf{W}) &= \hat{q}_{it}^{0}(\mathbf{X}^{low}) - \hat{P}_{jt}^{0}(\mathbf{X}^{high}) \left[\frac{\hat{q}_{it}^{0}(\mathbf{X}^{high}) - \hat{q}_{it}^{0}(\mathbf{X}^{low})}{\hat{P}_{jt}^{0}(\mathbf{X}^{high}) - \hat{P}_{jt}^{0}(\mathbf{X}^{low})}\right] \\
\hat{\pi}_{it}^{0}(1,S_i,\mathbf{W}) &= \hat{q}_{it}^{0}(\mathbf{X}^{low}) + \left[1 - \hat{P}_{jt}^{0}(\mathbf{X}^{high})\right] \left[\frac{\hat{q}_{it}^{0}(\mathbf{X}^{high}) - \hat{q}_{it}^{0}(\mathbf{X}^{low})}{\hat{P}_{jt}^{0}(\mathbf{X}^{high}) - \hat{P}_{jt}^{0}(\mathbf{X}^{low})}\right]
\end{aligned} \tag{28}$$

and

$$\hat{B}_{jt}^{0}(\mathbf{X}) = \frac{\hat{q}_{it}^{0}(\mathbf{X}) - \hat{\pi}_{it}^{0}(0,\mathbf{X})}{\hat{\pi}_{it}^{0}(1,\mathbf{X}) - \hat{\pi}_{it}^{0}(0,\mathbf{X})} \tag{29}$$

Finally, we construct the estimator of the integrated value function at period t.

$$\hat{V}_{it}^{\mathbf{B}}(\mathbf{X}) = \beta_i \sum_{\mathbf{X}_t} f_{it}^{\mathbf{B}}(\mathbf{X}_{t+1}|0,\mathbf{X})]\ \hat{V}_{it+1}^{\mathbf{B}}(\mathbf{X}_{t+1}) - \ln\left(1 - \hat{P}_{it}^{0}(\mathbf{X})\right) \tag{30}$$

It is straightforward to show that this estimator is consistent and asymptotically normal. The derivation of the asymptotic variance is cumbersome. In our empirical application we use the bootstrap method to obtain standard errors and confidence intervals for the estimates.

4.3 Estimation with parametric payoff function

In most applications, we assume a parametric specification of the payoff function. A very common class of parametric specifications is the linear in parameters model:

$$\pi_{it}(Y_{jt},\mathbf{X}_t) = z_{it}(Y_{jt},\mathbf{X}_t)\ \boldsymbol{\theta}_i \tag{31}$$

where $z_{it}(Y_{jt},\mathbf{X}_t)$ is a $1 \times K$ vector of known functions, and $\boldsymbol{\theta}_i$ is a $K \times 1$ vector of unknown structural parameters in player i's payoff function. Let $\boldsymbol{\theta}^0$ be $\{\boldsymbol{\theta}_i^0 : i = 1,2\}$. For instance, in the dynamic game of market entry and exit in the Example of section 2, the profit function in equation

(2) can be written as $\pi_i\left(Y_{jt}, S_{it}, W_t\right) = z_{it}(Y_{jt}, S_{it}, W_t)\ \boldsymbol{\theta}_i$, where the vector of parameters $\boldsymbol{\theta}_i$ is $(\theta_i^M,\ \theta_i^D,\ \theta_{i0}^{FC},\ \theta_{i1}^{FC},\ \theta_i^{EC})'$ and

$$z_{it}(Y_{jt}, S_{it}, W_t) = \{\ W_t(1 - Y_{jt}),\ W_t Y_{jt},\ -1,\ -\exp\{-S_{it}\}, -1\{S_{it} = 0\}\ \} \tag{32}$$

Given this specification, the model implies the following relationship:

$$q_{it}^0(\mathbf{X}) = z_{it}^{\mathbf{B}}\ (\mathbf{X})\ \boldsymbol{\theta}_i^0 \tag{33}$$

where $z_{it}^{\mathbf{B}}\ (\mathbf{X}) \equiv (1 - B_{jt}^0(\mathbf{X}))\ z_{it}(0, \mathbf{X}) + B_{jt}^0(\mathbf{X})\ z_{it}(1, \mathbf{X})$, and $\boldsymbol{\theta}_i^0$ is the true vector of parameters in the population.

To estimate $\boldsymbol{\theta}_i^0$ we propose a simple three steps method. The first two-steps are the same as for the nonparametric model.

Step 3: Given the estimates from step 2, we can apply a pseudo maximum likelihood method in the spirit of Aguirregabiria and Mira (2002, 2007) to estimate the structural parameters $\boldsymbol{\theta}^0$. Define the pseudo likelihood function:

$$Q(\boldsymbol{\theta}, \mathbf{B}, \mathbf{P}) \equiv \sum_{m=1}^{M} \sum_{t=1}^{T} \sum_{i=1}^{2} Y_{imt}\ \log \Lambda\left(\tilde{z}_{imt}^{\mathbf{B},\mathbf{P}}\ \boldsymbol{\theta}_i + \tilde{e}_{imt}^{\mathbf{B},\mathbf{P}}\right) + (1 - Y_{imt})\ \log\left(1 - \Lambda\left(\tilde{z}_{imt}^{\mathbf{B},\mathbf{P}}\ \boldsymbol{\theta}_i + \tilde{e}_{imt}^{\mathbf{B},\mathbf{P}}\right)\right)$$

