

Asymptotic Least Squares Estimators for Dynamic Games*

Martin Pesendorfer Philipp Schmidt-Dengler

London School of Economics and Political Sciences

First Version: July 2002; This Version: February 2006

Abstract

This paper considers the estimation problem in dynamic games with finite actions. We derive the equation system that characterizes the Markovian equilibria. The equilibrium equation system enables us to characterize conditions for identification. We consider a class of asymptotic least squares estimators defined by the equilibrium conditions. This class provides a unified framework for a number of well known estimators including Hotz and Miller (1993) and Aguirregabiria and Mira (2002). We show that these estimators differ in the weight they assign to individual equilibrium conditions. We derive the efficient weight matrix. A Monte Carlo study compares the finite sample performance of alternative estimators.

*This paper supersedes our earlier NBER working paper entitled “Identification and Estimation of Dynamic Games” from May 2003. We thank the editor Bernard Salanié and an anonymous referee, Mireia Jofre-Bonet, Alexandra Miltner, seminar participants at Alicante, Brown, Columbia, Duke, East Anglia, Lancaster, Leuven, Maryland, NYU, Queen’s, UCL, Toronto, Toulouse, Yale, Washington University, the NBER summer institute for helpful comments and Ariel Pakes for discussing the paper at the 2003 NBER summer institute. Martin Pesendorfer thanks the NSF under grant SES 0214222 for financial support.

1 Introduction

This paper builds on a well received literature for static single agent discrete choice models (McFadden (1978)) and their extension to Markov decision problems in Rust (1987), Hotz and Miller (1993) and Aguirregabiria and Mira (2002) among others. We also consider extensions of the single agent setting to multiple players.

The starting point of this paper is an equation system that provides necessary and sufficient conditions for a Markovian equilibrium. This set of equilibrium conditions enables us to illustrate conditions for identification. Based on these equilibrium conditions, we define a class of asymptotic least squares estimators. They are two step estimators. In the first step consistent estimates of the auxiliary choice and transition probabilities are obtained. In the second step the parameters of interest are inferred by using the set of equations characterizing the equilibrium choice probabilities.

The contribution of this paper is twofold: First, we suggest a unifying framework for the estimation of static and dynamic discrete choice models based on an equilibrium equation system. We show that a number of recently proposed estimators for dynamic models are asymptotic least squares estimators defined by the set of equilibrium conditions. The estimators differ in the weights they assign to individual equilibrium conditions. All estimators in the class are consistent and asymptotically normally distributed. They include the method of moments estimator, introduced by Hotz and Miller (1993) and the pseudo maximum likelihood estimator, Aguirregabiria and Mira (2002). We illustrate the associated weight matrices for these estimators.

The second contribution of this paper is to characterize the efficient estimator for dynamic games. We characterize the asymptotically optimal weight matrix.

The framework that we consider builds on the classic discrete choice models, see McFadden (1978), and includes a number of well known models of interaction. Prominent examples are the peer effect literature in public economics, see Brock and Durlauf (2001), the quantal response equilibrium in experimental economics, see McKelvey

and Palfrey (1995),¹ and the entry literature in industrial organization, see Bresnahan and Reiss (1990), and Seim (2005). Our framework also includes the substantial empirical literature on dynamic Markov decision models. A seminal contribution is Rust (1987), who proposes a nested fixed point algorithm to estimate the parameters of interest in dynamic models by maximum likelihood. The nested fixed point algorithm involves the calculation of the optimal choice probabilities and the value function for every parameter vector. Rust's estimator is efficient but computationally intensive. Rust (1994) shows that the estimator can be applied to estimate dynamic games. We illustrate that there is an asymptotic least squares estimator that is asymptotically equivalent to the maximum likelihood estimator.

Hotz and Miller (1993) introduce an elegant simplification consisting of a two step procedure. This two step approach forms the basis of a number of recent papers. In a first step the choice probabilities conditional on state variables are estimated. In a second step, the parameters of interest are inferred based on a set of moment conditions, one for each choice. Hotz and Miller's second step moment estimator involves choice specific moments interacted with instruments. We illustrate that the Hotz and Miller moment estimator is equivalent to an asymptotic least squares estimator. The choice and number of instruments matters, since together with the moment weight matrix they determine the weight assigned to individual equilibrium conditions. We show that instruments consisting of current and lagged state variables, as are typically used in applied work, may make it impossible to combine the equilibrium conditions efficiently. We characterize the form of the optimal instrument matrix.

Further, we show that the partial pseudo maximum likelihood estimator, introduced in Aguirregabiria and Mira (2002), is equivalent to a partial asymptotic least squares estimator with weights equal to the inverse of the covariance of the choice probabilities. The pseudo likelihood weights do not take into account the equilibrium

¹Bajari (1998) reformulates and estimates an auction model with discrete bids as a quantal response model.

conditions of dynamic games efficiently.

Recent working papers by Aguirregabiria and Mira (2004), Pakes, Ostrovsky and Berry (2004), and Bajari, Benkard and Levin (2004),² show that the pseudo likelihood estimator and the moment estimator can be applied to estimate dynamic games. We argue that the alternative estimators considered in these recent papers are asymptotic least squares estimators. We describe the corresponding weight matrices.

Our approach of first estimating equilibrium choices and beliefs and then using equilibrium conditions to infer the payoff parameters is related to the empirical auction literature. Elyakime, Laffont, Loisel and Vuong (1994) and Guerre, Perrigne, and Vuong (2000) estimate the distribution of equilibrium actions based on bid data. The distribution function estimates summarize bidders' beliefs and are used in a second step to infer bidders' valuations based on the equilibrium conditions manifested in the first order condition of optimal actions. Jofre-Bonet and Pesendorfer (2003) extend this approach to a dynamic setting.

While the unified framework allows us to conclude which estimator may be preferred asymptotically, this preference is less clear when the number of observations is small. We conduct a Monte Carlo study to examine the small sample performance of a number of these estimators. We consider the efficient asymptotic least squares estimator and compare it to the pseudo maximum likelihood estimator and a method of moments estimator based on the average probability of a given choice across states. Our Monte Carlo study reveals efficiency gains of the asymptotically efficient least squares estimator for moderate to large sample sizes.

Section 2 provides an illustrative example for a special case in which the dynamic element is absent. We illustrate our econometric approach based on the equilibrium conditions of a two-player two-action game. We describe the two step estimator and

²Bajari, Benkard and Levin (2004) also show that the moment estimator can include both discrete and continuous choices. Ryan (2004) and Collard-Wexler (2005) provide applications to the cement and ready-mix concrete industries.

illustrate the basic intuition for our estimation approach. Our subsequent analysis augments the setup to include settings in which agents rationally take the future implications of their actions into account.

Section 3 describes the elements of the dynamic game. We consider a model that is closest in spirit to the original contributions in Rust (1987), Hotz and Miller (1993) and much of the subsequent literature based on dynamic discrete choice. We assume that every period each player privately observes a vector of payoff shocks drawn from a known distribution function conditional on state variables. Players simultaneously choose an action from a finite set in order to maximize the sum of discounted future period payoffs.

Recently, alternative models have been considered in applied work in which players move sequentially, as in Einav (2003) and Schmidt-Dengler (2005), in which continuous choice variables are permitted, in which the game has a finite horizon, or in which an alternative error structure is imposed. These model formulations yield equilibrium conditions similar to ours. The subsequent analysis can be extended to these models with modifications in the appropriate places.

Section 4 establishes properties of the equilibrium. We show that there is a set of necessary and sufficient equilibrium conditions described by an equation system. The equation system forms the basis for the estimation problem. We show that a solution to the equation system, and thus an equilibrium, exists.

Section 5 analyzes the identification conditions for the model. There are two basic points emerging from the identification analysis. First, similar to the results on single agent discrete decision processes by Rust (1994) and Magnac and Thesmar (2002), we find that not all primitives of the model are identified. Second, we show that the degree of under-identification increases with the number of agents. We briefly describe restrictions that guarantee identification.

Section 6 discusses the estimation problem. We consider the class of asymptotic least squares estimators defined by the equilibrium equation system. We provide a

set of properties that apply to all estimators in the class. The asymptotically optimal weight matrix in the context of dynamic games is derived. We argue that the class of asymptotic least squares estimators provides a unified framework that encompasses a number of estimators proposed in the literature. Members of the class assign distinct weights to individual equilibrium conditions. We examine to what extent the weight matrices of estimators proposed in the literature differ from the optimal weight matrix.

Section 7 reports our results of a small sample Monte Carlo study. We compare the asymptotically optimal estimator to a number of estimators proposed recently in the literature.

Section 8 concludes.

2 Example

This section illustrates the intuition of the estimation approach in a simple setting. We consider a (static) two player game with two actions.³ We begin by describing the features of the game, then we illustrate the equilibrium conditions and finally we illustrate our estimation approach.

Set up. Consider a version of the “battle of the sexes” game in which a man and a woman decide whether to go to the opera or to the football match in the evening. Each player i has two possible choices: football and opera, $a_i \in \{\text{football}, \text{opera}\}$. The choices are made simultaneously. We augment this classic game with two elements: A publicly observed weather variable and a private mood variable. The weather variable describes whether there is snow, rain, or sunshine, $s \in \{1, 2, 3\}$. We shall assume that weather affects the consumption value of the open-air football match but not the consumption value of the indoor opera performance. Player i ’s payoff

³For other examples see the games studied in McKelvey and Palfrey (1995), Brock and Durlauf (2001), and Seim (2005).

depends on the players' choices and the weather in the following way:

$$\pi_i(a_i, a_j, s) \begin{cases} \theta_i^1 \cdot s & \text{if } a_i = a_j = \text{football}; \\ \theta_i^2 & \text{if } a_i = a_j = \text{opera}; \\ 0 & \text{otherwise.} \end{cases}$$

The term “battle of the sexes” has been attributed to this payoff structure as preference parameters may differ $\theta_1^1 < \theta_2^1, \theta_1^2 > \theta_2^2$, yet going to distinct events may be costly when $\theta_i^1, \theta_i^2 > 0$ for $i = 1, 2$.

The second element that we incorporate into the game is a mood variable. The mood is an opera specific preference shock ε_i drawn from the standard normal distribution function denoted by Φ . The shock ε_i measures player i 's preference for the opera relative to the football match on a given day. We assume that the preference shock is iid and enters additively yielding a total payoff of

$$\begin{cases} \pi_i(\text{football}, a_j, s) & \text{if } a_i = \text{football} \\ \pi_i(\text{opera}, a_j, s) + \varepsilon_i & \text{if } a_i = \text{opera} \end{cases}$$

We may view the mood shock ε_i as any element of preferences that is not known to the other player but is known to player i .

A pure strategy for player i in the augmented “battle of sexes” game is denoted by $a_i(s, \varepsilon_i)$ which is a function of the public state s and the private preference shock ε_i . We are interested in Bayesian Nash equilibria. A feature that distinguishes games from single agent decision problems is the possibility of multiple outcomes for the same (s, ε_i) . Indeed, when $\theta_i^1, \theta_i^2 > 0$ for $i = 1, 2$, then one equilibrium has the feature that players go more frequently to the opera, while in another equilibrium players go more frequently to the football match. This multiplicity of outcomes poses a difficulty for direct estimation approaches as the likelihood of observing an outcome will depend on which equilibrium is played. Furthermore, the set of equilibria may depend on the parameter values and may change as the parameter values are varied. Our estimation approach circumvents this difficulty by considering an indirect

inference approach based on “estimating equations.” We consider the equilibrium best response conditions as estimating equations. Next, we describe the equilibrium conditions in more detail. We then illustrate our estimator.

Equilibrium conditions. Let $\sigma_i(s)$ denote the belief of player i that player j will go to the opera in state s for $i = 1, 2$. The number $\sigma_i(s)$ equals the *ex ante* expected probability that player j will go to the opera, while the number $1 - \sigma_i(s)$ equals the *ex ante* expected probability that player j will go to the football match. In any equilibrium, it must be that going to the opera is a best response for some type ε_i , taking as given the beliefs about the opponent’s behavior:

$$\sigma_i(s) \cdot \theta_i^2 + \varepsilon_i \geq [1 - \sigma_i(s)] \cdot \theta_i^1 \cdot s \quad (1)$$

as the preference shock ε_i has support equal to the real line. Evaluating the best response condition *ex ante*, before the payoff shock ε_i is observed, we can describe the probability of going to the opera in state s for player i as:

$$p_i(s) = 1 - \Phi([1 - \sigma_i(s)] \cdot \theta_i^1 \cdot s - \sigma_i(s) \cdot \theta_i^2)$$

for $i = 1, 2$ and $s = 1, 2, 3$, which gives rise to an equation system that characterizes the optimal choices given the beliefs $\boldsymbol{\sigma} = (\sigma_i(s)_{i=1}^2)_{s=1}^3$. We may write the best response equation system compactly using vector notation as

$$\mathbf{p} - \Psi(\boldsymbol{\sigma}, \boldsymbol{\theta}) = 0. \quad (2)$$

where $\mathbf{p} = (p_i(s)_{i=1}^2)_{s=1}^3$ denotes the vector of choice probabilities, $\boldsymbol{\theta} = (\theta_1^1, \theta_1^2, \theta_2^1, \theta_2^2)$ denotes the parameter vector and Ψ is a multivariate function that characterizes the best responses.

