

Asymptotic Least Squares Estimators for Dynamic Games¹

MARTIN PESENDORFER and PHILIPP SCHMIDT-DENGLER
London School of Economics and Political Sciences

First version received September 2004; final version accepted September 2007 (Eds.)

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 those by Hotz and Miller (1993) and by 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 illustrates the small sample performance and computational feasibility of alternative estimators.

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 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 threefold: first, we suggest a unifying framework for the estimation of dynamic games with finite actions based on an equilibrium equation system. Static models with strategic interaction as well as single-agent static and dynamic discrete choice models arise as special cases within the proposed framework. 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), the pseudo-maximum likelihood estimator by Aguirregabiria and Mira (2002), and estimators for dynamic games recently considered in Aguirregabiria and Mira (2007), Pakes, Ostrovsky and Berry (2007), and in Bajari, Benkard and Levin (2007). We illustrate the associated weight matrices for these estimators.

1. This paper supersedes our earlier NBER Working Paper entitled "Identification and Estimation of Dynamic Games" from May 2003.

The second contribution of this paper is to provide sufficient conditions for the identification of the primitives of dynamic games.

Third, we 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),² and the entry literature in industrial organization, see Bresnahan and Reiss (1990) and Seim (2006). 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 conditions of dynamic games efficiently.

Recent working papers by Aguirregabiria and Mira (2007), Pakes *et al.* (2007) and Bajari *et al.* (2007)³ 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 pay-off 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

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

3. Bajari *et al.* (2007) also show that the moment estimator can include both discrete and continuous choices. Ryan (2006) and Collard-Wexler (2006) provide applications to the cement and ready-mix concrete industries.

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 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 pay-off 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 pay-offs.

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 analyses 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 underidentification 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. We begin by describing the features of the game, then we illustrate the equilibrium conditions, and finally, we illustrate our estimation approach.

Setup. Consider a static version duopoly game in which two firms have to decide whether to be active or not in a given period. Each firm i has two possible choices: be active or not active, $a_i \in \{\text{active}, \text{not active}\}$. The choices are made simultaneously. We augment this classic game with two elements: a publicly observed demand variable and a private profitability variable. The demand variable describes whether demand is low, medium, or high, $s \in \{1, 2, 3\}$. We shall assume that demand affects the profitability of being active. Firm i 's pay-off depends on the firms' choices and demand in the following way:

$$\pi_i(a_i, a_j, s) = \begin{cases} \theta_i^1 \cdot s & \text{if } a_i = \text{active}, a_j = \text{not active}; \\ \theta_i^2 \cdot s & \text{if } a_i = \text{active}, a_j = \text{active}; \\ 0 & \text{otherwise.} \end{cases}$$

The parameters $(\theta_1^1, \theta_1^2, \theta_2^1, \theta_2^2)$ describe the profit of an active firm per unit of demand with monopoly profit exceeding duopoly profit, $\theta_i^1 > \theta_i^2$ for $i = 1, 2$.

The second element that we incorporate into the game is a pay-off shock ε_i . This pay-off shock variable is firm specific and drawn from the standard normal distribution function denoted by Φ . The pay-off shock ε_i measures firm i 's preference for being active vs. not being active on a given day. We assume that the pay-off shock is i.i.d and enters additively yielding a total pay-off of

$$\begin{cases} \pi_i(\text{not active}, a_j, s) & \text{if } a_i = \text{not active} \\ \pi_i(\text{active}, a_j, s) + \varepsilon_i & \text{if } a_i = \text{active.} \end{cases}$$

We may view the pay-off shock ε_i as any element of period profits that is not known to the other firm but is known to firm i , such as privately observed demand for firm i 's product or the (negative) of the cost of setting up production. We will see in Section 3 that the shock distribution may additionally depend on the state variable and can represent an incumbent's scrap value or a potential entrant's entry cost.

A pure strategy for firm i in the augmented duopoly game is denoted by $a_i(s, \varepsilon_i)$, which is a function of the public state s and the private profitability 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, one equilibrium has the feature that firm i is active more often, while the other equilibrium involves firm j being active more frequently. 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 firm i that firm $j \neq i$ will be active in state s for $i = 1, 2$. The number $\sigma_i(s)$ equals the *ex ante* expected probability that firm j will be active, while the number $1 - \sigma_i(s)$ equals the *ex ante* expected probability that firm j will not be active. In any equilibrium, it must be that being active is a best response for some type ε_i , taking as given the beliefs about the opponent's behaviour:

$$[1 - \sigma_i(s)] \cdot \theta_i^1 \cdot s + \sigma_i(s) \cdot \theta_i^2 \cdot s + \varepsilon_i \geq 0, \tag{1}$$

as the pay-off shock ε_i has support equal to the real line. Evaluating the best response condition *ex ante*, before the pay-off shock ε_i is observed, we can describe the probability of being active in demand state s for firm i as

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

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

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

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:

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

The equilibrium equation (3) provides 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 pay-offs for the problem at hand. Our example has two firms, two actions, and three states, yielding a total of six best response equations in (3).

The privately observed pay-off shock plays an important role in the derivation of the equilibrium equations. Given a firm’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 0. 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 firms and a time series for the demand 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 that for a given vector of pay-off relevant state variables (s, ε_i) , firms 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 profitability variable ε_i is i.i.d 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 firm produces 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, $\widehat{p}_i(s)$, 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 pay-off parameters $\boldsymbol{\theta}$ as the only unknown elements:

$$\widehat{\mathbf{p}} - \Psi(\widehat{\mathbf{p}}, \boldsymbol{\theta}) = 0. \tag{4}$$

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

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

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 pay-offs, 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 pay-off 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 pay-off 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 pay-off 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 pay-off 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 Since the distribution of the shock depends on the state of a player, it may represent the entry cost for an inactive firm and the scrap value (the opportunity cost of staying in the market) for an incumbent as in Pakes *et al.* (2007). Independence of pay-off shock realizations over time is assumed to keep the number of pay-off relevant state variables small. With correlated shocks, a model of learning would emerge in which players would infer other players' private state based on past actions, increasing the dimensionality of the state space considerably.