$\tilde{z}_{imt}^{\mathbf{B},\mathbf{P}}$ is the sum of expected and discounted stream of $\{z_{it'}(Y_{jt'}, \mathbf{X}_{t'}) : t' = t,\ t+1,\ ...,\ T\}$ given that player i behaves according to the choice probabilities $P_{it'}(.)$ in $\mathbf{P}$, and player j behaves according to the probabilities $B_{jt'}(.)$ in $\mathbf{B}$. Similarly, $\tilde{e}_{imt}^{\mathbf{B},\mathbf{P}}$ is the sum of expected and discounted stream of $\{e(P_{it'}(\mathbf{X}_{t'})) : t' = t,\ t+1,\ ...,\ T\}$, and for the logit model $e(P_{it'}(\mathbf{X}_{t'})) = \gamma - \ln P_{it'}(\mathbf{X}_{t'})$ where γ is Euler's constant. From steps 1 and 2, we have consistent estimates of CCPs, $\hat{\mathbf{P}}^0$, and beliefs, $\hat{\mathbf{B}}^0$. Then, a consistent pseudo maximum likelihood estimator of $\boldsymbol{\theta}^0$ is defined as the value $\hat{\boldsymbol{\theta}}^{(1)}$ that maximizes $Q(\boldsymbol{\theta}, \hat{\mathbf{B}}^0, \hat{\mathbf{P}}^0)$. Note that the sample criterion function $Q(\boldsymbol{\theta}, \hat{\mathbf{B}}^0, \hat{\mathbf{P}}^0)$ is just the log likelihood function of a standard logit model with the restriction that the parameter of variable $\tilde{e}_{imt}^{\mathbf{B},\mathbf{P}}$ is equal to 1. The estimator is root-M consistent and asymptotically normal.

Given this parametric estimator of the payoff function, $\hat{\pi}_{it}(Y_{jt}, \mathbf{X}_t) = z_{it}(Y_{jt}, \mathbf{X}_t)\ \hat{\boldsymbol{\theta}}_i^{(1)}$, that is more precise than the original nonparametric estimator, we can also obtain a more precise estimator of players' CCPs and of players' beliefs. The new estimator of CCPs is based on the best response condition $P_{it}^0(\mathbf{X}) = \Lambda\left(\tilde{z}_{it}^{\mathbf{B}\ ,\mathbf{P}}\ (\mathbf{X})\ \boldsymbol{\theta}_i^0 + \tilde{e}_{it}^{\mathbf{B}\ ,\mathbf{P}}\ (\mathbf{X})\right)$ and it is equal to:

$$\hat{P}_{it}^{(1)}(\mathbf{X}) = \Lambda\left(\tilde{z}_{it}^{\hat{\mathbf{B}}\ ,\hat{\mathbf{P}}}\ (\mathbf{X})\ \hat{\boldsymbol{\theta}}_i^{(1)} + \tilde{e}_{it}^{\hat{\mathbf{B}}\ ,\hat{\mathbf{P}}}\ (\mathbf{X})\right) \tag{34}$$

where $\hat{\mathbf{B}}^0$ and $\hat{\mathbf{P}}^0$ represent the initial nonparametric estimators of beliefs and CCPs, respectively, and $\hat{\boldsymbol{\theta}}_i^0$ is the pseudo maximum likelihood estimator of the parameters in the payoff function. The new estimator of beliefs is based on the following expression for beliefs in terms of CCPs and q functions that is derived as an implication of the model:[16]

$$B_{jt}^0(\mathbf{X}) = P_{jt}^0(\mathbf{X}^{low}) + \left[P_{jt}^0(\mathbf{X}^{high}) - P_{jt}^0(\mathbf{X}^{low})\right] \left[\frac{q_{it}^0(\mathbf{X}) - q_{it}^0(\mathbf{X}^{low})}{q_{it}^0(\mathbf{X}^{high}) - q_{it}^0(\mathbf{X}^{low})}\right] \tag{35}$$

By construction, this expression is satisfied exactly by the initial nonparametric estimators of B_{jt}^0, P_{jt}^0, and q_{it}^0. Once we update our estimators of CCPs and of q functions to take into account the parametric restrictions, the initial estimator of beliefs no longer satisfies these restrictions. Therefore, we update our estimates of beliefs accordingly. The new estimator is:

$$\hat{B}_{jt}^{(1)}(\mathbf{X}) = \hat{P}_{jt}^{(1)}(\mathbf{X}^{low}) + \left[\hat{P}_{jt}^{(1)}(\mathbf{X}^{high}) - \hat{P}_{jt}^{(1)}(\mathbf{X}^{low})\right] \left[\frac{\hat{q}_{it}^{(1)}(\mathbf{X}) - \hat{q}_{it}^{(1)}(\mathbf{X}^{low})}{\hat{q}_{it}^{(1)}(\mathbf{X}^{high}) - \hat{q}_{it}^{(1)}(\mathbf{X}^{low})}\right] \tag{36}$$

where $\hat{P}_{jt}^{(1)}$ is the new parametric estimator of CCPs, and $\hat{q}_{it}^{(1)}$ is the new parametric estimator of the q_{it}^0 function that is $\hat{q}_{it}^{(1)}(\mathbf{X}) = [(1 - \hat{B}_{jt}^0(\mathbf{X}))\ z_{it}(0, \mathbf{X}) + \hat{B}_{jt}^0(\mathbf{X})\ z_{it}(1, \mathbf{X})]\ \hat{\boldsymbol{\theta}}_i^{(1)}$.

We can apply step 3 and the updating of CCPs and beliefs recursively to obtain a sequence of estimators $\{\hat{\boldsymbol{\theta}}^{(K)}, \hat{\mathbf{B}}^{(K)}, \hat{\mathbf{P}}^{(K)} : K \geq 1\}$. At each iteration K of this iterative procedure we perform tasks (i) to (v):