In equilibrium, beliefs are consistent, $p_i(s) = \sigma_j(s)$ for $i, j = 1, 2$, $i \neq j$ and $s = 1, 2, 3$, which yields a fixed point problem in terms of the *ex ante* choice probabilities:

$$\mathbf{p} = \Psi(\mathbf{p}, \boldsymbol{\theta}). \quad (3)$$

The equilibrium equations (3) provide a relationship between the equilibrium choice probabilities \mathbf{p} and the vector of model parameters $\boldsymbol{\theta}$. The function Ψ is determined by the assumptions on payoffs for the problem at hand. Our example has two players, two actions and three states yielding a total of six best response equations in (3).

The privately observed preference shock plays an important role in the derivation of the equilibrium equations. Given a player's beliefs, the decision problem in (1) is similar to a standard discrete choice problem. The unbounded support assumption yields a vector of equilibrium choice probabilities \mathbf{p} in which all components of the vector are bounded away from zero. Since all actions arise with positive probability, the equilibrium is characterized by a system of equations rather than a system of inequalities which emerges in the absence of the privately observed preference shock. The set of equilibrium equations thus allows us to adopt a simple estimation method.

Asymptotic least squares estimator. Suppose the available data consist of a time series of actions of repeated play of the same two individuals and a time series for the weather variable. The data set is thus summarized by $(a_i^t, a_j^t, s^t)_{t=1}^T$. Furthermore, we assume that the data are generated by a Markovian equilibrium which means for a given vector of payoff relevant state variables (s, ε_i) players make the same choices over time. Thus our time series approach gets around the multiplicity problem that would be present in a cross-sectional analysis. Given our assumption that the mood variable ε_i is iid normally distributed, the parameter vector that we wish to infer is $\boldsymbol{\theta}$.

Our estimator for $\boldsymbol{\theta}$ is based on the following idea. The time series data tell us the frequency with which every player goes to the opera for every realization of the weather variable. Given that errors are privately observed, players form beliefs based on the same information that is available to the econometrician. The data thus allow us to estimate the equilibrium choice probability, $\hat{p}_i(s)$, for all i and s . For example, a frequency estimator can be used. Next, the choice probability estimates can be substituted into the equilibrium equation system (3), and the resulting equation system

has the payoff parameters θ as the only unknown elements:

$$\hat{\mathbf{p}} - \Psi(\hat{\mathbf{p}}, \theta) = 0. \quad (4)$$

The equation system (4) has 6 equations and the unknown parameter vector θ has four elements. Our estimation approach is to estimate θ by using least squares on this equation system. For a given 6×6 weight matrix \mathbf{W} , the least squares problem is given by

$$\min_{\theta} [\hat{\mathbf{p}} - \Psi(\hat{\mathbf{p}}, \theta)]' \mathbf{W} [\hat{\mathbf{p}} - \Psi(\hat{\mathbf{p}}, \theta)].$$

A solution can be found by minimizing a quadratic form with the use of a suitably chosen weight matrix \mathbf{W} . Distinct choices of weight matrices will give rise to distinct estimates and thus define the class of asymptotic least squares estimators.

The estimator is simple to implement in a two step procedure. In the first step, the auxiliary estimates of choice probabilities can be characterized from the data by using the sample frequencies of choices. In the second step, the parameters of interest can be estimated using least squares on the equilibrium equation system (4).

In the following sections we specify a general dynamic discrete-choice multi-player model and characterize its equilibrium equations. We then define the class of asymptotic least squares estimators based on the equilibrium equations and show that all members of the class have nice asymptotic properties. We shall describe a weight matrix that is optimal in the sense that it yields estimates asymptotically equivalent to the maximum likelihood estimates. Finally, we also illustrate that a number of well known estimators are members of the class.

3 Model

This section describes the elements of the general model. We describe the sequencing of events, the period game, the transition function, the payoffs, the strategies and the equilibrium concept.

We consider a dynamic game with discrete time, $t = 1, 2, \dots, \infty$. The set of players is denoted by $\mathbf{N} = \{1, \dots, N\}$ and a typical player is denoted by $i \in \mathbf{N}$. The number of players is fixed and does not change over time. Players can choose from a finite set of $K + 1$ actions. Every period the following events take place:

States. Each player is endowed with state variable $s_i^t \in \mathbf{S}_i = \{1, \dots, L\}$ and a vector of payoff shocks $\boldsymbol{\varepsilon}_i^t \in \mathbb{R}^K$. The state variable s_i^t is publicly observed by all players and the econometrician. The payoff shock $\boldsymbol{\varepsilon}_i^t$ is privately observed by the player. The shock is not observed by other players or the econometrician.

The vector of all players' public state variables is denoted by $\mathbf{s}^t = (s_1^t, \dots, s_N^t) \in \mathbf{S} = \times_{j=1}^N \mathbf{S}_j$. Sometimes we use the notation $\mathbf{s}_{-i}^t = (s_1^t, \dots, s_{i-1}^t, s_{i+1}^t, \dots, s_N^t) \in \mathbf{S}_{-i}$ to denote the vector of states by players other than player i . The cardinality of the state space \mathbf{S} equals $m_s = L^N$.

We assume that the payoff shock $\boldsymbol{\varepsilon}_i^t$ is drawn independently from the strict monotone and continuous conditional distribution function $F(\cdot | s_i^t, \mathbf{s}_{-i}^t)$ defined on \mathbb{R}^K . The shock $\boldsymbol{\varepsilon}_i^t$ does not depend on the actions of other players and is independent from past $\boldsymbol{\varepsilon}_i^{t-\tau}$ for $\tau \geq 1$. We assume that $E[\boldsymbol{\varepsilon}_i^t | \boldsymbol{\varepsilon}_i^t \geq \boldsymbol{\varepsilon}]$ exists for all $\boldsymbol{\varepsilon} \in \mathbb{R}^K$, to ensure that the expected period return exists. The economic interpretation of the payoff shock may vary with the modelling context and has been attributed to a temporary productivity shock, a shock to opportunity costs and the mood of a player in the literature.⁴

Actions. Each player chooses an action $a_i^t \in A_i = \{0, 1, \dots, K\}$. All N players make their decisions simultaneously. The actions are taken after players observe the state and their idiosyncratic productivity shock. An action profile \mathbf{a}^t denotes the vector of joint actions in period t , $\mathbf{a}^t = (a_1^t, \dots, a_N^t) \in A = \times_{j=1}^N A_j$. Sometimes we use the notation $\mathbf{a}_{-i}^t = (a_1^t, \dots, a_{i-1}^t, a_{i+1}^t, \dots, a_N^t) \in A_{-i}$ to denote the actions by players

⁴An alternative interpretation in a static model is an idiosyncratic optimization error as in McKelvey and Palfrey (1995). In a dynamic model the latter interpretation would be internally inconsistent with optimal forward looking behavior, see Rust (1994).

other than player i . The cardinality of the action space A equals $m_a = (K + 1)^N$.

State transition. The state transition is described by a probability density function $g : A \times \mathbf{S} \times \mathbf{S} \longrightarrow [0, 1]$ where a typical element $g(\mathbf{a}^t, \mathbf{s}^t, \mathbf{s}^{t+1})$ equals the probability that state \mathbf{s}^{t+1} is reached when the current action profile and state are given by $(\mathbf{a}^t, \mathbf{s}^t)$. We require $\sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{a}, \mathbf{s}, \mathbf{s}') = 1$ for all $(\mathbf{a}, \mathbf{s}) \in A \times \mathbf{S}$. We frequently use the symbol \mathbf{G} to denote the $(m_a \cdot m_s) \times m_s$ dimensional transition matrix in which column $\mathbf{s}' \in \mathbf{S}$ consists of the vector of probabilities $\left[g(\mathbf{a}, \mathbf{s}, \mathbf{s}')_{\mathbf{a} \in A, \mathbf{s} \in \mathbf{S}} \right]$.

The period payoff of player i is collected at the end of the period after all actions are observed. The period payoff of player i consists of two components: (i) the action and state dependent period payoff, and (ii) the action dependent payoff shock realization. The period payoffs are given by:

$$\pi_i(\mathbf{a}^t, \mathbf{s}^t) + \sum_{k=1}^K \varepsilon_i^{tk} \cdot \mathbf{1}(a_i^t = k)$$

where $\mathbf{1}(a)$ equals one if event a occurs and zero otherwise. We sometimes use the symbol $\mathbf{\Pi}_i$ to denote the $(m_a \cdot m_s) \times 1$ dimensional period payoff vector defined by $\mathbf{\Pi}_i = \left[\pi_i(\mathbf{a}, \mathbf{s})_{\mathbf{a} \in A, \mathbf{s} \in \mathbf{S}} \right]$. The payoff shock affects actions $k > 0$ only. The assumption that action 0 is not affected by the payoff shock is not restrictive as only the payoff difference between alternative actions matters.

Game payoff. Players discount future payoffs with discount factor $\beta_i < 1$. The game payoff of player i equals the sum of discounted period payoffs. For expositional purposes we assume that the discount factor is common across players, $\beta_i = \beta$ for all i .

Following Maskin and Tirole (1994, 2001), we consider pure Markovian strategies $a_i(\varepsilon_i^t; s^t)$ in which an action for player i is a function of the player specific payoff shock and the state variables. We restrict attention to pure strategies and do not consider mixed strategies. The Markovian restriction requires that the action at time t is the same as the action at time t' , $a_i(\varepsilon_i^t; s^t) = a_i(\varepsilon_i^{t'}; s^{t'})$ if $(\varepsilon_i^t; s^t) = (\varepsilon_i^{t'}; s^{t'})$. The Markovian assumption allows us to abstract from calendar time and we subsequently

omit the time superscript.

Definition (Markov perfect equilibrium): A collection $(a, \sigma) = (a_1, \dots, a_N, \sigma_1, \dots, \sigma_N)$ is a Markov perfect equilibrium if

- (i) for all i , a_i is a best response to a_{-i} given the beliefs σ_i at all states $\mathbf{s} \in \mathbf{S}$;
- (ii) all players use Markovian strategies;
- (iii) for all i the beliefs σ_i are consistent with the strategies a .

Value function. We use the ex-ante value function V_i to express the discounted sum of future payoffs. It is defined as the discounted sum of future payoffs for player i before player-specific shocks are observed and actions are taken. Let $\sigma_i(\mathbf{a}|\mathbf{s})$ denote the conditional ex ante (before ε_i is observed) belief of player i that action profile \mathbf{a} will be chosen conditional on state \mathbf{s} . The ex ante value function can be written as,

$$\begin{aligned} V_i(\mathbf{s}; \sigma_i) &= \sum_{\mathbf{a} \in \mathbf{A}} \sigma_i(\mathbf{a}|\mathbf{s}) \left[\pi_i(\mathbf{a}, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{a}, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; \sigma_i) \right] \\ &\quad + \sum_{k=1}^K E_\varepsilon [\varepsilon_i^k | a_i = k] \cdot \sigma_i(a_i = k | \mathbf{s}), \end{aligned} \quad (5)$$

where E_ε denotes the expectation operator with respect to the player specific productivity shock. Equation (5) is satisfied at every state vector $\mathbf{s} \in \mathbf{S}$. Since the state space is finite, we can express it as a matrix equation:

$$\begin{aligned} \mathbf{V}_i(\sigma_i) &= \boldsymbol{\sigma}_i \boldsymbol{\Pi}_i + \mathbf{D}_i(\sigma_i) + \beta \boldsymbol{\sigma}_i \mathbf{G} \mathbf{V}_i(\sigma_i) \\ &= [\mathbf{I}_s - \beta \boldsymbol{\sigma}_i \mathbf{G}]^{-1} [\boldsymbol{\sigma}_i \boldsymbol{\Pi}_i + \mathbf{D}_i(\sigma_i)]. \end{aligned} \quad (6)$$

Equation (6) provides an expression for the *ex ante* value function associated with the beliefs σ_i . The terms on the right hand side are the discount factor, the choice probability matrix, the state transition matrix, the period return function, and the expected payoff shocks. Here $\mathbf{V}_i(\sigma_i) = [V_i(\mathbf{s}; \sigma_i)]_{\mathbf{s} \in \mathbf{S}}$ is the $m_s \times 1$ dimensional vector of expected discounted sum of future payoffs; $\boldsymbol{\sigma}_i$ is the $m_s \times (m_a \cdot m_s)$ dimensional matrix consisting of player i 's beliefs $\sigma_i(\mathbf{a}|\mathbf{s})$ in row \mathbf{s} , column (\mathbf{a}, \mathbf{s}) , and zeros in row \mathbf{s} , column $(\mathbf{a}, \mathbf{s}')$ with $\mathbf{s}' \neq \mathbf{s}$; $\mathbf{D}_i(\sigma_i) = [D_i(\mathbf{s}, \sigma_i)]_{\mathbf{s} \in \mathbf{S}}$ is the $m_s \times 1$ dimensional

vector of expected payoff shocks with element $D_i(\mathbf{s}, \sigma_i) = \sum_{k=1}^K E_\varepsilon [\varepsilon_i^k | a_i = k; \sigma_i] \cdot \sigma_i(a_i = k | \mathbf{s})$; and \mathbf{I}_s denotes the m_s dimensional identity matrix. The second line in equation (6) follows from the dominant diagonal property which implies that the matrix $[\mathbf{I}_s - \beta \boldsymbol{\sigma}_i \mathbf{G}]$ is invertible.