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 pay-off shock. An action profile \mathbf{a}^t denotes the vector of joint actions in period t , $\mathbf{a}^t = (a_1^t, \dots, a_N^t) \in \mathbf{A} = \times_{j=1}^N A_j$. Sometimes we use the notation $\mathbf{a}_{-i}^t = (a_1^t, \dots, a_{i-1}^t, a_{i+1}^t, \dots, a_N^t) \in \mathbf{A}_{-i}$ to denote the actions by players other than player i . The cardinality of the action space \mathbf{A} equals $m_a = (K + 1)^N$.

State transition. The state transition is described by a probability density function $g : \mathbf{A} \times \mathbf{S} \times \mathbf{S} \rightarrow [0, 1]$ where a typical element $g(\mathbf{a}^t, \mathbf{s}^t, \mathbf{s}^{t+1})$ equals the probability that state

4. An 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 behaviour, see Rust (1994).

s^{t+1} is reached when the current action profile and state are given by (\mathbf{a}^t, s^t) . We require $\sum_{s' \in \mathbf{S}} g(\mathbf{a}, s, s') = 1$ for all $(\mathbf{a}, s) \in \mathbf{A} \times \mathbf{S}$. We frequently use the symbol \mathbf{G} to denote the $(m_a \cdot m_s) \times m_s$ -dimensional transition matrix in which column $s' \in \mathbf{S}$ consists of the vector of probabilities $[g(\mathbf{a}, s, s')_{\mathbf{a} \in \mathbf{A}, s \in \mathbf{S}}]$.

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

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

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

Game pay-off. Players discount future pay-offs with discount factor $\beta_i < 1$. The game pay-off of player i equals the sum of discounted period pay-offs. 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 pay-off 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 $(\mathbf{a}, \boldsymbol{\sigma}) = (a_1, \dots, a_N, \sigma_1, \dots, \sigma_N)$ 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 $s \in \mathbf{S}$;
- (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 pay-offs. It is defined as the discounted sum of future pay-offs for player i before player-specific shocks are observed and actions are taken. Let $\sigma_i(\mathbf{a} | s)$ 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 s . The *ex ante* value function can be written as

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

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

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

Equation (6) provides an expression for the *ex ante* value function associated with the beliefs σ_i . The terms on the R.H.S. are the discount factor, the choice probability matrix, the state transition matrix, the period return function, and the expected pay-off 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 pay-offs; $\boldsymbol{\sigma}_i$ 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 s , column (\mathbf{a}, s) , and zeros in row s , column (\mathbf{a}, s') with $s' \neq s$; $\mathbf{D}_i(\sigma_i) = [D_i(\mathbf{s}, \sigma_i)]_{\mathbf{s} \in \mathbf{S}}$ is the $m_s \times 1$ -dimensional vector of expected pay-off shocks with element $D_i(\mathbf{s}, \sigma_i) = \sum_{k=1}^K E_\varepsilon [\varepsilon_i^k | a_i = k; \sigma_i | \sigma_i(a_i = k | \mathbf{s})]$; 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 limitations of our model assumptions. Section 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. Section 4.2 concludes with remarks on limitations and possible extensions of our model.

4.1. Equilibrium characterization

First, we examine when it is optimal to choose action a_i in state $(s, \boldsymbol{\varepsilon}_i)$. 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 pay-off shocks under action a_i with beliefs σ_i .

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

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

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

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 pay-off shock $\boldsymbol{\varepsilon}_i$ is observed. Doing so gives an expression for the *ex ante* optimal choice probability of player i given the beliefs σ_i :

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

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

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} :

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

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 L.H.S. and R.H.S. 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 pay-offs.

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 perfect 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 (2007) 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 (2006).

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.

Proposition 1 and Theorem 1 permit asymmetric pay-offs and strategies. Our framework takes heterogeneous behaviour by agents into account. It permits the possibility that players i and j adopt distinct strategies, $a_i(s_i, s_{-i}, \boldsymbol{\varepsilon}_i) \neq a_j(s'_j, s'_{-j}, \boldsymbol{\varepsilon}_j)$, even when the state variables are identical, $(s_i, s_{-i}) = (s'_j, s'_{-j})$ and $\boldsymbol{\varepsilon}_i = \boldsymbol{\varepsilon}_j$. Next, we shall illustrate that imposing symmetry on pay-offs, transition probabilities, and strategies does not affect the properties of the framework conceptually.

Symmetry restriction. Symmetry means that players i and j adopt identical strategies, receive identical pay-offs, and face identical transition probabilities when their state variables are identical. The game is symmetric if $\boldsymbol{\pi}_i(a^1, s^1) = \boldsymbol{\pi}_j(a^2, s^2)$ for all $i, j \in \mathbf{N}$, for all $(a^1, s^1), (a^2, s^2) \in \mathbf{A} \times \mathbf{S}$ such that $(a^1_i, a^1_{-i}) = (a^2_j, a^2_{-j})$, $(s^1_i, s^1_{-i}) = (s^2_j, s^2_{-j})$ and $g(a^1, s^1, s^{1'}) = g(a^2, s^2, s^{2'})$ for all $(a^1, s^1, s^{1'}), (a^2, s^2, s^{2'}) \in \mathbf{A} \times \mathbf{S} \times \mathbf{S}$ such that $(a^1_i, a^1_{-i}) = (a^2_j, a^2_{-j})$, $(s^1_i, s^1_{-i}) = (s^2_j, s^2_{-j})$ and $(s^{1'}_i, s^{1'}_{-i}) = (s^{2'}_j, s^{2'}_{-j})$. The strategies are symmetric if $a_i(s^1, \boldsymbol{\varepsilon}_i) = a_j(s^2, \boldsymbol{\varepsilon}_j)$ for all $i, j \in \mathbf{N}$, for all $s^1, s^2 \in \mathbf{S}$ such that $(s^1_i, s^1_{-i}) = (s^2_j, s^2_{-j})$, and for all $\boldsymbol{\varepsilon}_i, \boldsymbol{\varepsilon}_j \in \mathbb{R}^K$ such that $\boldsymbol{\varepsilon}_i = \boldsymbol{\varepsilon}_j$. A Markov perfect equilibrium is symmetric if (i) the game is symmetric and (ii) the strategies are symmetric.