(i) Update of preference parameters: $\hat{\boldsymbol{\theta}}^{(K)} = \arg\max_{\boldsymbol{\theta}}\ Q(\boldsymbol{\theta}, \hat{\mathbf{B}}^{(K-1)}, \hat{\mathbf{P}}^{(K-1)})$.
(ii) Update of CCP functions: $\hat{P}_{it}^{(1)}(\mathbf{X}) = \Lambda \left(\tilde{z}_{it}^{\hat{\mathbf{B}}\ K-\ ,\hat{\mathbf{P}}\ K-}(\mathbf{X})\ \hat{\boldsymbol{\theta}}_i^{(K)} + \tilde{e}_{it}^{\hat{\mathbf{B}}\ K-\ ,\hat{\mathbf{P}}\ K-} \right)$.
(iii) Update of value functions: $\hat{V}_{iT}^{(K)}(\mathbf{X}) = -\ln\left(1 - \hat{P}_{iT}^{(K)}(\mathbf{X})\right)$ and for $t < T$, $\hat{V}_{it}^{(K)}(\mathbf{X}) = -\ln\left(1 - \hat{P}_{it}^{(K)}(\mathbf{X})\right) + \beta_i \sum_{\mathbf{X}_t}\ f_{it}(\mathbf{X}_{t+1}|0, \mathbf{X})\ \hat{V}_{it+1}^{(K)}(\mathbf{X})$.
(iv) Update of q functions: $\hat{q}_{iT}^{(K)}(\mathbf{X}) = \Lambda^{-1}(\hat{P}_{it}^{(K)}(\mathbf{X})) - \beta_i \sum_{\mathbf{X}_t}\ [f_{it}(\mathbf{X}_{t+1}|1, \mathbf{X}) - f_{it}(\mathbf{X}_{t+1}|0, \mathbf{X})]\ \hat{V}_{it+1}^{(K)}(\mathbf{X}_{t+1})$.
(v) Update of beliefs functions:

$$\hat{B}_{jt}^{(K)}(\mathbf{X}) = \hat{P}_{jt}^{(K)}(\mathbf{X}^{low}) + \left[\hat{P}_{jt}^{(K)}(\mathbf{X}^{high}) - \hat{P}_{jt}^{(K)}(\mathbf{X}^{low})\right] \left[\frac{\hat{q}_{it}^{(K)}(\mathbf{X}) - \hat{q}_{it}^{(K)}(\mathbf{X}^{low})}{\hat{q}_{it}^{(K)}(\mathbf{X}^{high}) - \hat{q}_{it}^{(K)}(\mathbf{X}^{low})}\right]$$

[16] As shown in equation (12) above, for any vectors $\mathbf{X}^a$, $\mathbf{X}^b$, $\mathbf{X}^c$, and $\mathbf{X}^d$, we have that $\frac{q_i^0\ \mathbf{X}^a\ - q_i^0\ \mathbf{X}^b}{q_i^0\ \mathbf{X}^c\ - q_i^0\ \mathbf{X}^d} = \frac{B_j^0\ \mathbf{X}^a\ - B_j^0\ \mathbf{X}^b}{B_j^0\ \mathbf{X}^c\ - B_j^0\ \mathbf{X}^d}$. If we particularize this expression at the vectors $\mathbf{X}^a = \mathbf{X}$, $\mathbf{X}^b = \mathbf{X}^{low}$, $\mathbf{X}^c = \mathbf{X}^{high}$, and $\mathbf{X}^d = \mathbf{X}^{low}$, we have that $\frac{q_i^0\ \mathbf{X}\ - q_i^0\ \mathbf{X}^{low}}{q_i^0\ \mathbf{X}^{high}\ - q_i^0\ \mathbf{X}^{low}} = \frac{B_j^0\ \mathbf{X}\ - B_j^0\ \mathbf{X}^{low}}{B_j^0\ \mathbf{X}^{high}\ - B_j^0\ \mathbf{X}^{low}}$. By Assumption 5b, $B_j(\mathbf{X}^{low}) = P_j(\mathbf{X}^{low})$ and $B_j(\mathbf{X}^{high}) = P_j(\mathbf{X}^{high})$. Thus, solving for $B_j(\mathbf{X})$ we obtain $B_j(\mathbf{X}) = P_j(\mathbf{X}^{low}) + \left[P_j(\mathbf{X}^{high}) - P_j(\mathbf{X}^{low})\right] \left[\frac{q_i(\mathbf{X}) - q_i(\mathbf{X}^{low})}{q_i(\mathbf{X}^{high}) - q_i(\mathbf{X}^{low})}\right]$.

5 Empirical Application

We illustrate our model and methods with an application of a dynamic game of store location. There has been recently a significant interest in the estimation of game theoretic models of market entry and store location by retail firms. Most studies have assumed static games: see Mazzeo (2002), Seim (2006), Jia (2008), Nishida (2008), and Zhu and Singh (2009), among others. Holmes (2010) estimates a single-agent dynamic model of store location by Wal-Mart. Beresteanu and Ellickson (2005), Walrath (2008), and Suzuki (2010) propose and estimate dynamic games of store location.

We study store location of McDonalds (MD) and Burger King (BK) using data for the United Kingdom during the period 1991-1995. We divide the UK in local markets (districts) and study these companies' decision of how many stores, if any, to operate in each local market. The profits of a store in a market depends on local demand and cost conditions and on the degree of competition from other firms' stores and from stores of the same chain. There are sunk costs associated with opening a new store, and therefore this decision has implications for future profits. Firms are forward-looking and maximize the value of expected and discounted profits. Each firm has uncertainty about future demand and cost conditions in local markets. Firms also have uncertainty about the current and future behavior of the competitor. In this context, the standard assumption is that firms have rational expectations about other firms' strategies, and that these strategies constitute a Markov Perfect Equilibrium. Here we relax this assumption. The main question that we want to analyze in this empirical application is whether the beliefs of each of these companies about the store location strategy of the competitor are consistent with the actual behavior of the competitor.

5. Data and descriptive evidence

The dataset was collected by Otto Toivanen and Michael Waterson, who have used it in their paper Toivanen and Waterson (2005).[17] Our working sample is a five year panel that tracks 422 *local authority districts* (local markets), including the information on the stock and flow of MD and BK stores into each district. It also contains socio-economic variables at the district level such as population, density, age distribution, average rent, income per capita, local retail taxes, and distance to the UK headquarters of each of the firms. The *local authority district* is the smallest unit of local government in the UK, and generally consists of a city or a town sometimes with a

[17] We want to thank Otto Toivanen and Michael Waterson for generously sharing their data with us.

surrounding rural area. There are almost 500 *local authority districts* in Great Britain. Our working sample of 422 districts does not include those that belong to Greater London.[18] The median district in our sample has an area of 300 square kilometers and a population of 95,000 people.[19] Table 1 presents descriptive statistics for socio-economic and geographic characteristics of our sample of local authority districts.