Notice, that the value function can be evaluated at any vector of beliefs, not necessarily equilibrium beliefs. In the next section, we use this expression to characterize the properties of the equilibrium choices and beliefs.

4 Equilibrium Properties

This section characterizes properties of the dynamic equilibrium and provides remarks on the robustness of our model assumptions. Subsection 4.1 begins with a characterization of the equilibrium decision rule, and the equilibrium choice probabilities. We show that the equilibrium decision rule is characterized by an $N \cdot K \cdot m_s$ equation system. We show that an equilibrium exists. Subsection 4.2 concludes with remarks on the robustness of our conclusions to alternative model assumptions.

4.1 Equilibrium Characterization

First, we examine when it is optimal to choose action a_i in state $(\mathbf{s}, \varepsilon_i)$. Let $\boldsymbol{\theta}$ denote the parameter vector summarizing the model elements $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N, F, \beta, g)$. Let $u_i(a_i; \sigma_i, \boldsymbol{\theta})$ denote player i 's continuation value net of the payoff shocks under action a_i with beliefs σ_i :

$$u_i(a_i; \sigma_i, \boldsymbol{\theta}) = \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} \sigma_i(\mathbf{a}_{-i} | \mathbf{s}) \cdot \left[\pi_i(\mathbf{a}_{-i}, a_i, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in S} g(\mathbf{a}_{-i}, a_i, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; \sigma_i) \right] \quad (7)$$

where $\sigma_i(\mathbf{a}_{-i} | \mathbf{s})$ denotes the conditional beliefs of player i that players $-i$ choose an action profile \mathbf{a}_{-i} conditional on state \mathbf{s} .

It is optimal to choose action a_i given the beliefs σ_i if the continuation value under

action a_i exceeds the continuation value under all alternative actions:

$$u_i(a_i; \sigma_i, \boldsymbol{\theta}) + \varepsilon_i^{a_i} \geq u_i(a'_i; \sigma_i, \boldsymbol{\theta}) + \varepsilon_i^{a'_i} \quad \text{for all } a'_i \in \mathbf{A}_i \quad (8)$$

where $\varepsilon_i^0 = 0$. The optimality condition (8) characterizes the optimal decision rule up to a set of Lebesgue measure zero on which two or more alternative actions yield equal continuation values. We may evaluate condition (8) before the payoff shock ε_i is observed. Doing so, gives an expression for the ex ante optimal choice probability of player i given the beliefs σ_i :

$$\begin{aligned} p(a_i | \mathbf{s}, \sigma_i) &= \Psi_i(a_i, \mathbf{s}, \sigma_i; \boldsymbol{\theta}) \\ &= \int \prod_{k \in \mathbf{A}_i, k \neq a_i} \mathbf{1}(u_i(a_i; \sigma_i, \boldsymbol{\theta}) - u_i(k; \sigma_i, \boldsymbol{\theta}) \geq \varepsilon_i^k - \varepsilon_i^{a_i}) dF \end{aligned} \quad (9)$$

In matrix notation, equation (9) can be written as:

$$\mathbf{p} = \boldsymbol{\Psi}(\boldsymbol{\sigma}; \boldsymbol{\theta}) \quad (10)$$

where $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_N)$ denotes the $(N \cdot K \cdot m_s) \times 1$ dimensional vector of the optimal players' choice probabilities for all states, players and actions other than action 0; and $\boldsymbol{\sigma}$ denotes the $(N \cdot K \cdot m_s) \times 1$ dimensional vector of players' beliefs. We omit the choice probability for action 0 in equation (10) as it is already determined by the remaining choice probabilities, $p_i(0 | \mathbf{s}) = 1 - \sum_{k=1}^K p_i(k | \mathbf{s})$.

Equation system (10) determines the optimal choice probabilities as a function of players' beliefs. In equilibrium, beliefs are consistent leading to a fixed point problem in ex ante choice probabilities \mathbf{p} :

$$\mathbf{p} = \boldsymbol{\Psi}(\mathbf{p}; \boldsymbol{\theta}). \quad (11)$$

The following proposition characterizes the set of Markov Perfect equilibria. All proofs are given in the appendix.

Proposition 1 (*Characterization*) *In any Markov Perfect equilibrium, the probability vector \mathbf{p} satisfies equation (11). Conversely, any \mathbf{p} that satisfies equation (11) can be extended to a Markov perfect equilibrium.*

The proposition states that equation system (11) is a necessary and also a sufficient condition for any Markov equilibrium. It gives a system of $N \cdot K \cdot m_s$ equations characterizing the $N \cdot K \cdot m_s$ equilibrium ex ante choice probabilities. The necessity part stems from the optimality condition which says that the continuation value when taking action a_i is at least as large as the continuation value when taking action a'_i . The sufficiency part is established by showing that any \mathbf{p} that satisfies equation (11) can be extended to construct a decision rule based on condition (8) that constitutes a Markov perfect equilibrium under the beliefs \mathbf{p} .

Equation system (11) is not the only equation system that characterizes the set of equilibria. For example, any monotone transformation of the left hand side and right hand side variables in (11) can also be used as an equation system characterizing the equilibrium. This property will be convenient when we study identification of the model, where we will use a characterization that is linear in payoffs.

Next, we provide a few remarks on the existence, computation and multiplicity of equilibria. The following theorem establishes that an equilibrium exists.

Theorem 1 *A Markov equilibrium exists.*

The existence result follows from Brouwer's fixed point theorem. Equation (11) gives us a continuous function from $[0, 1]^{N \cdot K \cdot m_s}$ onto itself. By Brouwer's theorem it has a fixed point \mathbf{p} . Maskin and Tirole (1994, 2001) and Doraszelski and Satterthwaite (2003) provide alternative existence arguments for related problems.

Proposition 1 and Theorem 1 enable us to calculate equilibria numerically. They show that the equilibrium calculation is reduced to a fixed point problem in $[0, 1]^{N \cdot K \cdot m_s}$, which is solvable even for non-trivial m_s , K and N . Backwards solving algorithms, which calculate a new optimal policy and value function at each step of the iteration, are computationally complex, as is shown in Pakes and McGuire (2001). Proposition 1 shows that the fixed point problem in equation (11) gives an equivalent representation of equilibria which may facilitate the computation considerably. Hence, it is not

necessary to determine the optimal decision rule and value function at every step of the iteration. Numerical solution techniques for continuous time games are studied in Doraszelski and Judd (2004).

Markov perfect equilibria need not be unique. Section 2 provides an illustration for a game in which dynamic linkages are absent. Section 7 gives an example of multiple equilibria in a dynamic context.

4.2 Robustness

Next, we discuss the robustness of our framework to alternative model assumptions. We discuss the symmetry assumption which places tighter restrictions on players' strategies, the extension to correlated and asymmetrically distributed payoff shocks, and modification in the timing structure of the game. We conclude the section with a description of conceptual limitations of our framework.

Symmetry Restriction. Proposition 1 and Theorem 1 permit asymmetric payoffs and strategies. Our framework takes heterogeneous behavior by agents into account. It permits the possibility that players i and j adopt distinct strategies, $a_i(\mathbf{s}, \varepsilon_i) \neq a_j(\mathbf{s}', \varepsilon_j)$, even when the state variables are identical, $\mathbf{s} = \mathbf{s}'$ and $\varepsilon_i = \varepsilon_j$. Imposing symmetry on payoffs and strategies does not affect the properties of the framework conceptually. Symmetry requires identical payoff vectors $\Pi_i = \Pi$ for all i and identical decision rules, $a_i(\mathbf{s}, \varepsilon_i) = a(\mathbf{s}, \varepsilon_i)$ for all i . It can be verified that the above arguments remain valid under the symmetry restriction which leads us to the following Corollary:

Corollary 1 *Suppose $\Pi_i = \Pi$.*

- (i) *A symmetric Markov equilibrium exists.*
- (ii) *In any symmetric Markov Perfect equilibrium, the probability vector \mathbf{p} satisfies equation (11). Conversely, any \mathbf{p} that satisfies equation (11) can be extended to a symmetric Markov perfect equilibrium.*

The symmetry assumption reduces the number of equations in (11). The reduction in dimension reduces the complexity of the problem which may facilitate the numerical calculation of equilibria. Additionally, symmetry places a number of restrictions on the payoff vector Π which can be useful in the empirical analysis.

Payoff shock. The assumption regarding independence of the payoff shock $\varepsilon_i^{a_i}$ is an important assumption commonly made in this literature. It permits the use of the Markovian framework. Introducing serial correlation in the privately observed error would give rise to models of learning in which players form beliefs about other players' states based on past actions. To model these beliefs consistently, the state space would need to be amplified to include the set of all possible past actions. Introducing an additional payoff shock that is common to all players and observed by all players but not by the econometrician would give rise to a model in which the beliefs of players depend additionally on the common shock. It would thus be impossible to correctly infer the beliefs of players. For a detailed discussion of the independence assumption in Markovian decision problems see Rust (1994).

While independence of the payoff shocks is a crucial assumption, the properties of the model do not change when the payoff shocks are not identically distributed across actions and states. We maintain this assumption for expositional simplicity.

Structure of the Game. Static models, as in McKelvey and Palfrey (1995), Seim (2005), and others, are a special case of our framework with $\beta = 0$. Our framework can be extended to incorporate environments in which players move sequentially rather than simultaneously, as in Einav (2003), or in which there is a finite time horizon.

Action Space. Recall that the equilibrium was characterized by one equation for every action in every state. Characterization of a game with a continuous action space would require a continuum of equilibrium conditions for every state. In the pricing game studied in Jofre-Bonet and Pesendorfer (2003) the equilibrium condition corresponding to (11) becomes a set of first-order conditions.

Limitations. Our modelling framework and approach also has limitations. First, it

is better suited for situations in which there is a small number of players and states. In situations with many players, or states, the transition matrix G will be large and solving the equations in (11) may become computationally infeasible. Second, it may be difficult to simulate the equilibria even if estimation is possible as there can be multiple equilibria. Finding the set of fixed points of equation (11) may be computationally infeasible.

This section has characterized some properties of Markov perfect equilibria. We have provided a simple characterization of the equilibrium choice probabilities in form of an equation system that can be solved numerically. We have shown that the equation system has a solution which implies that a Markov perfect equilibrium exists. Finally, we illustrated that our model framework is robust to a number of variations in the model assumptions. We also have illustrated limitations to our approach. Next, we use the necessary and sufficient equilibrium condition (11) to address the question of identification.

5 Identification

This section examines the identification question. The model is identified if there exists a unique set of model primitives $(\Pi_1, \dots, \Pi_N, F, \beta, g)$ that can be inferred from a sufficiently rich data set characterizing choice and state transition probabilities.

We consider time series data generated by a single path of play and we exploit information contained in the repeated observations on the same set of players along the path of play. Our time series approach differs from the cross section approach frequently used in the single agent dynamic literature. The cross section approach does not extend to games without additional restrictive assumptions as, in contrast to single agent problems, games may have multiple equilibria.⁵ Estimation based on

⁵Identification of static entry models based on a cross-section data set is considered in Tamer (2003).

a single path of play gets around the multiplicity concerns inherent to Markovian equilibria: When the data are generated by a Markovian equilibrium, players make the same choices over time for a given vector of payoff relevant state variables (s, ε_i) .

We assume that time series data $(\mathbf{a}^t, \mathbf{s}^t)_{t=1}^T$ permit us to characterize the choice probabilities $p(\mathbf{a}|s)$ and the transition probabilities $g(\mathbf{a}, \mathbf{s}, \mathbf{s}')$ for any $\mathbf{s}, \mathbf{s}' \in \mathbf{S}, \mathbf{a} \in A$.

To study the identification problem, we examine the properties of equation system (11). The unknown elements in equation system (11) are the discount factor β , the distribution function F , and the $m_a \cdot m_s \cdot N$ parameters in $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N)$. Our approach follows the analysis of the single agent case by Magnac and Thesmar (2002) in that we fix (F, β) and analyze the identification of payoff parameters $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N)$. The following Proposition illustrates the need for identifying restrictions.

Proposition 2 *Suppose (F, β) are given. At most $K \cdot m_s \cdot N$ parameters can be identified.*

Proposition 2 states that not all the model elements in $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N, F, \beta)$ are non-parametrically identified. Even if the distribution function F and the discount factor β are known, not all parameters in $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N)$ can be identified. The reason is that there are a total of $K \cdot m_s \cdot N$ equations in (11) enabling identification of at most as many parameters while there are $m_a \cdot m_s \cdot N$ parameters in $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N)$. Thus at least $(m_a \cdot m_s - K \cdot m_s) \cdot N$ restrictions are needed. The non-identification result is similar in spirit to results obtained in the single agent dynamics literature, see Rust (1994) and Magnac and Thesmar (2002).⁶ Proposition 2 also illustrates that the degree of under-identification increases with the number of agents. The reason is that the number of required restrictions is exponential in N , while the number of equations is linear in N .