Corollary 1. *Suppose the game is symmetric.*

- (i) *A symmetric Markov perfect 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 pay-off vector $\boldsymbol{\Pi}$ and transition probability matrix \mathbf{G} , which can be useful in the empirical analysis.

Static models and single-agent models. Static models, as in McKelvey and Palfrey (1995), Seim (2006), and others, are a special case of our framework with $\beta = 0$. When the number of players $N = 1$, the framework corresponds to a single-agent discrete decision process as studied in Rust (1987, 1994).

4.2. Discussion

Next, we discuss some limitations and possible extensions of our framework. We illustrate that serially correlated pay-off shocks are not readily incorporated. We discuss possible extensions to continuous action spaces and describe computational limitations.

Pay-off shock. The assumption regarding conditional independence of the pay-off shock $\varepsilon_i^{a_i}$ is an important assumption commonly made in this literature. It permits the use of the Markovian framework. For a detailed discussion of the independence assumption in Markovian decision problems, see Rust (1994).

Permitting serial correlation in the privately observed shock 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. Doing so, may render the method computationally infeasible. Whether the assumption of conditional independence is appropriate depends on the environment studied. In some applications, a firm's scrap value, sunk entry cost, or productivity may be autocorrelated, and the proposed method may not be applicable.

Introducing a common pay-off shock component that is 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. In order to correctly infer the beliefs of players, the econometrician has to be able to integrate conditional choice probabilities over the distribution of common shocks.

Heterogeneity in the pay-off shock distribution can be incorporated by adding a player-specific dummy variable to the publicly observed state variable.

Action space. Recall that the equilibrium was characterized by one equation for every action in every state. A game with a continuous action space may yield a continuum of equilibrium conditions for every state. Such equilibrium conditions have been studied in special cases. For example, in the pricing game studied in Jofre-Bonet and Pesendorfer (2003), the equilibrium condition corresponding to (11) becomes a set of first order conditions.

Computational limitations. Our modelling framework and approach also has computational 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 \mathbf{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 the 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 described some limitations and possible extensions of our model framework. 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 $(\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_N, F, \beta, g)$ 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 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 pay-off relevant state variables $(s, \mathbf{\varepsilon}_i)$.⁶ Consequently, the Markovian assumption guarantees that a single time series has been generated by only one equilibrium. When considering cross-sectional data, one has to assume additionally that the same equilibrium is played in every path observed.

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 $g(\mathbf{a}, s, s')$ for any $s, s' \in \mathbf{S}, \mathbf{a} \in \mathbf{A}$.

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 $(\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_N)$. Our approach follows the analysis of the single-agent case by Magnac and Thesmar (2002) in that we fix (F, β) and analyse the identification of pay-off parameters $(\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_N)$. The following proposition illustrates the need for identifying restrictions.

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

6. Thus, the Markovian assumption rules out "equilibrium switches" in a single path of play.

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 $(\Pi_1, \dots, \Pi_N, F, \beta)$ are non-parametrically identified. Even if the distribution function F and the discount factor β are known, not all parameters in (Π_1, \dots, Π_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 (Π_1, \dots, Π_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).⁷ Proposition 2 also illustrates that the degree of underidentification 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 .

In many economic settings, the pay-off 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 pay-off parameters and ensure that the remaining pay-off 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

$$\mathbf{R}_i \cdot \Pi_i = \mathbf{r}_i. \tag{12}$$

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 pay-offs, which allows us to illustrate the identification problem for pay-off 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 pay-offs from action a_i and action 0. The indifference equation between action pairs is given by:

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

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 pay-off parameters. The linear equation system for player i is given by

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

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

7. Heckman and Navarro (2007) 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.

response system (11). Next, we augment the linear equation system (14) with the pay-off restrictions $\mathbf{R}_i \cdot \boldsymbol{\Pi}_i = \mathbf{r}_i$, yielding the $m_a \cdot m_s$ equations in $m_a \cdot m_s$ unknown pay-off parameters:

$$\begin{bmatrix} \mathbf{X}_i \\ \mathbf{R}_i \end{bmatrix} \boldsymbol{\Pi}_i + \begin{bmatrix} \mathbf{Y}_i \\ \mathbf{r}_i \end{bmatrix} = \bar{\mathbf{X}}_i \boldsymbol{\Pi}_i + \bar{\mathbf{Y}}_i \tag{15}$$

$$= 0.$$

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 rank $(\bar{\mathbf{X}}_i) = (m_a \cdot m_s)$, then $\boldsymbol{\Pi}_i$ is exactly identified.*

The proposition allows us to verify whether a set of linear pay-off 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 pay-off parameters are identified. If additional restrictions are imposed, then the pay-off 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 pay-off 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:*

$$\pi_i(a_i, \mathbf{a}_{-i}, s_i, \mathbf{s}_{-i}) = \pi_i(a_i, \mathbf{a}_{-i}, s_i, \mathbf{s}'_{-i}) \quad \forall \mathbf{a} \in \mathbf{A}, (s_i, \mathbf{s}_{-i}), (s_i, \mathbf{s}'_{-i}) \in \mathbf{S} \tag{16}$$

$$\pi_i(0, \mathbf{a}_{-i}, s_i) = r_i(\mathbf{a}_{-i}, s_i) \quad \forall \mathbf{a}_{-i} \in \mathbf{A}_{-i}, s_i \in \mathbf{S}_i \tag{17}$$