Table 2 presents descriptive statistics on the evolution of the number of stores for the two firms. Toivanen and Waterson present a detailed discussion of why the retail chain fast food hamburger industry in the UK during this period can be assumed as a duopoly of BK and MD. In 1990, MD had more than three times the number of stores of BK, and it was active in more than twice the number of local markets than BK. Conditional on being active in a local market, MD had also significantly more stores per market than BK. These differences between MD and BK have not declined significantly over the period 1991-1995. While BK has entered in more new local markets than MD (69 new markets for BK and 48 new markets for MD), MD has opened more stores (143 new stores for BK and 166 new stores for MD).

Table 3 presents the annual transition probabilities of market structure in local markets as described by the number of stores of the two firms. According to this transition matrix, opening a new store is an irreversible decision, i.e., no store closings are observed during this sample period. In Britain during our sample period, the fast food hamburger industry was still young and expanding, as shown by the large proportion of observations/local markets without stores (41.6%). Although there is significant persistence in every state, the less persistent market structures are those where BK is the leader. For instance, if the state is "$BK = 1$ & $MD = 0$", there is a 20% probability that the next year MD opens at least one store in the market. Similarly, when the state is "$BK = 2$ & $MD = 1$", the chances that MD opens one more store the next year are 31%.

Table 4 presents estimates of reduced form Probit models for the decision to open a new store. We obtain separate estimates for MD and BK. Our main interest is in the estimation of the effect of the previous year's number of stores (own stores and competitor's stores) on the probability of opening a new store. We include as control variables population, GDP per capita, population density, proportion of population 5-14, proportion population 15-29, average rent, and proportion of claimants of unemployment benefits. To control for unobserved local market heterogeneity we also

[18] The reason we exclude from our sample the districts in Greater London is that they do not satisfy the standard criteria of isolated geographic markets.

[19] As a definition of geographic market for the fast food retail industry, the district is perhaps a bit wide. However, an advantage of using district as definition of local market is that most of the markets in our sample are geographically isolated. Most districts contain a single urban area. And, in contrast to North America where many fast food restaurants are in transit locations, in UK these restaurants are mainly located in the centers of urban areas.

present two fixed effects estimations, one with county fixed effects and other with local district fixed effects. We only report estimates of the marginal effects associated with the dummy variables that represent previous year number of stores. The main empirical result from Table 4 is that, regardless of the set of control variables that we use, the own number of stores has a strong negative effect on the probability of opening a new store but the effect of the competitor's number of stores is either negligible or even positive. This finding is very robust to different specifications of the reduced form model and it is analogous to the result from the reduced form specifications in Toivanen and Waterson's paper. Controlling for unobserved heterogeneity using fixed effects reveals that the estimation without fixed effects suffers from a significant upward bias in the marginal effect of the number of own stores. However, the estimated marginal effect of the number of competitor's stores barely changes. The estimates show also a certain asymmetry between the two firms: the absence of response to the competitor's number of stores is more clear for BK than for MD. In particular, when BK has three stores in the market there is a significant reduction in MD's probability of opening a new store. That negative effect does not appear in the reduced form probit for BK.

This empirical evidence cannot be explained by a standard static model of store location by firms that sell substitute products. Here we explore three, non-mutually exclusive, explanations: (a) spillover effects; (b) forward looking behavior (dynamic game); and (c) biased beliefs about the behavior of the competitor.

(a) Spillover effects. The competitor's number of stores may have a positive spillover effect on the profit of a firm. There are several possible sources of this spillover effect. For example one firm may infer from another's decision to open a store in a particular market that market conditions are favorable (informational spillovers). Alternatively, one firm may benefit from another firm's entry through cost reductions, or from product expansion through advertising. Since we do not have price and quantity data at the level of local markets, we do not try to identify the source of the spillover effect. We include this effect in our specification of demand such that the natural interpretation in the context of our model is a product expansion coming from the advertising effect of retail stores. However, this should be interpreted as a shortcut or 'reduced form' specification of different possible spillover effects.

(b) Forward looking behavior. Opening a store is a partly irreversible decision that involves a significant sunk costs. Therefore, it is reasonable to consider that firms are forward looking when they make this decision. Moreover, dynamic strategic effects may help explain the apparent lack of competitive effects when we look at these decisions from the point of view of a static model of

entry. Suppose that firms anticipate, with some uncertainty, the total number of hamburger stores that a local market can sustain in the long-run given the size and the socioeconomic characteristics of the market. For simplicity, suppose that this number of "available slots" does not depend on the ownership of the stores because the products sold by the two firms are very close substitutes. In this context, firms play a 'racing' game to fill as many 'slots' as possible with their own stores. Diseconomies of scale and scope may generate a negative effect of the own number of stores on the decision of opening new stores. However, in this model, during most of the period of expansion the number of slots of the competitor has zero effect on the decision of opening a new store. Only when the market is filled or close to being filled do the competitor's stores have a significant effect on entry decisions.

(c) Biased beliefs. As mentioned in the Introduction, competition in actual oligopoly industries is often characterized by strategic uncertainty. Firms face significant uncertainty about the strategies of their competitors. In the context of our application, it may be the case that MD's or/and BK's beliefs overestimate the negative effect of the competitor's stores on the competitor's entry decisions. For instance, if MD has one store in a local market, BK may believe that the probability that MD opens a second store is close to zero. These over-optimistic beliefs about the competitor's behavior may generate an apparent lack of response of BK's entry decisions to the number of MD's stores.