⁶Heckman and Navarro (2005) obtain stronger identification results for a binary choice single agent model, because one action ('treatment') corresponds to an absorbing state and they have information on measured consequences of reaching this absorbing state.

In many economic settings the payoff parameters are the main object of interest. The researcher may be willing to assume a functional form for the distribution F and may obtain estimates for the discount factor β elsewhere. Next, we illustrate how to incorporate a (minimal) set of linear restrictions on the payoff parameters and ensure that the remaining payoff parameters are identified. Without loss of generality, we consider player i . Let \mathbf{R}_i be a $(m_a \cdot m_s - K \cdot m_s) \times (m_a \cdot m_s)$ dimensional restriction matrix and \mathbf{r}_i a $(m_a \cdot m_s - K \cdot m_s) \times 1$ dimensional vector such that

$$\mathbf{R}_i \cdot \mathbf{\Pi}_i = \mathbf{r}_i. \quad (12)$$

We determine whether the model is exactly identified with the restriction (12) in place.

It is convenient to rephrase the equilibrium conditions by using the optimality condition (8) for the action pair a_i and 0. The resulting equations are linear in period payoffs which allows us to illustrate the identification problem for payoff parameters in a simple way. Let $\bar{\varepsilon}_i^{a_i}(s)$ denote the realization of $\varepsilon_i^{a_i}$ that makes player i indifferent between actions a_i and 0 in state s . It corresponds to the realization of $\varepsilon_i^{a_i}$ at which equation (8) holds with equality when comparing the payoffs from action a_i and action 0. The indifference equation between action pairs is given by:

$$\begin{aligned} & \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} p(\mathbf{a}_{-i} | \mathbf{s}) \cdot \left[\pi_i(\mathbf{a}_{-i}, a_i, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in S} g(\mathbf{a}_{-i}, a_i, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; p) \right] + \bar{\varepsilon}_i^{a_i}(s) \\ = & \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} p(\mathbf{a}_{-i} | \mathbf{s}) \cdot \left[\pi_i(\mathbf{a}_{-i}, 0, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in S} g(\mathbf{a}_{-i}, 0, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; p) \right] \end{aligned} \quad (13)$$

Substituting the ex ante expected value function, equation (6), and expressing the above equation in matrix notation, we obtain a system of equations that is linear in the payoff parameters. The linear equation system for player i is given by:

$$\mathbf{X}_i(\mathbf{p}, g, \beta) \cdot \mathbf{\Pi}_i + \mathbf{Y}_i(\mathbf{p}, g, \beta) = 0. \quad (14)$$

where \mathbf{X}_i and \mathbf{Y}_i are $(K \cdot m_s) \times (m_a \cdot m_s)$ dimensional and $(K \cdot m_s) \times 1$ coefficient matrices determined by the choice probabilities, the transition probabilities and the

discount factor β . Lemma 1 in the appendix shows that \mathbf{p} satisfies (14) for all players i if and only if \mathbf{p} satisfies the best response system (11). Next, we augment the linear equation system (14) with the payoff restrictions $\mathbf{R}_i \cdot \mathbf{\Pi}_i = \mathbf{r}_i$, yielding the $m_a \cdot m_s$ equations in $m_a \cdot m_s$ unknown payoff parameters:

$$\begin{aligned} \begin{bmatrix} \mathbf{X}_i \\ \mathbf{R}_i \end{bmatrix} \mathbf{\Pi}_i + \begin{bmatrix} \mathbf{Y}_i \\ \mathbf{r}_i \end{bmatrix} &= \overline{\mathbf{X}}_i \mathbf{\Pi}_i + \overline{\mathbf{Y}}_i \\ &= 0. \end{aligned} \tag{15}$$

The identification problem is now reduced to determining whether the linear equation system has a unique solution.

Proposition 3 *Consider any player $i \in \mathbf{N}$ and suppose that (F, β) are given. If $\text{rank}(\overline{\mathbf{X}}_i) = (m_a \cdot m_s)$, then $\mathbf{\Pi}_i$ is exactly identified.*

The Proposition allows us to verify whether a set of linear payoff restrictions guarantees exact identification. It states that we need to verify whether the restriction together with the coefficient matrix \mathbf{X}_i satisfies a rank condition. Then, the payoff parameters are identified. If additional restrictions are imposed, then the payoff parameters are over-identified.

The set of restrictions imposed in equation (15) may depend on the economic problem at hand. We conclude this section with an example of two payoff restrictions that appear applicable in a number of settings.

Example: Suppose the number of state variables exceeds the number of actions, $L \geq K + 1$. Consider the following two restrictions:

$$\pi_i(a_i, \mathbf{a}_{-i}, s_i, \mathbf{s}_{-i}) = \pi_i(a_i, \mathbf{a}_{-i}, s_i, \mathbf{s}'_{-i}) \quad \forall \mathbf{a} \in \mathbf{A}, (s_i, \mathbf{s}_{-i}), (s_i, \mathbf{s}'_{-i}) \in \mathbf{S} \tag{R1}$$

$$\pi_i(0, \mathbf{a}_{-i}, s_i) = r_i(\mathbf{a}_{-i}, s_i) \quad \forall \mathbf{a}_{-i} \in \mathbf{A}_{-i}, s_i \in \mathbf{S}_i \tag{R2}$$

where \mathbf{r}_i is a $((K + 1)^{N-1} \cdot L) \times 1$ dimensional exogenous vector. *Restriction (R1)* says that period payoffs do not depend on the state variables of other firms. It fixes

$m_a \cdot (m_s - L)$ payoff parameters. Restriction (R1) is satisfied in games with adjustment costs such as entry or investment games. Restriction (R2) says that period payoffs under action $a_i = 0$ are fixed exogenously for every state s_i and action profile \mathbf{a}_{-i} . Restriction (R2) fixes $(K + 1)^{N-1} \cdot L$ payoff parameters. Restriction (R2) is satisfied in games in which one action, say action inactivity, implies zero payoffs. For given F and β , imposing restrictions (R1) and (R2) in equation (15) ensures identification of the payoff parameters provided the number of state variables exceeds the number of actions, $L \geq K + 1$, and the rank condition in Proposition 3 is satisfied. Observe also that a payoff normalization such as restriction (R2) is required for identification as the indifference analysis determines payoffs relative to the payoff under action $a_i = 0$ at every state only.

The illustration of the identification conditions and how to impose identifying restrictions allows us to proceed to the estimation problem. The next section proposes a class of asymptotic least squares estimators based on the equation system (11).

6 Estimation

We consider the class of asymptotic least squares estimators. This class provides a unified framework for a number of estimators proposed in the literature. We describe the asymptotically efficient weight matrix and discuss how a number of well known estimators can be recast as asymptotic least squares estimators.

6.1 Asymptotic Least Squares Estimators

Asymptotic least squares estimators are defined by using the equilibrium equation system equation (11) as estimating equations. They consist of two steps: In the first step, the auxiliary parameters consisting of the choice probabilities \mathbf{p} and the parameters entering the transition probability matrix \mathbf{G} are estimated. In the second

step, the parameters of interest are estimated by using weighted least squares on the equilibrium equation system.

Let $\boldsymbol{\theta} = (\boldsymbol{\theta}^\pi, \boldsymbol{\theta}^F, \beta, \boldsymbol{\theta}^g) \in \boldsymbol{\Theta} \subset \mathbb{R}^q$ denote an identified parameter vector.⁷ The parameter vector $\boldsymbol{\theta}$ may include parameters entering the period payoffs $\pi_i(\mathbf{a}, \mathbf{s}; \boldsymbol{\theta}^\pi)$ for all i , the transition probabilities $g(\mathbf{a}, \mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}^g)$, the distribution of payoff shocks $F(\varepsilon; \boldsymbol{\theta}^F)$ and the discount factor β . We introduce H auxiliary parameters consisting of the choice and state transition probabilities $\mathbf{p}(\boldsymbol{\theta})$ and $\mathbf{g}(\boldsymbol{\theta})$, with $H \leq (N \cdot K \cdot m_s) + (m_a \cdot m_s \cdot m_s)$. The auxiliary parameters are related to the parameters of interest through the system of $N \cdot K \cdot m_s$ implicit equations given by the equilibrium condition (11):

$$\mathbf{h}(\mathbf{p}, \mathbf{g}, \boldsymbol{\theta}) = \mathbf{p} - \Psi(\mathbf{p}, \mathbf{g}, \boldsymbol{\theta}) = 0 \quad (16)$$

The identification condition implies that the parameter vector $\boldsymbol{\theta}$ is determined without ambiguity from the auxiliary parameters \mathbf{p}, \mathbf{g} from the system of estimating equations (16). As explained in section 5 the equation system (16) is not the only representation of the equilibrium conditions. An equivalent representation exists in the space of expected discounted payoffs instead of choice probabilities, which is considered in Hotz, Miller, Sanders and Smith (1994) and is illustrated in Lemma 1 in the Appendix. The subsequent arguments can be made for any representation, but in the remainder of this section we select the representation given by equations (16).

Auxiliary parameters: We assume that estimators of the auxiliary parameters \mathbf{p}, \mathbf{g} exist and are consistent and asymptotically normally distributed. That is, we assume that there exists a sequence of estimators $(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T)$ of (\mathbf{p}, \mathbf{g}) such that

$$\begin{aligned} (\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T) &\longrightarrow (\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0)) \quad \text{a.s.} \\ \sqrt{T}((\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T) - (\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0))) &\xrightarrow{d} \mathcal{N}(0, \boldsymbol{\Sigma}(\boldsymbol{\theta}_0)) \end{aligned}$$

as $T \longrightarrow \infty$, where $\boldsymbol{\theta}_0$ denotes the true parameter vector.

⁷We include $\boldsymbol{\theta}^g$ in the parameter vector $\boldsymbol{\theta}$ to permit the possibility that a parametric restriction is placed on the state transition g .

There are a number of ways in which the auxiliary parameters consisting of the choice and transition probabilities can be estimated. Well known techniques include maximum likelihood or a Kernel smoother.⁸ A Monte Carlo study comparing these techniques in single agent dynamic models is conducted in Hotz, Miller, Sanders and Smith (1994). Maximum likelihood may be preferred if choices and state transitions are observed for all states, while Kernel smoothers may be preferred if some states are not observed. Let $p(k, i, s)$ denote the probability that player i selects action k in state s . Define the observed choice frequency by $n_{kis} = \sum_t \mathbf{1}(a_i^t = k, s^t = s)$ for all $i \in \mathbf{N}, k \in A_i, s \in S$ and observed state frequency by $n_{ass'} = \sum_t \mathbf{1}(a^t = a, s^t = s, s^{t+1} = s')$. By assumption, choices and states are multinomially distributed which implies that the maximum likelihood estimator equals the sample frequency:

$$\hat{p}(k, i, s) = \frac{n_{kis}}{\sum_{l \in \mathbf{A}_i} n_{lis}}, \quad \hat{g}(a, s, s') = \frac{n_{ass'}}{\sum_{s'' \in S} n_{ass''}}. \quad (17)$$

We denote with $\hat{\mathbf{p}} = (\hat{p}(k, i, s)_{i \in \mathbf{N}, k \in \mathbf{A}_i, s \in S})$ the vector of sample frequencies. Billingsley (1961) establishes that the maximum likelihood estimator is consistent and asymptotically efficient with a normal limiting distribution.

The asymptotic least squares estimating principle consists of estimating the parameters of interest $\boldsymbol{\theta}$ by forcing the constraints

$$\mathbf{h}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta}) = \hat{\mathbf{p}}_T - \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta}) = 0 \quad (18)$$

to be satisfied approximately. Consider initially the case in which the payoff parameter vector is identified exactly. The asymptotic least squares problem then has an explicit solution. Further, a simple argument based on the delta method shows that the explicit solution is an efficient estimator of $\boldsymbol{\theta}$, see Pesendorfer and Schmidt-Dengler (2003). With an over-identified parameter vector there are more equations than

⁸Grund (1993) provides conditions such that a Kernel estimator of cell probabilities has the assumed large sample properties.

parameters and there is no unique way to select a solution. A solution depends on the weight individual equations are given.

Let \mathbf{W}_T be a symmetric positive definite matrix of dimension $(N \cdot K \cdot m_s) \times (N \cdot K \cdot m_s)$ that may depend on the observations. An asymptotic least squares estimator associated with \mathbf{W}_T is a solution $\tilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$ to the problem

$$\min_{\boldsymbol{\theta} \in \Theta} [\hat{\mathbf{p}}_T - \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})]' \mathbf{W}_T [\hat{\mathbf{p}}_T - \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})]. \quad (19)$$

Thus, the asymptotic least squares estimator $\tilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$ brings the constraint $\hat{\mathbf{p}}_T - \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})$ closest to zero in the metric associated with the scalar product defined by \mathbf{W}_T .