where \mathbf{r}_i is a $((K + 1)^{N-1} \cdot L) \times 1$ -dimensional exogenous vector. *Restriction (16)* says that period pay-offs do not depend on the state variables of other firms. It fixes $m_a \cdot (m_s - L)$ pay-off parameters. *Restriction (16)* is satisfied in games with adjustment costs such as entry or investment games. *Restriction (17)* says that period pay-offs under action $a_i = 0$ are fixed exogenously for every state s_i and action profile \mathbf{a}_{-i} . *Restriction (17)* fixes $(K + 1)^{N-1} \cdot L$ pay-off parameters. *Restriction (17)* is satisfied in games in which one action, say action inactivity, implies zero pay-offs. For given F and β , imposing restrictions (16) and (17) in equation (15) ensures identification of the pay-off 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 pay-off normalization such as restriction (17) is required for identification as the indifference analysis determines pay-offs relative to the pay-off 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 $\boldsymbol{\theta} = (\boldsymbol{\theta}^\pi, \boldsymbol{\theta}^F, \beta, \boldsymbol{\theta}^g) \in \Theta \subset \mathbb{R}^q$ denote an identified parameter vector.⁸ The parameter vector $\boldsymbol{\theta}$ may include parameters entering the period pay-offs $\pi_i(\mathbf{a}, s; \boldsymbol{\theta}^\pi)$ for all i , the transition probabilities $g(\mathbf{a}, s, s'; \boldsymbol{\theta}^g)$, the distribution of pay-off shocks $F(\boldsymbol{\varepsilon}; \boldsymbol{\theta}^F)$, 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).

$$\mathbf{h}(\mathbf{p}, \mathbf{g}, \boldsymbol{\theta}) = \mathbf{p} - \Psi(\mathbf{p}, \mathbf{g}, \boldsymbol{\theta}) = 0. \tag{18}$$

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 (18). As explained in Section 5, the equation system (18) is not the only representation of the equilibrium conditions. An equivalent representation exists in the space of expected discounted pay-offs 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 (18).

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

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

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

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.⁹ A Monte Carlo study comparing these techniques in single-agent dynamic models is conducted in Hotz *et al.* (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 $i \in \mathbf{N}, k \in \mathbf{A}_i, s \in \mathbf{S}$ and observed state frequency by $n_{ass'} = \sum_t \mathbf{1}(a^t = \mathbf{a}, s^t = s, s^{t+1} = s')$. By assumption, choices and states are multinomially distributed, which implies that the maximum likelihood estimator equals the sample frequency:

$$\widehat{p}(k, i, s) = \frac{n_{kis}}{\sum_{l \in \mathbf{A}_i} n_{lis}}, \quad \widehat{g}(\mathbf{a}, s, s') = \frac{n_{ass'}}{\sum_{s'' \in \mathbf{S}} n_{ass''}}. \tag{19}$$

8. We 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 .

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

We denote with $\widehat{\mathbf{p}} = (\widehat{p}(k, i, s))_{i \in \mathbf{N}, k \in \mathbf{A}_i, s \in \mathbf{S}}$ the vector of sample frequencies. Billingsley (1968) 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

$$\mathbf{h}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta}) = \widehat{\mathbf{p}}_T - \boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta}) = 0 \tag{20}$$

to be satisfied approximately. Consider initially the case in which the pay-off 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 overidentified parameter vector, there are more equations than parameters, and there is no unique way to select a solution. A solution depends on the weight assigned to individual equations.

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 $\widetilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$ to the problem

$$\min_{\boldsymbol{\theta} \in \Theta} [\widehat{\mathbf{p}}_T - \boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})]' \mathbf{W}_T [\widehat{\mathbf{p}}_T - \boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})]. \tag{21}$$

Thus, the asymptotic least squares estimator $\widetilde{\boldsymbol{\theta}}_T(\mathbf{W}_T)$ brings the constraint $\widehat{\mathbf{p}}_T - \boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})$ 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 $T \rightarrow \infty$, $\mathbf{W}_T \rightarrow \mathbf{W}_0$ a.s. where \mathbf{W}_0 is a non-stochastic positive definite matrix.
- A4: $\boldsymbol{\theta}$ satisfies $[\mathbf{p}(\boldsymbol{\theta}_0) - \boldsymbol{\Psi}(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta})]' \mathbf{W}_0 [\mathbf{p}(\boldsymbol{\theta}_0) - \boldsymbol{\Psi}(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta})] = 0$ implies that $\boldsymbol{\theta} = \boldsymbol{\theta}_0$.
- A5: the functions π, g, F are twice continuously differentiable in θ .
- A6: the matrix $[\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta}_0)]' \mathbf{W}_0 [\nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi}(\mathbf{p}(\boldsymbol{\theta}_0), \mathbf{g}(\boldsymbol{\theta}_0), \boldsymbol{\theta}_0)]$ 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 Gourieroux, Monfort and Trognon (1985) for asymptotic least squares estimators (see also Gourieroux and Monfort, 1995, theorem 9.1).

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

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

as $T \rightarrow \infty$, where

$$\begin{aligned} \boldsymbol{\Omega}(\boldsymbol{\theta}_0) &= (\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \mathbf{W}_0 \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi})^{-1} \nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \mathbf{W}_0 [(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi}] \boldsymbol{\Sigma} \\ &\quad [(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \boldsymbol{\Psi}]' \mathbf{W}_0 \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi} (\nabla_{\boldsymbol{\theta}} \boldsymbol{\Psi}' \mathbf{W}_0 \nabla_{\boldsymbol{\theta}'} \boldsymbol{\Psi})^{-1} \end{aligned} \tag{22}$$

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 .¹⁰

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*

$$[(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi] \Sigma [(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi]' \quad \text{and}$$

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

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 et al. (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

$$\Omega(\mathbf{W}_0^*) = \left(\nabla_{\theta} \Psi' \left([(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi] \Sigma [(\mathbf{I} : \mathbf{0}) - \nabla_{(p,g)'} \Psi]' \right)^{-1} \nabla_{\theta'} \Psi \right)^{-1}. \quad (23)$$

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 $\alpha_k \in \{0, 1\}$ for all k . Let \mathbf{T}_{is} denote the set of observations for player i in state s . The orthogonality condition is defined as

$$E[Z \otimes [\alpha_{is} - \Psi_{is}(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \theta)]] = 0, \quad (24)$$

10. If two distinct auxiliary estimators, $\hat{\mathbf{p}}_T^1$ and $\hat{\mathbf{p}}_T^2$, are employed in the equilibrium conditions (20), $\hat{\mathbf{p}}_T^1 - \Psi(\hat{\mathbf{p}}_T^2, \hat{\mathbf{g}}_T, \theta) = 0$, then the matrix Σ in Proposition 4 refers to the asymptotic covariance matrix associated with the vector of auxiliary parameters $(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T) = (\hat{\mathbf{p}}_T^1, \hat{\mathbf{p}}_T^2, \hat{\mathbf{g}}_T)$.