5.2 Model

Consider two retail chains competing in a local market. Each firm sells a differentiated product using its stores. Let $Y_{it} \in \{0, 1, ..., K\}$ be the number of stores of firm i at period t. We abstract from store location within a local market and assume that every store of the same firm has the same demand. Let q_{it} be the quantity sold by all the stores of firm i. The demand for all the stores of firm i is:

$$q_{it} = \begin{cases} W_t \;\; (1 + b_{it} - p_{it}) & \text{if} \quad Y_{jt} = 0 \\ \\ W_t \left(\dfrac{Y_{it}}{Y_{it} + Y_{jt}} \right) (1 + [b_{it} + \delta_i \; b_{jt}] - [p_{it} - p_{jt}]) & \text{if} \quad Y_{jt} > 0 \end{cases} \tag{37}$$

where: W_t is a measure of market size and it is exogenous; b_{it} is the 'quality' of product i at period t;[20] p_{it} is the 'price' of product i at period t; and $\delta_i \geq -1$ is a parameter that captures the net effect of the quality of firm j on the demand of firm i. This net effect is positive ($\delta_i > 0$) if the spillover

[20] This 'quality' is just the willingness to pay for the product of the average consumer in the market.

effect from an increase in the quality of the competitor is larger than the competitive effect, and it is negative ($\delta_i < 0$) otherwise. When there is not any positive spillover effect we have that $\delta_i = -1$. In that case, if the two firms have the same qualities and prices, they share the market size W_t in proportion to their number of stores. A firm with better quality or/and lower price can get a larger proportion of the market. Production costs are linear in the quantity produced, i.e., $C_{it} = c_i \; q_{it}$, where c_i is firm i's constant marginal cost. The variable profit of firm i is $VP_{it} = (p_{it} - c_i)q_{it}$. Given market size and qualities at period t, firms compete in prices ala Bertrand to maximize current variable profit.

If firm i is a monopolist in the market (i.e., $Y_{jt} = 0$), then the profit-maximizing price is $p_{it} = c_i + 1 + b_{it}$ and the variable profit is $VP_{it} = W_t \left([1 + b_{it} - c_i]/2\right)^2$. If both firms are active in the market (i.e., $Y_{it} > 0$ and $Y_{jt} > 0$), then the Bertrand equilibrium price is[21] $p_{it} = c_i + 1 + [b_{it} + \delta_i b_{jt}] - (1/3)(\Delta c + \Delta B_t)$, where $\Delta B_t \equiv [b_{it} + \delta_i b_{jt}] - [b_{jt} + \delta_j b_{it}]$ and $\Delta c \equiv c_i - c_j$, and the equilibrium variable profit is $VP_{it} = W_t \; Y_{it}(Y_{it} + Y_{jt})^{-1} \; (1 + [b_{it} + \delta_i b_{jt}] - (1/3)(\Delta c + \Delta B_t))^2$.

The visibility of a retail firm in a local market increases with its number of stores. We assume that a firm's quality increases with the number of stores that the firm has in the market. There are at least two ways in which the number of stores in the market affects the willingness to pay of the average consumer. First, an increase in the number of stores implies a reduction in consumer transportation cost to visit a store of the chain. Second, stores are like 'advertisements' in the sense that they increase the awareness of local consumers about the retail chain. We assume the following specification:

$$b_{it} = b_i^{(0)} + b_i^{(1)} \; Y_{it} \tag{38}$$

where $b_i^{(0)} \geq 0$ is a parameter that represents the 'exogenous' quality of firm i in every local market.[22] And $b_i^{(1)} \geq 0$ is a parameter that measures the effect of the number of stores in the local market on the 'quality' of the firm in that market.

Firm i is active in the market at period t if Y_{it} is strictly positive. In order to distinguish decision and state variables, we use the variable S_{it} to represent the number of stores at period $t-1$, i.e., $S_{it} \equiv Y_{it-1}$. Every period, the two firms know the 'stocks' of stores in the market, S_{it}

[21] The first order condition of profit maximization for firm i is $p_i - c_i = B_i - \Delta p$, where $B_i \equiv 1 + b_i + \delta_i b_j$, and $\Delta p \equiv p_i - p_j$. Similarly, the first order condition of profit maximization for firm j is $p_j - c_j = B_j + \Delta p$. The difference between these first two conditions implies that $\Delta p = (1/3)(\Delta c + \Delta B)$ where $\Delta B \equiv B_i - B_j$ and $\Delta c \equiv c_i - c_j$. Solving this formula for Δp into the first order condition of firm i, we get the following expression for the profit-maximizing price: $p_i = c_i + B_i - (1/3)(\Delta c + \Delta B)$.

[22] By 'exogenous' quality here we mean that it is the part of quality that does not depend on the number of stores in the market.

and S_{jt}, and choose simultaneously the new number of stores. Firm i's total profit function is:

$$\begin{aligned} \Pi_{it} &= VP_{it} - 1\{Y_{it} > 0 \text{ and } S_{it} = 0\}\ \theta_i^{EC} \\ &- 1\{Y_{it} > 0\}\ \left[\theta_{0i}^{FC} - \theta_{1i}^{FC} Y_{it} - \theta_{2i}^{FC} Y_{it}^2\right] \\ &- 1\{Y_{it} > S_{it}\}\ \varepsilon_{it} \end{aligned} \tag{39}$$

where $1\{.\}$ is the indicator function, and θ_i^{EC}, θ_{0i}^{FC}, θ_{1i}^{FC} and θ_{2i}^{FC} are parameters in the cost function. θ_{0i}^{EC} is an entry cost that is paid the first time that the firm opens a store in the local market. θ_{0i}^{FC} is a lump-sum cost associated with having any positive number of stores in the market. The function $\theta_{1i}^{FC} Y_{it} + \theta_{2i}^{FC} Y_{it}^2$ takes into account that operating costs increase with the number of stores in a linear or quadratic form. The variable ε_{it} is a private information shock in the cost of opening a new store, and it is i.i.d. normally distributed.

Note that our model implies the exclusion restriction that, given Y_{jt}, the profit of firm i does not depend on $S_{jt} = Y_{jt-1}$. That is, a firm's current profit depends on his own and his opponents current number of stores in the market, but given these variables it does not depend on the number of stores of the competitors at period $t-1$. Of course a firm's beliefs about the probability distribution of Y_{jt} depends on S_{jt}.