Assumptions: Next, we state a number of assumptions that allow us to characterize the asymptotic properties of asymptotic least squares estimators. The assumptions are the following:

A1: Θ is a compact set.

A2: the true value $\boldsymbol{\theta}_0$ is in the interior of Θ .

A3: as $T \longrightarrow \infty$, $\mathbf{W}_T \longrightarrow \mathbf{W}_0$ a.s. where \mathbf{W}_0 is a non-stochastic positive definite matrix.

A4: $\boldsymbol{\theta}$ satisfies $[\mathbf{p}(\boldsymbol{\theta}_0) - \Psi(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta})]' \mathbf{W}_0 [\mathbf{p}(\boldsymbol{\theta}_0) - \Psi(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta})] = 0$ implies that $\boldsymbol{\theta} = \boldsymbol{\theta}_0$.

A5: the functions π, g, F are twice continuously differentiable in θ .

A6: the matrix $[\nabla_{\boldsymbol{\theta}} \Psi(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta}_0)]' \mathbf{W}_0 [\nabla_{\boldsymbol{\theta}'} \Psi(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta}_0)]$ is non-singular.

Assumptions A1-A3, A5, and A6 are standard technical conditions to ensure the problem is well behaved. Assumption A4 ensures that the parameter vector is identified.

The following Proposition shows that given the assumptions above, the asymptotic least squares estimator is consistent and asymptotically normally distributed. The Proposition follows from results developed in Gouriéroux, Monfort and Trognon

(1985) for asymptotic least squares estimators (see also Gouriou and Monfort, 1995, Theorem 9.1).

Proposition 4 *Under assumptions A1-A6, the asymptotic least squares estimator $\tilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$ exists, strongly converges to $\boldsymbol{\theta}_0$ and is asymptotically normally distributed with*

$$\sqrt{T} \left(\tilde{\boldsymbol{\theta}}_T(\mathbf{W}_T) - \boldsymbol{\theta}_0 \right) \xrightarrow{d} \mathcal{N}(0, \boldsymbol{\Omega}(\boldsymbol{\theta}_0))$$

as $T \longrightarrow \infty$, where

$$\begin{aligned} \boldsymbol{\Omega}(\boldsymbol{\theta}_0) = & \left(\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \mathbf{W}_0 \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi} \right)^{-1} \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \mathbf{W}_0 \left[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi} \right] \boldsymbol{\Sigma} \\ & \cdot \left[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi} \right]' \mathbf{W}_0 \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi} \left(\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \mathbf{W}_0 \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi} \right)^{-1} \end{aligned} \quad (20)$$

where $\mathbf{0}$ denotes a $(N \cdot K \cdot m_s) \times (m_a \cdot (m_s)^2)$ matrix of zeros and the various matrices are evaluated at $\boldsymbol{\theta}_0$, $\mathbf{p}(\boldsymbol{\theta}_0)$, and $\mathbf{g}(\boldsymbol{\theta}_0)$.

The Proposition shows that there are a number of consistent and asymptotically normally distributed estimators, each one corresponding to a particular sequence of matrices \mathbf{W}_T .⁹

Next, we address the optimal choice of matrix \mathbf{W}_T .

6.2 Efficient Asymptotic Least Squares Estimator

In this section we describe the optimal weight matrix. We make an additional assumption:

A7: *the matrices*

$$\begin{aligned} & \left[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi} \right] \boldsymbol{\Sigma} \left[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi} \right]' \text{ and} \\ & \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \left(\left[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi} \right] \boldsymbol{\Sigma} \left[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi} \right]' \right)^{-1} \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi} \end{aligned}$$

⁹If two distinct auxiliary estimators, $\hat{\mathbf{p}}_T^1$ and $\hat{\mathbf{p}}_T^2$, are employed in the equilibrium conditions (18), $\hat{\mathbf{p}}_T^1 - \boldsymbol{\Psi}(\hat{\mathbf{p}}_T^2, \hat{\mathbf{g}}_T, \boldsymbol{\theta}) = 0$, then the matrix matrix $\boldsymbol{\Sigma}$ in Proposition 4 refers to the asymptotic covariance matrix associated with the vector of auxiliary parameters $(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T) = (\hat{\mathbf{p}}_T^1, \hat{\mathbf{p}}_T^2, \hat{\mathbf{g}}_T)$.

are non-singular where $\mathbf{0}$ denotes a $(N \cdot K \cdot m_s) \times (m_a \cdot (m_s)^2)$ matrix of zeros and $\Sigma, \nabla_\theta \Psi$ and $\nabla'_{(p,g)} \Psi$ are evaluated at θ_0 .

The optimal weight matrix for asymptotic least squares estimators follows from Gourieroux, Monfort and Trognon (1985). The following Proposition states the result:

Proposition 5 *Under assumptions A1-A7, the best asymptotic least squares estimators exist. They correspond to sequences of matrices \mathbf{W}_T^* converging to*

$$\mathbf{W}_0^* = \left([(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi] \Sigma [(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi]' \right)^{-1}.$$

Their asymptotic variance covariance matrices are equal and given by

$$\Omega(\mathbf{W}^*) = \left(\nabla_\theta \Psi' \left([(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi] \Sigma [(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi]' \right)^{-1} \nabla_{\theta'} \Psi \right)^{-1} \quad (21)$$

Proposition 5 establishes that the asymptotically optimal weights \mathbf{W}_0^* depend on the derivative of the estimating equations with respect to the auxiliary parameters (\mathbf{p}, \mathbf{g}) and the covariance matrix of the auxiliary parameters Σ .

The next subsection shows that a number of well known estimators are members of the class of asymptotic least squares estimators.

6.3 Examples of Asymptotic Least Squares Estimators

This section considers moment estimators and pseudo maximum likelihood estimators.

6.3.1 Moment Estimator

A moment estimator is proposed in Hotz and Miller (1993). It is based on the moment condition for the $K \times 1$ dimensional random vector of choices $\alpha_{is} = (\alpha_1, \dots, \alpha_K)$ for player i in state s multiplied by an instrument vector. A realization α^t of the random variable α_{is} is a $K \times 1$ dimensional vector with entry k equal to 1 if action k is adopted and 0 otherwise, so that $\sum_{k=1}^K \alpha_k^t \in \{0, 1\}$ and $\hat{\alpha}_k \in \{0, 1\}$ for all k . Let \mathbf{T}_{is} denote

the set of observations for player i in state s . The orthogonality condition is defined as:

$$E [Z \otimes [\boldsymbol{\alpha}_{is} - \Psi_{is}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})]] = 0 \quad (22)$$

where Z is a $J \times 1$ dimensional instrument vector exogenous to the term in the second square bracket, and \otimes denotes the Kronecker product.¹⁰ It is easy to see that the moment estimator based on the moment condition (22) is an asymptotic least squares estimator when setting $Z^t = Z^{is}$ for all $t \in \mathbf{T}_{is}$ and for all $i \in \mathbf{N}, s \in \mathbf{S}$. The reason is that the inner sum of the sample analog of (22) can be rewritten as a weighted average of the estimating equation:

$$\frac{1}{NT} \sum_{i \in \mathbf{N}, s \in \mathbf{S}} \sum_{t \in \mathbf{T}_{is}} Z^t \otimes [\boldsymbol{\alpha}^t - \Psi_{is}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})] = \frac{1}{NT} \sum_{i \in \mathbf{N}, s \in \mathbf{S}} n_{is} [Z^{is} \otimes [\hat{\mathbf{p}}_{is} - \Psi_{is}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})]] .$$

as $\sum_{t \in \mathbf{T}_{is}} \boldsymbol{\alpha}^t$ equals the vector of sample frequencies $\hat{\mathbf{p}}_{is} = (n_{1is}/n_{is}, \dots, n_{Kis}/n_{is})$ defined in equation (17). So, the estimating equation of the equivalent asymptotic least squares estimator is $\hat{\mathbf{p}} - \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta}) = 0$.¹¹ The weights assigned to the individual estimating equations depend on the vector of instruments \mathbf{Z} , and the choice of a $(J \cdot K)$ square weight matrix associated with the moment conditions.¹²

Current and lagged state variables are selected as instruments in a number of applications including Hotz and Miller (1993), Slade (1998) and Aguirregabiria (1999). These instrument matrices differ from the optimal weight matrix as they generate a set of weighted moment averages with weights determined by the level of the state variables. In contrast, the optimal choice would single out a separate moment condition for every state. One way to achieve efficiency is to set the number of instruments

¹⁰Hotz, Miller, Sanders and Smith (1994) consider an alternative representation of the equilibrium in the payoff space. As explained earlier, this representation is equivalent to our formulation in the choice probability space.

¹¹Observe, that the frequency estimator $\hat{\mathbf{p}}$ is used in the first appearance of the auxiliary parameter in the estimating equation while a possibly distinct parameter $\hat{\mathbf{p}}_T$ may appear in the second place.

¹²The condition $Z^t = Z^{is}$ for all $t \in T_{is}$ and for all $i \in \mathbf{N}, s \in \mathbf{S}$ is not necessary for the moment estimator to be an asymptotic least squares estimator. If $Z^t \neq Z^{is}$ for some $t \in T_{is}$, then the frequency estimator $\hat{\mathbf{p}}$ would be replaced by a different estimator $\tilde{\mathbf{p}}$.

equal to $J = m_s \cdot N$, where the element (i, s) of Z^t is equal to $\frac{1}{n_{is}}$, if $t \in \mathbf{T}_{is}$ and zero if $t \notin \mathbf{T}_{is}$. In this manner $N \cdot m_s \cdot K$ moment conditions are generated, corresponding to the equilibrium conditions in (18). It remains to choose the $(m_s \cdot N)$ square weight matrix appropriately equal to \mathbf{W}_0^* as defined in Proposition 5, and the moment estimator corresponds to the efficient asymptotic least squares estimator.

Hansen (1982) and Chamberlain (1987) derive the optimal choice of instruments

trovsky and Berry (2004),¹⁴ propose moment estimators that are asymptotic least squares estimators with weight matrices equalling the identity matrix or a diagonal matrix with elements equal to the inverse of the number of observations per state. For over-identified models as considered in these papers, these weight matrices do not take the equilibrium equation system or the covariance of the auxiliary parameters into account optimally.

6.3.2 Pseudo Maximum Likelihood Estimator

The pseudo maximum likelihood estimator, or PML, maximizes the pseudo log likelihood. Aguirregabiria and Mira (2002) consider the *partial* pseudo log likelihood conditional on the transition probability estimates $\hat{\mathbf{g}}_T$. It is given by:

$$\ell = \sum_{s \in \mathbf{S}} \sum_{i \in \mathbf{N}} \sum_{k \in \mathbf{A}_i} n_{kis} \log \Psi_{kis}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})$$

where Ψ_{kis} is evaluated at $(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T)$. The first order condition of the *partial* pseudo maximum likelihood estimator is equivalent to the first order condition of the asymptotic least squares estimator defined in (19) with weight matrix \mathbf{W}_T^{ml} converging to the inverse of the covariance matrix of the choice probabilities $\boldsymbol{\Sigma}_p^{-1}$. The reason is that the first order condition is given by

$$\frac{\partial \ell}{\partial \boldsymbol{\theta}} = (\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}') \boldsymbol{\Sigma}_p^{-1}(\boldsymbol{\Psi}) [\hat{\mathbf{p}} - \boldsymbol{\Psi}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})].$$

as is shown in the appendix. Observe that for auxiliary estimates not equalling the sample frequency, $\hat{\mathbf{p}}_T \neq \hat{\mathbf{p}}$, the PML uses two distinct auxiliary parameter estimators of the choice probabilities, $\hat{\mathbf{p}}$ and $\hat{\mathbf{p}}_T$, and the equivalent asymptotic least squares estimator solves the equation system $\hat{\mathbf{p}} - \boldsymbol{\Psi}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta}) = 0$. The first appearance of \mathbf{p}

¹⁴Pakes, Ostrovsky and Berry (2004) apply two moment estimators in their Monte Carlo study of an entry model. The first moment estimator is defined by the average entry (and exit) rate across a fixed subset of states; and the second equals the entry and exit rate at a state weighted by one over the number of observations for that state, and again defined on a fixed subset of states.

in the equation system is evaluated at $\hat{\mathbf{p}}$ and the second at $\hat{\mathbf{p}}_T$. The intuition is that the PML sets the sample frequency $\hat{\mathbf{p}}$ equal to the pseudo likelihood Ψ evaluated at $(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})$. This feature is shared with the moment estimator.¹⁵

The optimal weight matrix of the *partial* asymptotic least squares estimator conditional on the transition probability estimates $\hat{\mathbf{g}}_T$ differs from the one of the *partial* pseudo maximum likelihood estimator, which equals Σ_p^{-1} , by a term that accounts for the derivative of the estimating equations with respect to the auxiliary parameters \mathbf{p} . It is manifested in the term $\nabla_{p'} \Psi$. Aguirregabiria and Mira (2002) show that the derivative $\nabla_p \Psi$ vanishes in the single agent case, as the continuation value, which enters Ψ and is defined in equation (7), achieves a maximum at the optimal choice probabilities \mathbf{p} . The pseudo maximum likelihood estimator is then optimal. In the multi agent case, as is shown in Aguirregabiria and Mira (2004), the derivative $\nabla_p \Psi$ need not vanish, as player j 's choices do not maximize player i 's continuation value. The weight matrix of the conditional pseudo maximum likelihood estimator is then not optimal.