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.¹¹ It is easy to see that the moment estimator based on the moment condition (24) 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 analogue of (24) can be rewritten as a weighted average of the estimating equation:

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

as $\sum_{t \in \mathbf{T}_{is}} \alpha^t$ equals the vector of sample frequencies $\hat{\mathbf{p}}_{is} = (n_{1is}/n_{is}, \dots, n_{Kis}/n_{is})$ defined in equation (19). So, the estimating equation of the equivalent asymptotic least squares estimator is $\hat{\mathbf{p}} - \Psi(\hat{\mathbf{p}}_T, \hat{\mathbf{g}}_T, \theta) = 0$.¹² The weights assigned to the individual estimating equations depend on the vector of instruments Z and the choice of a $(J \cdot K)$ square weight matrix associated with the moment conditions.¹³

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 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 0 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 (20). 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 and weight matrix for a given set of moment conditions. They correspond to the optimal asymptotic least squares weight matrix \mathbf{W}_0^* if a moment condition is considered for every agent and at every state. Notice though that this enlarged set of moment conditions is distinct from those considered in Hotz and Miller (1993).

In the case of exact identification, the asymptotic least squares estimator is optimal regardless of the choice of weight matrix. Jofre-Bonet and Pesendorfer (2003) consider a continuous-choice dynamic pricing model in which the equilibrium equation system is linear in the cost parameters of interest. They impose restrictions for exact identification. Their estimator can be recast in our model as an asymptotic least square estimator with weight matrix equal to the identity matrix.

Our approach could be extended to general continuous action games by replacing equilibrium equation system (20) by a continuum of conditions. The corresponding asymptotic least squares estimator minimizes the norm of these conditions in a Hilbert space. Carrasco and Florens (2000) establish consistency and asymptotic normality. Efficient estimation of an overidentified model, however, would require the inversion of a covariance operator in an infinite dimensional Hilbert space.

11. Hotz *et al.* (1994) consider an alternative representation of the equilibrium in the pay-off space. As explained earlier, this representation is equivalent to our formulation in the choice probability space.

12. Observe 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.

13. The 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 $\tilde{\mathbf{p}}$.

Recent working papers by Bajari *et al.* (2007)¹⁴ and Pakes *et al.* (2007)¹⁵ 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 overidentified 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 $\widehat{\mathbf{g}}_T$. It is given by

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

where Ψ_{kis} is evaluated at $(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T)$. 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 (21) 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

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

as is shown in the appendix. Observe that for auxiliary estimates not equalling the sample frequency, $\widehat{\mathbf{p}}_T \neq \widehat{\mathbf{p}}$, the PML uses two distinct auxiliary parameter estimators of the choice probabilities, $\widehat{\mathbf{p}}$ and $\widehat{\mathbf{p}}_T$, and the equivalent asymptotic least squares estimator solves the equation system $\widehat{\mathbf{p}} - \boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta}) = 0$. The first appearance of \mathbf{p} in the equation system is evaluated at $\widehat{\mathbf{p}}$ and the second at $\widehat{\mathbf{p}}_T$. The intuition is that the PML sets the sample frequency $\widehat{\mathbf{p}}$ equal to the pseudo-likelihood $\boldsymbol{\Psi}$ evaluated at $(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})$. This feature is shared with the moment estimator.¹⁶

The optimal weight matrix of the *partial* asymptotic least squares estimator conditional on the transition probability estimates $\widehat{\mathbf{g}}_T$ differs from the one of the *partial* pseudo-maximum likelihood estimator, which equals $\boldsymbol{\Sigma}_p^{-1}$, 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 $\nabla_{\mathbf{p}'} \boldsymbol{\Psi}$. Aguirregabiria and Mira (2002) show that the derivative $\nabla_{\mathbf{p}} \boldsymbol{\Psi}$ vanishes in the single-agent case, as the continuation value, which enters $\boldsymbol{\Psi}$ 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 (2007), the derivative $\nabla_{\mathbf{p}} \boldsymbol{\Psi}$ 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

14. Bajari *et al.* (2007) consider a least squares estimator based on a criterion function that squares the difference between the expected discounted pay-off associated with the observed choice and any other choice times an indicator function that equals one if the difference is negative. They use the identity matrix as a weight matrix. Bajari *et al.* (2007) also permit continuous-choice variables, but define the estimator on a finite set of choices.

15. Pakes *et al.* (2007) 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.

16. Aguirregabiria 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 $\boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})$, evaluated at the PML estimate $\widehat{\boldsymbol{\theta}}_1$, defines a new auxiliary parameter estimate, $\widehat{\mathbf{p}}_1 = \boldsymbol{\Psi}(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \widehat{\boldsymbol{\theta}}_1)$. This process can be iterated.

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.

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;¹⁷ (ii) the pseudo-maximum likelihood estimator, PML; (iii) the estimator using the identity matrix as weights, LS-I; and (iv) the *k*-step iterated pseudo-maximum likelihood estimator, k-PML, proposed by Aguirregabiria and Mira (2002, 2007). 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 *et al.* (1994), Aguirregabiria and Mira (2002, 2007), and Pakes *et al.* (2007) for the effect of alternative auxiliary estimators.