5.3 Estimation of the structural model

As described in section 5.1 above, we do not observe store closings in our sample. Furthermore, for almost all the observations with store openings the number of new stores is one. Therefore, we assume that $Y_{imt} \in \{S_{imt}, S_{imt}+1\}$ or equivalently, $Y_{imt} - S_{imt} \in \{0, 1\}$. Variable $Y_{imt} - S_{imt}$ is the binary indicator of the event "firm i opens a new store in market m at year t". The maximum value of S_{imt} in the sample is 13, and we assume that the set of possible values of S_{imt} is $\{0, 1, ..., 15\}$. Therefore, the state space $\mathcal{X}$ is $\{0, 1, ..., 15\} \times \{0, 1, ..., 15\}$ that has 256 grid points. . We assume that market characteristics are constant over time. The measure of market size W_m is total population in the district. For some specifications, we allow the cost of investment to depend on market characteristics such as average rent, retail taxes, population density, or average income. Therefore, each market has its own vector of players' CCPs. The dimension of the vectors $\mathbf{P}_i$ in this model is equal to $108,032$, i.e., 422 markets times 256 states.

Tables 5 and 6 present estimates of the structural model under the assumption that firms are myopic, $\beta = 0$, and under the assumption that firms are forward looking, $\beta = 0.95$, respectively. We report two different sets of point estimates: estimates using a simple two-step Pseudo Maximum Likelihood method where the estimator of (equilibrium) players' beliefs in the first step is a non-

parametric frequency estimator; and estimates using the Nested Pseudo Likelihood (NPL) method proposed in Aguirregabiria and Mira (2007). The NPL method imposes the equilibrium restrictions in the sample (i.e., the estimated beliefs should be equal to the estimated best response probabilities), while the two-step method only satisfies the equilibrium restrictions asymptotically. The NPL estimator has smaller asymptotic variance and finite sample bias than the two-step method. There are very substantial differences between two models, particularly in the estimates of the parameters that capture cannibalization and competition effects. While these effects have the 'wrong' sign in the myopic model, the signs are the expected ones in the forward looking model. All the parameter estimates in the forward looking model have the expected signs and have reasonable magnitudes. Therefore, it seems that forward looking behavior explains part of the puzzle in the reduced form estimates.

Table 7 presents results of our test of equilibrium beliefs. We implement separate tests for MD and BK. We impose the restrictions that beliefs for $S_{jt} = 0$ and $S_{jt} = 3$ are unbiased.

6 Conclusion

This paper studies a class of dynamic games of incomplete information where players' beliefs about the other players' actions may not be in equilibrium. We present new results on identification, estimation, and inference of structural parameters and beliefs in this class of games when the researcher does not have data on elicited beliefs. Specifically, we derive sufficient conditions under which payoffs and beliefs are point identified. These conditions then lead naturally to a two-step estimator of payoffs and beliefs, which we show can be extended to provide a sequence of estimators with asymptotic variances and finite sample biases that decline monotonically. We also present a procedure for testing the null hypothesis that beliefs are in equilibrium. We illustrate our model and methods with an empirical application of a dynamic game of store location by McDonalds and Burger King. Our main interest in this application is to explain a puzzling empirical regularity, that the probability a firm opens a new store in a local market depends negatively on the number of stores it currently has open in the location, and does not depend (or may even positively depend) on the number of stores its competitor currently has open in the location. In the context of our model we explore three alternative (but not mutually exclusive) explanations for these: cross-firm spillovers, forward looking behavior, and out of equilibrium (i.e., biased) beliefs. We find empirical evidence for the hypotheses of forward looking behavior and biased beliefs.

References

[1] Aguirregabiria, (2010): "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior," *Journal of Business and Economic Statistics,* 28(2), 201–218.

[2] Aguirregabiria, V. and P. Mira (2007) "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75, 1–53.

[3] Aguirregabiria, V. and P. Mira (2009) "Structural Estimation and Counterfactual Experiments in Games when the Data Come from Multiple Equilibria," manuscript. Department of Economics. University of Toronto.

[4] Aradillas-Lopez, A. (2009): "Pairwise-Difference Estimation of Incomplete Information Games," Working Paper, Princeton University.

[5] Aradillas-Lopez, A. and E. Tamer (2008): "The Identification Power of Equilibrium in Games," *Journal of Business and Economic Statistics*, 26 (3), 261-283.

[6] Bajari, P., L. Benkard and J. Levin (2006): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331-1370.

[7] Bajari, P. and H. Hong (2005): "Semiparametric Estimation of a Dynamic Game of Incomplete Information. Manuscript," Duke University.

[8] Bajari, P., H. Hong, J. Kreiner, and D. Nekipelov (2009): "Estimating Static Models of Strategic Interaction," Journal of Business and Economic Statistics, forthcoming.

[9] Beresteanu, A. and P. Ellickson (2005): "The Dynamics of Retail Oligopolies," manuscript. Duke University.

[10] Besanko, D., and U. Doraszelski (2004): "Capacity Dynamics and Endogenous Asymmetries in Firm Size," *RAND Journal of Economics,* 35, 23-49.

[11] Besanko, D., U. Doraszelski, L. Lu, and M. Satterthwaite (2010): "On the role of demand and strategic uncertainty in capacity investment and disinvestment dynamics," *International Journal of Industrial Organization,* 28(4), 383-389.

[12] Camerer, C. (2003): "Behavioral Game Theory: Experiments in Strategic Interaction," Princeton University Press: Princeton NJ.

[13] Camerer, C. F., T-H Ho, and J-K Chong (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119, 861-898.

[14] Goldfarb, A., and M. Xiao (2011): "Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets," *American Economic Review*, forthcoming.

[15] Holmes, T. (2010): "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, forthcoming.