So far, we have described large sample properties of asymptotic least squares estimators based on the equilibrium equation system. We have shown that a number of estimators proposed in the literature are asymptotic least squares estimators. The estimators differ in the weight they assign to individual equilibrium equations. In each case, we have characterized the associated weight matrix. In the next section, we conduct a Monte Carlo study to compare the performance of a number of estimators when the sample size is small.

¹⁵Aguirregabiria and Mira (2002) also introduce *the iterated pseudo maximum likelihood estimator*, or k-PML. It is based on the observation that the choice probability vector $\Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})$, evaluated at the PML estimate $\hat{\boldsymbol{\theta}}_1$, defines a new auxiliary parameter estimate, $\hat{\mathbf{p}}_1 = \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \hat{\boldsymbol{\theta}}_1)$. This process can be iterated.

7 Monte Carlo Study

This section examines the practical effect of alternative weight matrices in a Monte Carlo study. We also illustrate the multiplicity of equilibria inherent to games.

We compare the performance of four asymptotic least squares estimators: These estimators are: (i) the efficient estimator, LS-E;¹⁶ (ii) the pseudo maximum likelihood estimator, PML; (iii) the estimator using the identity matrix as weights, LS-I; (iv) the k-step iterated pseudo maximum likelihood estimator, k-PML, proposed by Aguirregabiria and Mira (2002,2004). We fix the number of iterations in the k-step procedures at 20, although convergence may be achieved at earlier iterations.

We select a simple model design and assess the performance of the estimators in symmetric and asymmetric equilibria with distinct number of observations yielding a total of 12 Monte Carlo experiments. We use the frequency estimator for the auxiliary parameters. We refer the interested reader to the Monte Carlo studies by Hotz and Miller (1994), Aguirregabiria and Mira (2002,2004) and Pakes, Ostrovsky and Berry (2004) for the effect of alternative auxiliary estimators.

We describe the design of the Monte Carlo study in more detail in subsection 7.1. Subsection 7.2 reports the results and comments on the findings.

7.1 Design

To keep the study simple and transparent, we consider a setting with two players, binary actions $\{0, 1\}$ and binary states $\{0, 1\}$. We have conducted exercises for problems involving more players, states and actions yielding similar results at increased computational costs. The payoff structure and model parameters are selected to

¹⁶The efficient weight matrix is calculated by using a two step procedure. In the initial step, the profit parameters are estimated consistently by using the identity matrix as weights. The efficient weight matrix is then constructed by using the first stage estimates and is used in the second stage to obtain the estimates.

imitate the empirical entry application in Pesendorfer and Schmidt-Dengler (2003). Variations in the model parameters achieve results similar to the ones reported here, and we report results for our baseline specification only.

The specification has the following features: The distribution of the profitability shocks F is the standard normal. The discount factor is fixed at 0.9. The state transition law is given by $s_i^{t+1} = a_i$. Period payoffs are symmetric and are parametrized as follows:

$$\pi(a_i, a_j, s_i) = \begin{cases} 0 & \text{if } a_i = 0; s_i = 0 \\ x & \text{if } a_i = 0; s_i = 1 \\ \pi^1 + c & \text{if } a_i = 1; a_j = 0; s_i = 0 \\ \pi^2 + c & \text{if } a_i = 1; a_j = 1; s_i = 0 \\ \pi^1 & \text{if } a_i = 1; a_j = 0; s_i = 1 \\ \pi^2 & \text{if } a_i = 1; a_j = 1; s_i = 1 \end{cases}$$

where $x = 0.1$; $c = -0.2$; $\pi^1 = 1.2$; and $\pi^2 = -1.2$. The period payoffs can be interpreting as stemming from a game with switching costs and/or as entry/exit game. A player that selects action 1 receives monopoly profits π^1 if she is the only active player, and she receives duopoly profits π^2 otherwise. Additionally, a player that switches states from 0 to 1 incurs the entry cost c ; while a player that switches from 1 to 0 receives the exit value x .

Multiplicity. The game illustrates the possibility of multiple equilibria which is a feature inherent to games. We illustrate five equilibria, four of which are asymmetric and one of which is symmetric. The equilibria have the following distinguishing features: In equilibrium (i), player one is more likely to choose action 0 than action 1 in all states. The ex ante probability vectors for both players are given by: $p(a_1 = 0|s_1, s_2) = (0.27, 0.39, 0.20, 0.25)'$, $p(a_2 = 0|s_2, s_1) = (0.72, 0.78, 0.58, 0.71)'$ where the order of the elements in the probability vectors correspond to the state vector $(s_1, s_2) = ((0, 0), (0, 1), (1, 0), (1, 1))$.

In equilibrium (ii), player two is more likely to choose action 0 than action 1 in

all states with the exception of state $(1, 0)$. The probability vectors are given by $p(a_1 = 0|s_1, s_2) = (0.38, 0.69, 0.17, 0.39)'$, $p(a_2 = 0|s_2, s_1) = (0.47, 0.70, 0.16, 0.42)'$.

Equilibrium (iii) is symmetric. The probability vectors are given by $p(a_1 = 0|s_1, s_2) = (0.42, 0.70, 0.16, 0.41)'$, $p(a_2 = 0|s_2, s_1) = (0.42, 0.70, 0.16, 0.41)'$. Two additional equilibria have the property that the identities of players 1 and 2 in equilibria (i) and (ii) are reversed. We consider equilibria (i)-(iii) in the subsequent analysis.

Identification and Estimation. The game has four distinct states and permits the exact identification of at most four parameters per player. We impose the restrictions $x = 0.1$, and that the parameters are identical for both players. We estimate three parameters: c , π^1 , and π^2 .

The simulated data are generated by randomly drawing a time series of actions from the calculated equilibrium choice probabilities described above for each of the equilibria (i)-(iii) respectively. The initial state is taken as $(0, 0)$ and we start the sampling process after 250 periods. The length of the time series is varied in the experiment with T equalling 100, 1,000, 10,000 and 100,000. For each design we conduct 1,000 repetitions of the experiment in order to obtain a distribution of the estimates. For each estimator we report the mean, standard deviation and mean squared error based on the simulated distribution of parameter estimates.

7.2 Results

Tables 1-3 summarize our Monte Carlo results for equilibria (i)-(iii). In total there are 12 specifications. We report the mean, the standard deviation and the mean squared error, MSE. The MSE is summed over the three parameters and is scaled by a factor of 100.

[Tables 1-3 about here]

The LS-E estimator is selected by the MSE criterion in eight of 12 cases as the preferred estimator. Overall the LS-E is the best performing estimator for all three

equilibria. The LS-E does not perform well in small sample sizes with T equalling 100 observations, but is preferred in moderate to large sample sizes with T equalling 1,000, 10,000 and 100,000 observations. A possible reason for the poor performance with small T are imprecisions in the estimated optimal weight matrix.

The PML is ranked second according to the MSE criterion in seven of 12 specifications. The PML performs better than the LS-E for small sample sizes with T equalling 100 and worse for larger sample sizes. A reason for the good performance in small sample sizes may be that the covariance matrix of the auxiliary parameters is more accurately estimated than the optimal weight matrix. The PML performs better than the LS-I and k-PML in equilibrium (ii), but performs worse than the LS-I and k-PML in equilibrium (i). Both the PML and the LS-I perform similarly to the LS-E in equilibrium (iii) for moderate to large sample sizes. This may be a special feature of equilibrium (iii) as the equilibrium choice probabilities are symmetric and close to one half. As a result, the weight matrices of the LS-E, PML and LS-I have similar properties.

The LS-I is ranked third in nine of 12 specifications according to the MSE criterion. It performs better than the PML in equilibrium (i) for sample sizes with T larger than 100, but worse than the LS-E and k-PML. Overall, the MSE criterion does not select the LS-I in any specification as the preferred estimator. The weak performance may be attributable to a larger standard deviation of the estimates. The LS-I does have a smaller bias than the LS-E and PML in equilibria (ii) and (iii) for small sample sizes with T equalling 100 and 1000.

The k-PML is ranked 4th on six of 12 specifications according to the MSE criterion. It has a low MSE for T equalling 100, but this low MSE is attributable to a low standard deviation. The k-PML is severely biased for small, moderate and large T in equilibria (ii) and (iii). The bias increases when we consider more than 20 iterations. The k-PML performs well in equilibrium (i) with a MSE only about 20-30 higher than that of the LS-E for moderate to large sample sizes with T equalling at least 1,000.

The Monte Carlo study shows that the efficient LS-E estimator may be the preferred estimator for moderate to large sample sizes. For small sample sizes the PML or the LS-I may be preferred alternatives, as they do not require the estimation of the optimal weight matrix. The iterated k-PML performs reasonably well in equilibrium (i), but has a strong bias in equilibria (ii) and (iii).

8 Concluding Remarks

This paper considers the class of asymptotic least squares estimators for dynamic games. The basis of our analysis is an equation system characterizing the equilibrium choice probabilities. This equilibrium equation system permits us to derive identification conditions of the underlying model parameters. Asymptotic least squares estimators minimize the weighted distance between estimated choice probabilities and choice probabilities implied by the model. We show that a number of well known estimators fall into the class of asymptotic least squares estimators. These estimators differ in the weight they assign to individual equilibrium conditions. We characterize the efficient weight matrix. Asymptotic least squares estimators for dynamic games are simple to implement. We conducted a small scale Monte Carlo study illustrating the small sample performances of alternative weight matrices. We found benefits to using the efficient weight matrix in moderate and large sample sizes.

There are several directions for future research. One is to relax the assumption that the econometrician has access to the same information on state variables as the players. It appears possible that our proposed estimator is extendable to such environments, as long as consistent estimates of players' beliefs and state transitions are obtainable. Another extension is to explore a game with continuous action space. Such a game may entail a continuum of equilibrium conditions, and as discussed in section 6, efficient estimation may require the inversion of a covariance operator in an infinite dimensional Hilbert space.

9 Appendix

This appendix contains proofs.

Proof of Proposition 1. First, we show that equation (11) must be satisfied in any equilibrium. An action $a_i \in \mathbf{A}_i$ is optimal for player i in state s given beliefs σ_i if it yields a continuation value at least as large as any other action $a'_i \in \mathbf{A}_i$:

$$u_i(a_i; \sigma_i, \boldsymbol{\theta}) + \varepsilon_i^{a_i} \geq u_i(a'_i; \sigma_i, \boldsymbol{\theta}) + \varepsilon_i^{a'_i} \text{ for all } a'_i \in \mathbf{A}_i \quad (23)$$

where $\varepsilon_i^0 = 0$. The optimality condition (23) can be expressed before the payoff shock ε_i is observed. Doing so, gives an expression for the probability that player i chooses action $a_i \in \mathbf{A}_i$ in state s given the beliefs σ_i :

$$\begin{aligned} p_i(a_i | \mathbf{s}) &= \Psi_i(a_i, \mathbf{s}, \sigma_i, \boldsymbol{\theta}) \\ &= \int \prod_{k \in \mathbf{A}_i, k \neq a_i} \mathbf{1}(u_i(a_i; \sigma_i, \boldsymbol{\theta}) - u_i(k; \sigma_i, \boldsymbol{\theta}) \geq \varepsilon_i^k - \varepsilon_i^{a_i}) dF \end{aligned} \quad (24)$$

Equation (24) is a necessary equilibrium condition that must be satisfied at every state $s \in \mathbf{S}$, for every player i , and for every action $a_i = 1, \dots, K$, yielding a total of $m_s \cdot K \cdot N$ equations, one for each possible state, action and player. We may compactly write equation (24) in matrix form as:

$$\mathbf{p} = \boldsymbol{\Psi}(\boldsymbol{\sigma}; \boldsymbol{\theta}) \quad (25)$$

where $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_N)$ denotes the $(N \cdot K \cdot m_s) \times 1$ dimensional vector of the optimal players' choice probabilities for all states, players and actions other than action 0; and $\boldsymbol{\sigma}$ denotes the $(N \cdot K \cdot m_s) \times 1$ dimensional vector of players' beliefs. We omit the choice probability for action 0 in equation (25) as it is already determined by the remaining choice probabilities,

$$p_i(0 | \mathbf{s}) = 1 - \sum_{k=1}^K p_i(k | \mathbf{s}).$$

In equilibrium, beliefs must be consistent yielding equation (11).