We describe the design of the Monte Carlo study in more detail in Section 7.1. Section 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 pay-off structure and model parameters are selected to 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 pay-offs 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 pay-offs can be interpreted as stemming from a game with switching costs and/or as entry/exit game. A player who selects

17. The 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.

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 who switches states from 0 to 1 incurs the entry cost c , while a player who 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 1 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 2 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, 1000, 10,000, and 100,000. For each design, we conduct 1000 repetitions of the experiment in order to obtain a distribution of the estimates. For each estimator, we report the mean, S.D., and mean squared error (MSE) 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 MSE. The MSE is summed over the three parameters and is scaled by a factor of 100.

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 1000, 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 half. As a result, the weight matrices of the LS-E, PML, and LS-I have similar properties.

TABLE 1
Monte Carlo results, equilibrium (i)

<i>T</i>	Estimator	<i>c</i>	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
1000	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)

<i>T</i>	Estimator	<i>c</i>	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
1000	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

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 fourth in 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)

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
1000	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

TABLE 4
CPU times

N	Model parameters							
	2	2	5	5	10	10	15	15
Demand levels	2	10	2	10	2	10	2	10
CPU times								
Solve model	0.98	2.70	0.66	34.46	5.24	582.70	42.31	4720.82
Estimate \hat{p}	1.12	1.12	1.12	1.19	1.13	1.62	1.21	3.52
	(0.03)	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.77)
LS-I	0.14	0.04	0.04	0.06	0.02	0.04	0.03	0.10
	(0.19)	(0.00)	(0.01)	(0.01)	(0.00)	(0.01)	(0.00)	(0.02)
LS-E	0.05	0.21	0.08	3.47	0.73	48.21	3.42	288.65
	(0.02)	(0.03)	(0.01)	(0.01)	(0.02)	(0.74)	(0.02)	(48.92)
PML	0.04	0.03	0.05	0.08	0.02	0.11	0.03	0.20
	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.02)	(0.01)	(0.03)
k-PML	0.86	0.67	0.62	1.76	0.88	9.97	1.91	42.75
	(0.01)	(0.04)	(0.01)	(0.01)	(0.02)	(0.11)	(0.02)	(9.77)

Note: For each set of model parameters, the model was solved once and simulated five times. Reported central processing unit (CPU) times are averages in seconds. S.E. are in parentheses.

with a MSE that is only about 20–30 per cent higher than that of the LS-E for moderate to large sample sizes with T equalling at least 1000.

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).

Table 4 reports central processing unit (CPU) times for finding an equilibrium numerically and for estimating the parameters using alternative estimators. We consider a Cournot Model

with a stochastic demand component that rotates the slope of the demand curve.¹⁸ We restrict the attention to symmetric equilibria in order to keep the burden of solving for the equilibrium at reasonable levels. The table illustrates that the CPU time required to estimate the model parameters increases as the number of firms and states increases, but at a slower rate than the time required to compute the equilibria. For large state spaces, the computational burden of finding an equilibrium is substantially higher than the burden of estimating the parameters. There is considerable variation in CPU times across alternative estimators. Part of the variation may be software and algorithm specific, but part may be attributable to the fact that computing the efficient weight matrix becomes increasingly burdensome when the state space increases as the size of the derivative matrix $\nabla_p \Psi$ increases. Consequently, the PML may be computationally less burdensome than the LS-E for large state spaces.

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.

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 $a'_i \in \mathbf{A}_i$:

$$u_i(a_i; \sigma_i, \theta) + \varepsilon_i^{a_i} \geq u_i(a'_i; \sigma_i, \theta) + \varepsilon_i^{a'_i} \text{ for all } a'_i \in \mathbf{A}_i, \tag{A.1}$$

18. In particular, we consider a pay-off function that arises from Cournot competition with linear demand function $P = M \cdot (a - bQ)$ with $a, b > 0$ and $M \in \{M_1, M_2, \dots, M_{\overline{M}}\}$ follows a Markov process with the typical element of the $\overline{M} \times \overline{M}$ transition matrix given by

$$t_{ij} = 0.5 \cdot \mathbf{1}_{\{i=j\}} + 0.25 \cdot \mathbf{1}_{\{i,j=1\}} + .25 \cdot \mathbf{1}_{\{i,j=\overline{M}\}} + .25 \cdot \mathbf{1}_{\{|i-j|=1\}}.$$

The fixed cost of production is equal to f_c , and the marginal cost of production is zero. Firms receive a scrap value of $x > 0$ when becoming inactive. We estimate the composite Cournot parameter $(\frac{a^2}{b})$ and the fixed cost f_c . The advantage of this design is that results are comparable, because varying the number of firms and possible demand states, leaves the number of parameters unchanged.

where $\varepsilon_i^0 = 0$. The optimality condition (A.1) can be expressed before the pay-off 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 :

$$\begin{aligned}
 p_i(a_i | s) &= \Psi_i(a_i, s, \sigma_i, \boldsymbol{\theta}) \\
 &= \int \prod_{k \in \mathbf{A}_i, k \neq a_i} \mathbf{1}(u_i(a_i; \sigma_i, \boldsymbol{\theta}) - u_i(k; \sigma_i, \boldsymbol{\theta}) \geq \varepsilon_i^k - \varepsilon_i^{a_i}) dF.
 \end{aligned}
 \tag{A.2}$$

Equation (A.2) 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 (A.2) in matrix form as

$$\mathbf{p} = \boldsymbol{\Psi}(\boldsymbol{\sigma}; \boldsymbol{\theta})
 \tag{A.3}$$

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 (A.3) as it is already determined by the remaining choice probabilities,

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

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 (A.1). 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:

$$\begin{aligned}
 a_i(\boldsymbol{\varepsilon}_i, s) &= \min\{k \in K_i(\boldsymbol{\varepsilon}_i, s)\} \\
 \text{where } K_i(\boldsymbol{\varepsilon}_i, s) &= \{k \in \mathbf{A}_i \mid u_i(k; p_i, \boldsymbol{\theta}) + \varepsilon_i^k \geq u_i(k'; p_i, \boldsymbol{\theta}) + \varepsilon_i^{k'} \text{ for all } k' \in \mathbf{A}_i\}
 \end{aligned}$$

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. \parallel

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. \parallel

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. \parallel

Lemma 1 (Equilibrium Characterization Linear in Pay-offs). *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*

$$\mathbf{X}_i(\mathbf{p}; g, \beta)\Pi_i + \mathbf{Y}_i(\mathbf{p}; g, \beta) = 0 \quad \text{for all } i \in \mathbf{N}
 \tag{A.4}$$

and any \mathbf{p} that solves (A.4) also satisfies (11).