[16] Jia, P. (2008): "What Happens When Wal-Mart Comes to Town: An empirical analysis of the Discount Industry," *Econometrica*, 76, 1263-1316.

[17] Kasahara, H. and K. Shimotsu (2009): "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, 77, 135-175.

[18] Khan, S., and E. Tamer (2009): "Irregular Identification, Support Conditions, and Inverse Weight Estimation," manuscript. Duke University and Northwestern University.

[19] Magnac, T. and Thesmar, D. (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801-816.

[20] Matzkin, R. (1992): "Nonparametric and Distribution-Free Estimation of the Binary Choice and the Threshold Crossing Models", *Econometrica*, 60, 239-270.

[21] Matzkin, R. (1994): "Restrictions of Economic Theory in Nonparametric Methods," *Handbook of Econometrics*, Vol. 4, edited by C.F. Engel and D.L. McFadden, Elsevier.

[22] Mazzeo, M. (2002): "Product choice and oligopoly market structure," *Rand Journal of Economics*, 33(2), 221–242.

[23] Morris. S., and H. Song (2002): "Measuring Strategic Uncertainty," manuscript. Department of Economics. Princeton University.

[24] Nishida, M. (2008): "Estimating a model of strategic store network choice," Working Paper, University of Chicago.

[25] Pesendorfer, M. and Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games," *The Review of Economic Studies*, 75, 901-928.

[26] Samuelson, P. (1938): "A Note on the Pure Theory of Consumers' Behaviour," *Economica,* 5, 61-71.

[27] Seim, K. (2006): "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," *The RAND Journal of Economics*, 37(3), 619–640.

[28] Suzuki, J. (2010): "Land Use Regulation as a Barrier to Entry: Evidence from the Texas Lodging Industry," manuscript. Department of Economics. University of Toronto.

[29] Tamer, E. (2003): "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria," *Review of Economic Studies*, 70, 147-165.

[30] Toivanen, O., and M. Waterson (2005): "Market Structure and Entry: Where's the Beef?," RAND Journal of Economics, 36(3), 680-699.

[31] Walrath, M. (2008): "Religion as an industry: estimating a strategic entry model for churches," manuscript. Department of Economics. University of St. Thomas.

[32] Zhu, T. and V. Singh (2009): "Spatial competition with endogenous location choices: An application to discount retailing," *Quantitative Marketing and Economics*, 7, 1–35.

Table

Descriptive Statistics on Local Markets (Year 99)

422 local authority districts (excluding Greater London districts)

Variable	Median	Std. Dev.	Pctile 5%	Pctile 95%
Area (thousand square km)	0.30	0.73	0.03	1.67
Population (thousands)	94.85	93.04	37.10	280.50
Children: Age 5- 4 (%)	12.43	1.00	10.74	14.07
Young: 5-29 (%)	21.24	2.46	17.80	25.17
Pensioners: 65-74 (%)	9.01	1.50	6.89	11.82
GDP per capita (thousand £)	92.00	12.14	74.40	112.70
Claimants of UB / Population ratio (%)	2.75	1.27	1.24	5.11
Avg. Weekly Rent per dwelling (£)	25.31	10.61	19.11	35.07
Council tax (thousand £)	0.24	0.05	0.11	0.31
Number of BK stores	0.00	0.62	0.00	1.00
Number of MD stores	1.00	1.16	0.00	3.00

Table 2
Evolution of the Number of Stores
422 local authority districts (excluding Greater London districts)

	Burger King					
	990	**99**	**992**	**993**	**994**	**995**
#Markets with Stores	71	98	104	118	131	150
Change in #Markets with Stores	-	17	6	14	13	19
# of Stores	79	115	128	153	181	222
Change in # of Stores	-	36	13	25	28	41
Mean #Stores per Market (Conditional on #Stores>0)	1.11	1.17	1.23	1.30	1.38	1.48
	McDonalds					
	990	**99**	**992**	**993**	**994**	**995**
#Markets with Stores	206	213	220	237	248	254
Change in #Markets with Stores		7	7	17	11	6
# of Stores	281	316	344	382	421	447
Change in # of Stores		35	28	38	39	26
Mean #Stores per Market (Conditional on #Stores>0)	1.36	1.49	1.56	1.61	1.70	1.76

Table 3

Transition Probability Matrix for Market Structure

Annual Transitions. Market structure: BK=x & MD=y, where x and y are number of stores

	% **Market Structure at t+**								
	BK=0 MD=0	BK=0 MD=1	BK=0 MD$\geq$ 2	BK=1 MD=0	BK=1 MD=1	BK=1 MD$\geq$ 2	BK$\geq$ 2 MD=0	BK$\geq$ 2 MD=1	BK$\geq$ 2 MD$\geq$ 2
BK=0 &MD=0	**95.**	3.6	0.2	1.0	-	-	-	0.1	-
BK=0 &MD=1	-	**87.2**	4.2	-	7.4	1.0	-	-	1.4
BK=0 &MD$\geq$ 2	-	-	**82.7**	-	-	15.8	-	-	1.4
BK=1 &MD=0	-	-	-	**76.0**	18.0	2.0	4.0	-	-
BK=1 &MD=1	-	-	-	-	**87.**	8.1	-	3.3	1.4
BK=1 &MD$\geq$ 2	-	-	-	-	-	**86.5**	-	-	13.5
BK$\geq$ 2 &MD=0	-	-	-	-	-	-	**84.6**	15.4	-
BK$\geq$ 2 &MD=1	-	-	-	-	-	-	-	**69.0**	31.0
BK$\geq$ 2 &MD$\geq$ 2	-	-	-	-	-	-	-	-	**00.0**
Frequency	41.6	23.3	6.6	2.2	10.9	8.8	0.6	1.4	4.5