Next, we show that any \mathbf{p} that satisfies equation (11) can be extended to construct a Markov perfect equilibrium. Let the beliefs of player i equal p_i and define the decision rule of player i based on condition (23). That is, suppose player i chooses action $a_i \in \mathbf{A}_i$ in state s when the continuation value under action a_i is at least as large as the continuation value under any action $a'_i \in \mathbf{A}_i$. In case of equal continuation values of alternative actions we assume the player selects the smallest action:

$$a_i(\varepsilon_i, s) = \min\{k \in K_i(\varepsilon_i, s)\}$$

$$\text{where } K_i(\varepsilon_i, s) = \{k \in \mathbf{A}_i | u_i(k; p_i, \boldsymbol{\theta}) + \varepsilon_i^k \geq u_i(k'; p_i, \boldsymbol{\theta}) + \varepsilon_i^{k'} \text{ for all } k' \in \mathbf{A}_i\}$$

and $\varepsilon_i^0 = 0$. By construction, the decision rule and the beliefs are an equilibrium as the decision rule is optimal given the beliefs p_i for all i , and the beliefs are consistent. This completes the proof. ■

Proof of Theorem 1. We need to show that equation (11) has a solution. The choice probabilities p are contained in the unit interval. The function Ψ is continuous in p . Brouwer's fixed point theorem implies that there exists a fixed point p of the function Ψ . By Proposition 1, the fixed point corresponds to an equilibrium. ■

Proof of Corollary 1. The arguments in the proofs of Proposition 1 and Theorem 1 do not rely on asymmetry. The same arguments remain valid with the symmetry assumption in place. ■

Lemma 1 (Equilibrium Characterization Linear in Payoffs) *There exist $(\mathbf{X}_i(\mathbf{p}; g, \beta), \mathbf{Y}_i(\mathbf{p}; g, \beta))_{i=1}^N$ such that any \mathbf{p} that satisfies (11) solves*

$$\mathbf{X}_i(\mathbf{p}; g, \beta)\Pi_i + \mathbf{Y}_i(\mathbf{p}; g, \beta) = 0 \quad \text{for all } i \in \mathbf{N} \quad (26)$$

and any \mathbf{p} that solves (26) also satisfies (11).

Proof of Lemma 1. We shall consider any player i and show that there exists an equivalent representation to the best response equations in (11). The representation is linear in payoffs.

The alternative representation is based on the optimality condition (8) for the action pair a_i and 0. Let $\bar{\varepsilon}_i^{a_i}(\mathbf{s})$ be the type that is indifferent between actions a_i and 0 in state s . The equation characterizing the indifferent type is linear in payoffs and given by:

$$\begin{aligned} & \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} p(\mathbf{a}_{-i}|\mathbf{s}) \cdot \left[\pi_i(\mathbf{a}_{-i}, a_i, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{a}_{-i}, a_i, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; p) \right] + \bar{\varepsilon}_i^{a_i}(\mathbf{s}) \\ = & \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} p(\mathbf{a}_{-i}|\mathbf{s}) \cdot \left[\pi_i(\mathbf{a}_{-i}, 0, \mathbf{s}) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{a}_{-i}, 0, \mathbf{s}, \mathbf{s}') V_i(\mathbf{s}'; p) \right] \end{aligned}$$

By transitivity any type $(\varepsilon_i^{a_i}, \varepsilon_i^{a'_i})$ will prefer action a_i over action a'_i in state \mathbf{s} if $\varepsilon_i^{a_i} - \bar{\varepsilon}_i^{a_i}(\mathbf{s}) > \varepsilon_i^{a'_i} - \bar{\varepsilon}_i^{a'_i}(\mathbf{s})$. Thus, the equilibrium decision rule for player i with type ε_i in state \mathbf{s} can be written as

$$a_i(\varepsilon_i, \mathbf{s}) = \begin{cases} k & \text{if } \varepsilon_i^k > \bar{\varepsilon}_i^k(\mathbf{s}) \text{ and for all } k' \neq k: \varepsilon_i^k - \varepsilon_i^{k'} > \bar{\varepsilon}_i^k(\mathbf{s}) - \bar{\varepsilon}_i^{k'}(\mathbf{s}); \\ 0 & \text{if } \varepsilon_i^k < \bar{\varepsilon}_i^k(\mathbf{s}) \text{ for all } k. \end{cases} \quad (27)$$

and the choice probability is given by:

$$p(a_i = k | s) = \Pr \left(\varepsilon_i^k > \bar{\varepsilon}_i^k(s) \text{ and for all } k' \neq k: \varepsilon_i^k - \varepsilon_i^{k'} > \bar{\varepsilon}_i^k(s) - \bar{\varepsilon}_i^{k'}(s) \right) \quad (28)$$

for all $i \in \mathbf{N}, k \in A_i, s \in S$. Observe that there is a one-to-one relationship between choice probabilities and indifferent types based on equation (28), as is originally shown in Hotz and Miller (1993). Our earlier working paper, Pesendorfer and Schmidt-Dengler (2003), includes also a proof for this statement. Thus, any vector of choice probabilities implies a unique vector of indifferent types and vice versa.

To obtain the representation stated in the Lemma we write the equations characterizing the indifferent types in matrix notation. We introduce the following notation. Let $\bar{\varepsilon}_i = [\bar{\varepsilon}_i^k(\mathbf{s})]_{s \in \mathbf{S}, k=1, \dots, K}$ be the $(m_s \cdot K) \times 1$ dimensional vector of indifferent types and let \mathbf{P}_{-i} be the $m_s \times ((K+1)^{N-1} \cdot m_s)$ dimensional matrix consisting of choice probability $p(\mathbf{a}_{-i}|\mathbf{s})$ in row \mathbf{s} column $(\mathbf{a}_{-i}, \mathbf{s})$, and zeros in row \mathbf{s} column $(\mathbf{a}_{-i}, \mathbf{s}')$ with

$\mathbf{s}' \neq \mathbf{s}$; and define the $(K \cdot m_s) \times (m_a \cdot m_s)$ dimensional matrix $\mathbf{P}_i(\bar{\varepsilon})$ as:

$$\mathbf{P}_i(\bar{\varepsilon}) = \begin{matrix} & \begin{matrix} 0 & 1 & 2 & \dots & K \end{matrix} \\ \begin{matrix} -\mathbf{P}_{-i} & \mathbf{P}_{-i} & 0 & \dots & 0 \\ -\mathbf{P}_{-i} & 0 & \mathbf{P}_{-i} & \dots & 0 \\ \vdots & 0 & 0 & \ddots & 0 \\ -\mathbf{P}_{-i} & 0 & 0 & \dots & \mathbf{P}_{-i} \end{matrix} \end{matrix}$$

We can re-state the indifference equations for player i as:

$$\mathbf{P}_i(\bar{\varepsilon}) \boldsymbol{\Pi}_i + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} \mathbf{V}_i + \bar{\boldsymbol{\varepsilon}}_i = 0 \quad (29)$$

where $\bar{\boldsymbol{\varepsilon}} = [\bar{\varepsilon}_i]_{i \in N}$ is the $(K \cdot m_s \cdot N) \times 1$ dimensional vector of indifferent types. Using the expression for the ex ante expected value function, equation (6), we obtain the following equation system:

$$\begin{aligned} & [\mathbf{P}_i(\bar{\varepsilon}) + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} [\mathbf{I}_s - \beta \mathbf{p}_i(\bar{\varepsilon}) \mathbf{G}]^{-1} \mathbf{p}_i] \boldsymbol{\Pi}_i \\ & + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} [\mathbf{I}_s - \beta \mathbf{p}_i(\bar{\varepsilon}) \mathbf{G}]^{-1} \mathbf{D}_i(\bar{\varepsilon}) + \bar{\boldsymbol{\varepsilon}}_i = 0 \end{aligned} \quad (30)$$

Equation system (30) consists of $m_s \cdot K$ equations for player i . Notice also that the ex ante expected payoff shock can be uniquely written in terms of indifferent types, as:

$$D_i(\mathbf{s}) = \sum_{k=1}^K \int_{\bar{\varepsilon}_i^k(s)}^{\infty} \varepsilon^k \prod_{k' \geq 1, k' \neq k} F(\varepsilon^k + \bar{\varepsilon}^{k'} - \bar{\varepsilon}^k) f(\varepsilon^k) d\varepsilon^k.$$

Thus, we have shown that for any player i , there is an equilibrium characterization (30) consisting of $K \cdot m_s$ equations of the form stated in the Lemma and in which

$$\mathbf{X}_i = [\mathbf{P}_i(\bar{\varepsilon}) + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} [\mathbf{I}_s - \beta \mathbf{p}_i(\bar{\varepsilon}) \mathbf{G}]^{-1} \mathbf{p}_i]$$

and

$$\mathbf{Y}_i = \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} [\mathbf{I}_s - \beta \mathbf{p}_i(\bar{\varepsilon}) \mathbf{G}]^{-1} \mathbf{D}_i(\bar{\varepsilon}) + \bar{\boldsymbol{\varepsilon}}_i$$

and the unknown parameters are in Π_i .

Now, take a vector \mathbf{p} that satisfies (11). We can recover the indifferent types $\bar{\varepsilon}$ based on equation (28). By optimality of \mathbf{p} the indifferent types must satisfy the indifference equations (26).

Conversely, take a vector $\bar{\varepsilon}$ that satisfies (26). We can calculate the implied ex ante choice probabilities based on equation (28). The choice probabilities will satisfy (11) as they are the ex ante expected values of equilibrium decision rules satisfying the equilibrium conditions (26). This completes the proof. ■

Proof of Proposition 2. The equation system (26) characterized in Lemma 1 consists of $K \cdot m_s \cdot N$ equations. Hence, at most $K \cdot m_s \cdot N$ parameters in (Π_1, \dots, Π_N) can be identified. ■

Proof of Proposition 3. Consider the equation system (26) in Lemma 1. By construction, the coefficient matrix \mathbf{X}_i is of dimension $(K \cdot m_s) \times (m_a \cdot m_s)$ and the vector \mathbf{Y}_i is of dimension $(K \cdot m_s) \times 1$. Augmenting the linear equation system (14) with a set of $m_a \cdot m_s - K \cdot m_s$ payoff restrictions $\mathbf{R}_i \cdot \Pi_i = \mathbf{r}_i$, yields the $m_a \cdot m_s$ equations in $m_a \cdot m_s$ unknown payoff parameters Π_i

$$\begin{bmatrix} \mathbf{X}_i \\ \mathbf{R}_i \end{bmatrix} \Pi_i + \begin{bmatrix} \mathbf{Y}_i \\ \mathbf{r}_i \end{bmatrix} = \bar{\mathbf{X}}_i \Pi_i + \bar{\mathbf{Y}}_i \quad (31)$$

$$= 0. \quad (32)$$

The identification problem is now reduced to determining whether the linear equation system has a unique solution. Standard arguments for linear equation systems establish that there exists exactly one solution if the rank of the $(m_a \cdot m_s) \times (m_a \cdot m_s)$ matrix $\bar{\mathbf{X}}_i$ is equal to $(m_a \cdot m_s)$. Then Π_i is exactly identified. ■

Proof of Proposition 4. Assumption A5 implies that the function h is twice continuously differentiable. All the assumptions in Theorems 9.1 and 9.2 on pages 280 and 281 in Gourieroux and Monfort (1995) are satisfied. The result follows from those theorems. ■

Proof of Proposition 5. The result follows from property 9.1 on page 282 in Gourieroux and Monfort (1995). ■

Proof that the pseudo maximum likelihood estimator is equivalent to an asymptotic least squares estimator. To see the equivalence, observe that the first order condition of the pseudo maximum likelihood estimator is given by:

$$\begin{aligned}
\frac{\partial \ell}{\partial \boldsymbol{\theta}} &= \sum_{s \in \mathbf{S}} \sum_{i \in \mathbf{N}} \sum_{k=1}^K \left[\frac{n_{kis}}{\Psi_{kis}} - \frac{n_{0is}}{\Psi_{0is}} \right] \frac{\partial \Psi_{kis}}{\partial \boldsymbol{\theta}} \\
&= \sum_{s \in \mathbf{S}} \sum_{i \in \mathbf{N}} \sum_{k=1}^K \left[\sum_{l=0}^K n_{lis} \right] \left[\frac{\hat{p}_{kis} - \Psi_{kis}}{\Psi_{kis}} - \frac{\hat{p}_{0is} - \Psi_{0is}}{\Psi_{0is}} \right] \frac{\partial \Psi_{kis}}{\partial \boldsymbol{\theta}} \\
&= \sum_{s \in \mathbf{S}} \sum_{i \in \mathbf{N}} \sum_{k=1}^K \left[\sum_{l=0}^K n_{lis} \right] \left[\frac{1}{\Psi_{kis}} (\hat{p}_{kis} - \Psi_{kis}) + \sum_{l=1}^K \frac{1}{\Psi_{0is}} (\hat{p}_{lis} - \Psi_{lis}) \right] \frac{\partial \Psi_{kis}}{\partial \boldsymbol{\theta}} \\
&= (\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}') \boldsymbol{\Sigma}_p^{-1}(\boldsymbol{\Psi}) [\hat{\mathbf{p}} - \boldsymbol{\Psi}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})]
\end{aligned}$$

where the second equality uses the definition of the frequency estimator $\hat{p}_{kis} = n_{kis} / \sum_{l=0}^K n_{lis}$ from equation (17), and augments the first term in the second square bracket by minus one and the second term by plus one; the third equality uses the definitions $\hat{p}_{0is} = 1 - \sum_{k=1}^K \hat{p}_{kis}$ and $\Psi_{0is} = 1 - \sum_{k=1}^K \Psi_{kis}$. The term $\boldsymbol{\Sigma}_p^{-1}(\boldsymbol{\Psi})$ in the fourth equality denotes the inverse of the covariance matrix of the choice probability vector $\boldsymbol{\Psi}$ evaluated at $(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \boldsymbol{\theta})$. To see the fourth equality, observe that the inverse of the covariance matrix $\boldsymbol{\Sigma}_{pis}$ of a multinomial distributed random variable with probabilities $(\Psi_{0is}, \Psi_{1is}, \dots, \Psi_{Kis})$ equals

$$\boldsymbol{\Sigma}_{pis}^{-1} = \text{diag}(1/\Psi_{1is}, \dots, 1/\Psi_{Kis}) + 1/\Psi_{0is} \cdot \mathbf{e}\mathbf{e}'$$

where \mathbf{e} is a $K \times 1$ dimensional vector given by $\mathbf{e} = (1, 1, \dots, 1)'$. The expression for the inverse is given for example in Tanabe and Sagae (1992). The inverse of the covariance matrix $\boldsymbol{\Sigma}_p$ is then given by the block-diagonal matrix $\boldsymbol{\Sigma}_p^{-1} = \text{diag}(\boldsymbol{\Sigma}_{p11}^{-1}, \dots, \boldsymbol{\Sigma}_{pNm_s}^{-1})$. This follows from our assumption that the choices are independently distributed across states and players. ■