Proof. 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 pay-offs.

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

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

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

$$a_i(\varepsilon_i, s) = \begin{cases} k & \text{if } \varepsilon_i^k > \bar{\varepsilon}_i^k(s) \text{ and for all } k' \neq k : \varepsilon_i^k - \varepsilon_i^{k'} > \bar{\varepsilon}_i^k(s) - \bar{\varepsilon}_i^{k'}(s); \\ 0 & \text{if } \varepsilon_i^k < \bar{\varepsilon}_i^k(s) \text{ for all } k, \end{cases} \tag{A.5}$$

and the choice probability is given by

$$p(a_i = k | s) = \Pr\left(\varepsilon_i^k > \bar{\varepsilon}_i^k(s) \text{ and for all } k' \neq k : \varepsilon_i^k - \varepsilon_i^{k'} > \bar{\varepsilon}_i^k(s) - \bar{\varepsilon}_i^{k'}(s)\right) \tag{A.6}$$

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 (A.6), 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 $\bar{\varepsilon}_i = [\bar{\varepsilon}_i^k(s)]_{s \in S, k=1, \dots, K}$ be the $(m_s \cdot K) \times 1$ -dimensional vector of indifferent types. Let \mathbf{P} denote the $m_s \times (m_a \cdot m_s)$ -dimensional matrix consisting of choice probability $p(a | s)$ in row s column (a_{-i}, s) and zeros in row s column (a, s') with $s' \neq s$. Further, let \mathbf{P}_{-i} be the $m_s \times ((K + 1)^{N-1} \cdot m_s)$ -dimensional matrix consisting of choice probability $p(a_{-i} s)$ in row s column (a_{-i}, s) and zeros in row s column (a_{-i}, s') with $s' \neq s$, and define the $(K \cdot m_s) \times (m_a \cdot m_s)$ -dimensional matrix $\mathbf{P}_i(\bar{\varepsilon})$ as

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

We can restate the indifference equations for player i as

$$\mathbf{P}_i(\bar{\varepsilon}) \Pi_i + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} \mathbf{V}_i + \bar{\varepsilon}_i = 0, \tag{A.7}$$

where $\bar{\varepsilon} = [\bar{\varepsilon}_i]_{i \in \mathbf{N}}$ 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) and substituting equilibrium beliefs \mathbf{P} for σ_i , we obtain the following equation system:

$$[\mathbf{P}_i(\bar{\varepsilon}) + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} [\mathbf{I}_s - \beta \mathbf{P}(\bar{\varepsilon}) \mathbf{G}]^{-1} \mathbf{P}_i] \Pi_i + \beta \mathbf{P}_i(\bar{\varepsilon}) \mathbf{G} [\mathbf{I}_s - \beta \mathbf{P}(\bar{\varepsilon}) \mathbf{G}]^{-1} \mathbf{D}_i(\bar{\varepsilon}) + \bar{\varepsilon}_i = 0. \tag{A.8}$$

Equation system (A.8) consists of $m_s \cdot K$ equations for player i . Notice also that the *ex ante* expected pay-off shock can be uniquely written in terms of indifferent types, as

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

Thus, we have shown that for any player i , there is an equilibrium characterization (A.8) 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

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

and the unknown parameters are in Π_i .

Now, take a vector \mathbf{p} that satisfies (11). We can recover the indifferent types $\bar{\boldsymbol{\varepsilon}}$ based on equation (A.6). By optimality of \mathbf{p} , the indifferent types must satisfy the indifference equations (A.4).

Conversely, take a vector $\bar{\boldsymbol{\varepsilon}}$ that satisfies (A.4). We can calculate the implied *ex ante* choice probabilities based on equation (A.6). The choice probabilities will satisfy (11) as they are the *ex ante* expected values of equilibrium decision rules satisfying the equilibrium conditions (A.4). This completes the proof. \parallel

Proof of Proposition 2. The equation system (A.4) characterized in Lemma 1 consists of $K \cdot m_s \cdot N$ equations. Hence, at most $K \cdot m_s \cdot N$ parameters in $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_N)$ can be identified. \parallel

Proof of Proposition 3. Consider the equation system (A.4) 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 $m_a \cdot m_s - K \cdot m_s$ pay-off restrictions $\mathbf{R}_i \cdot \boldsymbol{\Pi}_i = \mathbf{r}_i$ yields the $m_a \cdot m_s$ equations in $m_a \cdot m_s$ unknown pay-off parameters $\boldsymbol{\Pi}_i$

$$\begin{bmatrix} \mathbf{X}_i \\ \mathbf{R}_i \end{bmatrix} \boldsymbol{\Pi}_i + \begin{bmatrix} \mathbf{Y}_i \\ \mathbf{r}_i \end{bmatrix} = \bar{\mathbf{X}}_i \boldsymbol{\Pi}_i + \bar{\mathbf{Y}}_i \tag{A.9}$$

$$= 0. \tag{A.10}$$

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 $\boldsymbol{\Pi}_i$ is exactly identified. \parallel

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. \parallel

Proof of Proposition 5. The result follows from property 9.1 on p. 282 in *Gourieroux and Monfort (1995)*. \parallel

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

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

where the second equality uses the definition of the frequency estimator $\widehat{p}_{kis} = n_{kis} / \sum_{l=0}^K n_{lis}$ from equation (19) 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 $\widehat{p}_{0is} = 1 - \sum_{k=1}^K \widehat{p}_{kis}$ 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 $(\widehat{\mathbf{p}}_T, \widehat{\mathbf{g}}_T, \boldsymbol{\theta})$. 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. \parallel

Acknowledgements. We thank the editor Bernard Salanié and two anonymous referees, Mireia Jofre-Bonet, Alexandra Miltner, seminar participants at Alicante, Brown, Columbia, Duke, East Anglia, Lancaster, Leuven, Maryland, Munich, NYU, Queen's, UCL, Toronto, Toulouse, Yale, Washington University, the NBER summer institute, and the CEPR IO conference for helpful comments. We also thank Ariel Pakes and Peter Davis for their discussions of the paper. Martin Pesendorfer thanks the NSF under grant SES 0214222 for financial support.