Table 4
Reduced Form Probits for the Decision to Open a Store

	Estimated Marginal Effects[1] ($\Delta P(x)$ when dummy goes from 0 to 1)					
	Burger King			McDonalds		
Explanatory Variable	**No FE**	**County FE**	**District FE**	**No FE**	**County FE**	**District FE**
Own number of stores at t-						
Dummy: Own #stores = 1	-0.021**	-0.036**	-0.885**	-0.035**	-0.045**	-0.550**
	(0.005)	(0.007)	(0.063)	(0.010)	(0.012)	(0.056)
Dummy: Own #stores = 2	-0.023**	-0.030**	-0.210*	-0.047**	-0.060*	-0.757**
	(0.004)	(0.005)	(0.085)	(0.006)	(0.008)	(0.041)
Dummy: Own #stores $\geq$ 3	-0.019**	-0.027**	-0.056	-0.043**	-0.053**	-0.816**
	(0.005)	(0.005)	(0.036)	(0.006)	(0.008)	(0.038)
Competitor's number of stores at t-						
Dummy: Comp.'s #stores = 1	0.032**	0.037*	-0.025	0.020	0.032*	0.052**
	(0.011)	(0.014)	(0.055)	(0.013)	(0.018)	(0.073)
Dummy: Comp.'s #stores = 2	0.045*	0.052*	-0.017	0.041	0.076	-0.007**
	(0.023)	(0.029)	(0.031)	(0.029)	(0.046)	(0.093)
Dummy: Comp.'s #stores $\geq$ 3	0.089*	0.101*	0.011	-0.041**	-0.050**	-0.104**
	(0.048)	(0.059)	(0.084)	(0.007)	(0.009)	(0.020)
Pred. Prob. Y=1 at mean X	0.024	0.027	0.014	0.045	0.054	0.085
Time dummies	YES	YES	YES	YES	YES	YES
Control variables[2]	YES	YES	YES	YES	YES	YES
County Fixed Effects	NO	YES	NO	NO	YES	NO
District Fixed Effects	NO	NO	YES	NO	NO	YES
Number of Observations[3]	2110	1715	535	2110	1855	640
Number of Local Districts[3]	422	343	107	422	371	128
log likelihood	-371.89	-340.26	-110.54	-467.46	-449.02	-198.50
Pseudo R-square	0.229	0.252	0.624	0.159	0.161	0.441

Note 1: Estimated Marginal Effects are evaluated at the mean value of the rest of the explanatory variables.
Note 2: Every estimation includes as control variables log of population, log of GDP per capita, log of population density, proportion population 5-14, proportion population 15-29, average rent, and proportion of claimants of unemployment benefits.
Note 3: Fixed effects estimations do not include districts for which the dependent variable does not have enough time variation.

Table 5
Myopic Game of Entry for McDonalds and Burger King
Under the Assumption that Players' Beliefs are in Equilibrium

Data: 422 markets, 2 firms, 5 years = 4,220 observations

	$\beta = 0.00$ (not estimated)			
	Two Step Estimates		**NPL Estimates**	
	Burger King	**McDonalds**	**Burger King**	**McDonalds**
Variable Profits:				
θ_0^{VP}	4.904 (1.070)*	7.909 (2.289)*	4.864 (1.081)*	7.898 (2.287)*
θ_1^{VP} cannibalization	2.005 (0.869)*	3.510 (0.659)*	2.035 (0.831)*	3.466 (0.647)*
θ_2^{VP} competition	0.014 (0.046)	0.032 (0.051)	0.016 (0.044)	0.037 (0.053)
Fixed Costs:				
θ_0^{FC} fixed	0.378 (0.212)*	0.806 (0.248)*	0.374 (0.212)*	0.808 (0.247)*
θ_1^{FC} linear	3.099 (0.436)*	2.662 (0.405)*	3.103 (0.436)*	2.659 (0.405)*
θ_2^{FC} quadratic	-0.054 (0.064)	0.085 (0.041)	-0.052 (0.063)	0.087 (0.041)
Pseudo R-square	0.154		0.154	
Log-Likelihood	-895.5		-895.4	
Distance $\|\|P^K - P^{K-}\|\|$			0.00	
# NPL iterations	1		5	

Table 6
Dynamic Game of Entry for McDonalds and Burger King
Under the Assumption that Players' Beliefs are in Equilibrium

Data: 422 markets, 2 firms, 5 years = 4,220 observations

	$\beta = 0.95$ (not estimated)			
	Two Step Estimates		**NPL Estimates**	
	Burger King	**McDonalds**	**Burger King**	**McDonalds**
Variable Profits:				
θ_0^{VP}	0.5849 (0.1077)*	0.8303 (0.2968)*	1.098 (0.2169)*	0.9737 (0.3091)*
θ_1^{VP} cannibalization	-0.2096 (0.0552)*	-0.0024 (0.0392)	-0.0765 (0.0725)	0.2874 (0.0986)*
θ_2^{VP} competition	-0.0110 (0.0029)*	0.0008 (0.0027)	-0.0129 (0.0065)*	-0.0074 (0.0073)
Fixed Costs:				
θ_0^{FC} fixed	0.0784 (0.0213)*	0.0822 (0.0332)*	0.0788 (0.0307)*	0.0773 (0.0261)*
θ_1^{FC} linear	0.0790 (0.0420)*	0.1076 (0.0400)*	0.1509 (0.0282)*	0.1302 (0.0185)*
θ_2^{FC} quadratic	-0.0078 (0.0059)	-0.0034 (0.0023)	-0.0054 (0.0026)*	0.0001 (0.016)
Pseudo R-square	0.323		0.146	
Log-Likelihood	-655.7		-893.4	
Distance $\|\|P^K - P^{K-}\|\|$	4831.26		0.00	
# NPL iterations	1		31	

Table 7
Estimated Bias in BK Beliefs
Difference Between B_{MD} and P_{MD}

	Stores of BK	
	0	
Stores of MD	**-0. 7 (0.04)**	-0.10 (0.06)
2	-0.08 (0.07)	-0.06 (0.10)

Estimated Bias in MD Beliefs
Difference Between B_{BK} and P_{BK}

	Stores of MD	
	0	
Stores of BK	-0.03 (0.05)	0.02 (0.04)
2	0.03 (0.10)	0.04 (0.12)