References

- [1] Aguirregabiria, V. (1999), The Dynamics of Markups and Inventories in Retailing Firms, *Review of Economic Studies*, 66, 275-308.
- [2] Aguirregabiria, V., and Mira P. (2002), Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models, *Econometrica*, 70, 1519-1544.
- [3] Aguirregabiria, V., and Mira P. (2004), Sequential Estimation of Dynamic Discrete Games, technical report, Boston University, forthcoming *Econometrica*.
- [4] Bajari, P. (1998), Econometrics of Sealed-Bid Auctions, 1998 Proceedings of the Business and Economic Statistics Section of the American Statistical Association, 41-49.
- [5] Bajari, P., Benkard, L. and Levin (2004), Estimating Dynamic Models of Imperfect Competition, technical report, Stanford University.
- [6] Billingsley, P. (1968), *Convergence of Probability Measures*, John Wiley, New York.
- [7] Bresnahan, T.F., and P. Reiss (1990), Entry in Monopoly Markets, *Review of Economic Studies*, 57, 531-553.
- [8] Brock, W., and Durlauf S. (2001), Discrete Choice with Social Interactions, *Review of Economic Studies*, 68, 235-260.
- [9] Carrasco, M., and Florens, J.-P. (2000), Generalization of GMM to a Continuum of Moment Conditions, *Econometric Theory*, 16, 797-834.
- [10] Chamberlain, G. (1987), Asymptotic Efficiency in Estimation with Conditional Moment Restrictions, *Journal of Econometrics*, 34, 305-334.

- [11] Collard-Wexler, A. (2005), Plant Turnover and Demand Fluctuations in the Ready-Mix Concrete Industry, technical report, Northwestern University.
- [12] Doraszelski. U., and M. Satterthwaite (2003), Foundations of Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity, technical report, Northwestern University and Hoover Institution.
- [13] Doraszelski. U. and K.L. Judd (2004), Breaking the Curse of Dimensionality with Continuous-Time, Finite-State Stochastic Games, technical report, Hoover Institution.
- [14] Einav, L. (2003), Not All Rivals Look Alike: Estimating an Equilibrium Model of The Release Date Timing Game, technical report, Stanford.
- [15] Elyakime, B., J.J. Laffont, P. Loisel, and Q. Vuong (1994), First-Price Sealed-Bid auctions with Secret Reservation Prices, *Annales d'Economie et de Statistique*, 34, 115-141.
- [16] Gouriéroux, C. and A. Monfort (1995), *Statistics and Econometric Models: Volume I*, Cambridge University Press.
- [17] Gouriéroux, C., Monfort A., and A. Trognon (1985), Moindres Carres Asymptotiques, *Annales de l'INSEE*, 58, 91-121.
- [18] Grund, B. (1993), Kernel Estimators for Cell Probabilities, *Journal of Multivariate Analysis*, 46, 283-308.
- [19] Guerre, E., I. Perrigne, and Q. Vuong (2000), Optimal Nonparametric Estimation of First-Price Auctions, *Econometrica*, 68, 525-574.
- [20] Hansen, L. (1982), Large Sample Properties of Methods of Moments Estimators, *Econometrica*, 50, 1029-1054.

- [21] Heckman J., and S. Navarro (2005), Dynamic Discrete Choice and Dynamic Treatment Effects, technical working paper No 316, National Bureau of Economic Research.
- [22] Hotz, J., and R.A. Miller (1993), Conditional Choice Probabilities and the Estimation of Dynamic Models, *Review of Economic Studies*, 60, 497-529.
- [23] Hotz, J., Miller, R.A., Sanders S., and J. Smith (1994), A Simulation Estimator for Dynamic Models of Discrete Choice, *Review of Economic Studies*, 61, 265-289.
- [24] Jofre-Bonet, M., and M. Pesendorfer (2003), Estimation of a Dynamic Auction Game, *Econometrica*, 71, 1443-1489.
- [25] Magnac T., and D. Thesmar (2002), Identifying Dynamic Discrete Decision Processes, *Econometrica*, 70, 801-816.
- [26] Maskin, E., and J. Tirole (1994), Markov Perfect Equilibrium, technical report, Harvard University.
- [27] Maskin, E., and J. Tirole (2001), Markov Perfect Equilibrium I: Observable Actions, *Journal of Economic Theory*, 100, 191-219.
- [28] McFadden, D. (1978), Modelling the Choice of Residential Location, in A. Karlquist et al., eds., *Spatial Interaction Theory and Planning Models*, Amsterdam: North Holland, 1978.
- [29] McKelvey, R.D., and T.R. Palfrey (1995), Quantal Response Equilibria in Normal Form Games, *Games and Economic Behavior*, 10, 6-38.
- [30] Pakes, A., Ostrovsky, T., and S. Berry (2004), Two Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples), technical report, Harvard University.

- [31] Pesendorfer, M., and Schmidt-Dengler, P. (2003), Identification and Estimation of Dynamic Games, NBER Working Paper No. w9726.
- [32] Rust. J. (1987), Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher, *Econometrica*, 57, 1027-1058.
- [33] Rust. J. (1994), Estimation of Dynamic Structural Models, Problems and Prospects, Part I: Discrete Decision Processes, in C. Sims and J. Laffont (eds.), *Advances in Econometrics: 6th World Congress, Volume II*, Cambridge University Press.
- [34] Ryan, S. (2004), The Costs of Environmental Regulation in a Concentrated Industry, technical report, Duke University.
- [35] Schmidt-Dengler, P. (200), The Timing of New Technology Adoption: The Case of MRI, technical report, London School of Economics.
- [36] Seim, K. (2005), An Empirical Model of Firm Entry with Endogenous Product-Type Choices, technical report, Stanford University, forthcoming, *RAND Journal of Economics*.
- [37] Slade, M. (1998), Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices, *Review of Economic Studies*, 65, 87-107.
- [38] Tamer, E. (2003), Incomplete Simultaneous Response Models with Multiple Equilibria, *Review of Economic Studies*, 70, 147-165.
- [39] Tanabe, K., and M. Sagae (1992), An Exact Cholesky Decomposition and the Generalized Inverse of the Variance-Covariance Matrix of the Multinomial Distribution, with Applications, *Journal of the Royal Statistical Society. Series B (Methodological)*, 54, 211-219.

Table 1: Monte Carlo Results, Equilibrium (i).

T	Estimator	c		π^1		π^2		MSE
100	LS-E	-0.116	(0.768)	1.424	(1.472)	-1.367	(1.108)	406.617
	LS-I	-0.292	(0.400)	1.086	(0.493)	-1.052	(0.511)	70.732
	PML	-0.260	(0.310)	1.083	(0.344)	-1.067	(0.372)	38.738
	k-PML	-0.231	(0.163)	1.225	(0.309)	-1.186	(0.229)	17.655
1,000	LS-E	-0.195	(0.030)	1.207	(0.077)	-1.208	(0.067)	1.156
	LS-I	-0.213	(0.104)	1.185	(0.093)	-1.189	(0.134)	3.791
	PML	-0.209	(0.111)	1.187	(0.107)	-1.191	(0.129)	4.061
	k-PML	-0.204	(0.040)	1.197	(0.090)	-1.202	(0.061)	1.334
10,000	LS-E	-0.200	(0.008)	1.199	(0.022)	-1.200	(0.018)	0.089
	LS-I	-0.203	(0.032)	1.197	(0.028)	-1.196	(0.042)	0.363
	PML	-0.202	(0.034)	1.196	(0.034)	-1.197	(0.040)	0.396
	k-PML	-0.201	(0.012)	1.199	(0.029)	-1.200	(0.018)	0.132
100,000	LS-E	-0.200	(0.003)	1.200	(0.007)	-1.200	(0.005)	0.009
	LS-I	-0.201	(0.010)	1.200	(0.009)	-1.200	(0.013)	0.036
	PML	-0.201	(0.011)	1.200	(0.011)	-1.199	(0.013)	0.042
	k-PML	-0.200	(0.004)	1.200	(0.009)	-1.200	(0.006)	0.013

Table 2: Monte Carlo Results, Equilibrium (ii).

T	Estimator	c		π^1		π^2		MSE
100	LS-E	-0.391	(0.559)	1.047	(0.415)	-0.967	(0.577)	93.139
	LS-I	-0.235	(0.509)	1.094	(0.480)	-1.025	(0.678)	99.243
	PML	-0.322	(0.448)	0.954	(0.399)	-0.907	(0.573)	84.802
	k-PML	-0.499	(0.193)	1.031	(0.245)	-0.734	(0.186)	46.639
1,000	LS-E	-0.219	(0.153)	1.181	(0.126)	-1.169	(0.192)	7.759
	LS-I	-0.208	(0.173)	1.191	(0.140)	-1.176	(0.207)	9.319
	PML	-0.222	(0.165)	1.165	(0.139)	-1.151	(0.207)	9.377
	k-PML	-0.491	(0.060)	0.998	(0.065)	-0.749	(0.052)	33.937
10,000	LS-E	-0.198	(0.021)	1.199	(0.025)	-1.202	(0.030)	0.192
	LS-I	-0.202	(0.056)	1.198	(0.046)	-1.197	(0.066)	0.959
	PML	-0.205	(0.054)	1.193	(0.045)	-1.192	(0.067)	0.953
	k-PML	-0.493	(0.029)	0.991	(0.027)	-0.750	(0.036)	33.535
100,000	LS-E	-0.200	(0.002)	1.200	(0.007)	-1.200	(0.005)	0.008
	LS-I	-0.201	(0.017)	1.200	(0.014)	-1.199	(0.020)	0.088
	PML	-0.201	(0.016)	1.199	(0.014)	-1.199	(0.020)	0.087
	k-PML	-0.491	(0.032)	0.991	(0.025)	-0.752	(0.047)	33.245

Table 3: Monte Carlo Results, Equilibrium (iii).

T	Estimator	c	π^1	π^2	MSE
100	LS-E	-0.439 (0.794)	1.036 (0.596)	-0.928 (0.963)	206.942
	LS-I	-0.240 (0.491)	1.101 (0.476)	-1.008 (0.627)	90.878
	PML	-0.326 (0.433)	0.964 (0.393)	-0.896 (0.541)	79.729
	k-PML	-0.504 (0.190)	1.037 (0.229)	-0.722 (0.169)	46.474
1,000	LS-E	-0.219 (0.179)	1.181 (0.145)	-1.173 (0.212)	9.925
	LS-I	-0.209 (0.180)	1.189 (0.147)	-1.179 (0.211)	9.934
	PML	-0.227 (0.170)	1.156 (0.141)	-1.147 (0.207)	9.733
	k-PML	-0.511 (0.062)	0.992 (0.066)	-0.734 (0.052)	36.835
10,000	LS-E	-0.201 (0.053)	1.198 (0.042)	-1.200 (0.063)	0.858
	LS-I	-0.200 (0.057)	1.199 (0.046)	-1.201 (0.065)	0.952
	PML	-0.202 (0.053)	1.195 (0.043)	-1.197 (0.064)	0.883
	k-PML	-0.510 (0.025)	0.989 (0.022)	-0.733 (0.030)	36.082
100,000	LS-E	-0.201 (0.017)	1.199 (0.013)	-1.198 (0.020)	0.085
	LS-I	-0.200 (0.018)	1.200 (0.014)	-1.199 (0.021)	0.096
	PML	-0.201 (0.017)	1.199 (0.014)	-1.199 (0.021)	0.091
	k-PML	-0.508 (0.031)	0.990 (0.020)	-0.736 (0.045)	35.700