REFERENCES

- AGUIRREGABIRIA, V. (1999), "The Dynamics of Markups and Inventories in Retailing Firms", *Review of Economic Studies*, **66**, 275–308.
- 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.
- AGUIRREGABIRIA, V. and MIRA, P. (2007), "Sequential Estimation of Dynamic Discrete Games", *Econometrica*, **75**, 1–53.
- BAJARI, P. (1998), "Econometrics of Sealed-Bid Auctions", *Proceedings of the Business and Economic Statistics Section of the American Statistical Association*, 41–49.
- BAJARI, P., BENKARD, L. and LEVIN, J. (2007), "Estimating Dynamic Models of Imperfect Competition", *Econometrica*, **75**, 1331–1370.
- BILLINGSLEY, P. (1968), *Convergence of Probability Measures* (New York: John Wiley).
- BRESNAHAN, T. F. and REISS, P. (1990), "Entry in Monopoly Markets", *Review of Economic Studies*, **57**, 531–553.
- BROCK, W. and DURLAUF, S. (2001), "Discrete Choice with Social Interactions", *Review of Economic Studies*, **68**, 235–260.
- CARRASCO, M. and FLORENS, J.-P. (2000), "Generalization of GMM to a Continuum of Moment Conditions", *Econometric Theory*, **16**, 797–834.
- CHAMBERLAIN, G. (1987), "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions", *Journal of Econometrics*, **34**, 305–334.
- COLLARD-WEXLER, A. (2006), "Plant Turnover and Demand Fluctuations in the Ready-Mix Concrete Industry" (Technical Report, New York University).
- DORASZELSKI, U. and JUDD, K. L. (2006), "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games" (Technical Report, Harvard University).
- DORASZELSKI, U. and SATTERTHWAITE, M. (2007), "Computable Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity" (Technical Report, Harvard University).
- ELYAKIME, B., LAFFONT, J. J., LOISEL, P. and VUONG, Q. (1994), "First-Price Sealed-Bid Auctions with Secret Reservation Prices", *Annales d'Economie et de Statistique*, **34**, 115–141.
- GOURIEROUX, C. and MONFORT, A. (1995), *Statistics and Econometric Models*, Vol. 1 (Cambridge: Cambridge University Press).
- GOURIEROUX, C., MONFORT, A. and TROGNON, A. (1985), "Moindres Carres Asymptotiques", *Annales de l'INSEE*, **58**, 91–121.
- GRUND, B. (1993), "Kernel Estimators for Cell Probabilities", *Journal of Multivariate Analysis*, **46**, 283–308.
- GUERRE, E., PERRIGNE, I. and VUONG, Q. (2000), "Optimal Nonparametric Estimation of First-Price Auctions", *Econometrica*, **68**, 525–574.
- HANSEN, L. (1982), "Large Sample Properties of Methods of Moments Estimators", *Econometrica*, **50**, 1029–1054.
- HECKMAN, J. and NAVARRO, S. (2007), "Dynamic Discrete Choice and Dynamic Treatment Effects", *Journal of Econometrics*, **136**, 341–396.
- HOTZ, J. and MILLER, R. A. (1993), "Conditional Choice Probabilities and the Estimation of Dynamic Models", *Review of Economic Studies*, **60**, 497–529.
- HOTZ, J., MILLER, R. A., SANDERS, S. and SMITH, J. (1994), "A Simulation Estimator for Dynamic Models of Discrete Choice", *Review of Economic Studies*, **61**, 265–289.
- JOFRE-BONET, M. and PESENDORFER, M. (2003), "Estimation of a Dynamic Auction Game", *Econometrica*, **71**, 1443–1489.
- MAGNAC, T. and THESMAR, D. (2002), "Identifying Dynamic Discrete Decision Processes", *Econometrica*, **70**, 801–816.
- MASKIN, E. and TIROLE, J. (1994), "Markov Perfect Equilibrium" (Technical Report, Harvard University).
- MASKIN, E. and TIROLE, J. (2001), "Markov Perfect Equilibrium I: Observable Actions", *Journal of Economic Theory*, **100**, 191–219.
- McFADDEN, D. (1978), "Modelling the Choice of Residential Location", in A. Karlquist *et al.* (eds.) *Spatial Interaction Theory and Planning Models* (Amsterdam: North Holland) 75–96.
- McKELVEY, R. D. and PALFREY, T. R. (1995), "Quantal Response Equilibria in Normal Form Games", *Games and Economic Behavior*, **10**, 6–38.
- PAKES, A. and MCGUIRE, P. (2001), "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality", *Econometrica*, **69**, 1261–1281.
- PAKES, A., OSTROVSKY, T. and BERRY, S. (2007), "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)", *RAND Journal of Economics*, Summer.
- PESENDORFER, M. and SCHMIDT-DENGLER, P. (2003), "Identification and Estimation of Dynamic Games" (NBER Working Paper No. w9726, Cambridge).

- RUST, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica*, **57**, 1027–1058.
- 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: Sixth World Congress*, Vol. 2 (Cambridge: Cambridge University Press), 119–170.
- RYAN, S. (2006), "The Costs of Environmental Regulation in a Concentrated Industry" (Technical Report, MIT).
- SEIM, K. (2006), "An Empirical Model of Firm Entry with Endogenous Product-Type Choices (Technical Report, Stanford University)", *RAND Journal of Economics*, **37**, 619–640.
- SLADE, M. (1998), "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices", *Review of Economic Studies*, **65**, 87–107.
- TAMER, E. (2003), "Incomplete Simultaneous Response Models with Multiple Equilibria", *Review of Economic Studies*, **70**, 147–165.
- TANABE, K. and SAGAE, M. (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.