

Equilibrium multiplicity in dynamic games: Testing and estimation

TAISUKE OTSU AND MARTIN PESENDORFER

London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK.

Email: t.otsu@lse.ac.uk, m.pesendorfer@lse.ac.uk

First version received: 22 October 2021; final version accepted: 31 January 2022.

Summary: This paper surveys the recent literature on dynamic games estimation when there is a concern of equilibrium multiplicity. We focus on the questions of testing for equilibrium multiplicity and estimation in the presence of multiplicity.

Keywords: *equilibrium multiplicity, dynamic game, partial identification, set inference.*

1. INTRODUCTION

During the last twenty years a substantial body of empirical works on dynamic games has emerged. The theoretical underpinnings are given by the pioneering works by Abreu et al. (1990) and Maskin and Tirole (1988), which emphasize that equilibrium outcomes are typically not a singleton. Abreu et al. (1990) propose a method to calculate the set of subgame perfect equilibria, while Maskin and Tirole (1988) illustrate equilibrium multiplicity in the context of a Markov perfect dynamic pricing game. Some empirical works have departed from the classical theory by imposition of additional ad hoc simplifying assumptions requiring that a unique and identical equilibrium is played in the cross section of markets. This paper focuses on the literature that remains with the classical theory permitting equilibrium multiplicity.

The pioneering paper by Jovanovic (1989) describes the problem of equilibrium multiplicity for statistical inference. Multiplicity arises when a complete specification of the environment does not give rise to a unique determination of the endogenous variables. Jovanovic shows that equilibrium selection will place restrictions on observables that complete the model. If an equilibrium selection mechanism is included into the model, then this requires an extra assumption, which may give rise to the possibility of misspecification; see de Paula (2013).

Tamer (2003) treats equilibrium multiplicity as an incomplete model. Incompleteness refers to the property that the model may have multiple outcomes for some (or all) parameter values, see also Bresnahan and Reiss (1991). Leaving the equilibrium selection mechanism unspecified, and assuming random sampling, the incomplete model gives rise to a set of inequalities, or bounds, that need to be satisfied by a set of parameter values. The influential works by Tamer (2003) and Ciliberto and Tamer (2009) show that these bounds contain useful information that in general leads to set identification of parameters. There is an emerging body of works that applies these ideas to inference in one-shot static games with equilibrium multiplicity; see de Paula (2013) and Aradillas-López (2020) for excellent surveys.

This paper reviews recent contributions on econometric methods under equilibrium multiplicity in the context of dynamic Markov games. Data generated by dynamic Markov games have the

feature that multiple observations over time from play of one equilibrium are observable. This distinguishing feature gives rise to an advantage over one-shot static games which can facilitate the inference problem posed by multiplicity concerns. In particular, if the time series dimension is large, then estimation can proceed using the time series for one Markov equilibrium, see Jofre-Bonet and Pesendorfer (2003). Our aim in this paper is to consider settings in which the time series dimension is short, and the researcher wishes to pool data across markets. Our focus will be on describing recent approaches that develop statistical testing procedures for poolability and parameter inference under the presence of equilibrium multiplicity using a cross section of markets. For a general discussion on backgrounds and applications of dynamic games estimation, we refer a detailed survey by Aguirregabiria et al. (2021).

The setup for dynamic Markov games is described in Section 2.1. Well-known properties are highlighted in Section 2.2. Traditional estimation approaches, such as the classic minimum distance approach based on the set of Markov perfect equilibrium conditions, are described in Section 2.3. We also review the moment inequality approach by Bajari et al. (2007). Then we discuss two important issues for these traditional approaches: effects of iterations on the asymptotic properties of different estimation methods, and accommodation of unobserved market heterogeneity. These traditional estimation methods typically focus on the case where the data used for estimation are generated from a single equilibrium.

In the following sections we discuss two strands of recent advances on econometrics for dynamic games. First, Section 3 discusses hypothesis testing for homogeneity of the data generating process or equilibrium multiplicity. We describe the asymptotic tests by Otsu et al. (2016) that examine homogeneity of the distribution of state variables, where asymptotics may refer to the number of time periods or the number of markets getting large. We also discuss recent advances on homogeneity or multiplicity testing, such as an approximately exact test by Bugni et al. (2021) and rank test based on serial correlations due to multiplicity by de Paula and Tang (2021). Second, Section 4 discusses robust inference methods under equilibrium multiplicity. We review conventional approaches for multiplicity robust inference methods and discuss two recent advances: set inference for partially identified parameters by Otsu et al. (2021), which is based on the duality approach in Schennach (2014), and nonparametric point identification analysis by Luo et al. (2022). Finally, Section 5 discusses some directions for future research.

2. SETUP AND TRADITIONAL METHODS

This section describes our setup and well-known implications, and reviews traditional estimation methods.

2.1. Setup

Consider the dynamic Markov game setup in discrete time $t = 1, 2, \dots$. We describe the setup for one market. The same setup applies for all markets $j = 1, \dots, M$.

Players. A typical player is denoted by $i = 1, \dots, N$. The number of players is fixed and does not change over time. Every period the econometrician observes a profile of states and actions described as follows.

States. Each player is endowed with two state variables (s_i^t, ε_i^t) with $s_i^t \in \mathbf{S}_i$ in finite support and $\varepsilon_i^t \in \mathbb{R}^K$ at each period t . The state variable s_i^t is publicly observed by all players. We maintain the assumption that the econometrician observes s_i^t but does not observe ε_i^t . We assume that ε_i^t is

drawn from a continuous cumulative distribution function F with support \mathbb{R}^K , observed privately by player i at the beginning of period t , and is independent from the past variables $(s_i^{t'}, \varepsilon_i^{t'})$ for all $t' < t$ and $i = 1, \dots, N$. The vector of all players' public state variables is denoted by $\mathbf{s}^t = (s_1^t, \dots, s_N^t) \in \mathbf{S} = \times_i \mathbf{S}_i$ whose cardinality is $|\mathbf{S}|$.

Actions. Each player chooses an action $a_i^t \in \mathbf{A}_i = \{0, 1, \dots, K\}$ in finite support at each period t . The decisions are made after the state (s_i^t, ε_i^t) is observed. The decisions can be made simultaneously or sequentially. The decision may also be taken after an idiosyncratic random utility (or a random profit shock) is observed. We leave the details of the decision process unspecified. The vector of joint actions in period t is denoted by $\mathbf{a}^t = (a_1^t, \dots, a_N^t) \in \mathbf{A} = \times_i \mathbf{A}_i$ whose cardinality is $|\mathbf{A}|$. We assume actions are publicly observed by all players and the econometrician. We also use the notation $\mathbf{a} = (a_i, \mathbf{a}_{-i})$ where \mathbf{a}_{-i} denotes the actions of all players other than player i . While most of our exposition considers the finite action space \mathbf{A}_i , we shall sometimes remark on the continuous action settings in which $\mathbf{A}_i = \mathbb{R}^K$.

Choice probability matrix. Let $\sigma(\mathbf{a}|\mathbf{s}) = \Pr\{\mathbf{a}^t = \mathbf{a} | \mathbf{s}^t = \mathbf{s}\}$ denote the conditional probability that an action profile \mathbf{a} will be chosen conditionally on a state \mathbf{s} . Throughout the paper we restrict attention to Markov perfect equilibrium outcomes, which imply that σ is time invariant and that given $\mathbf{s}^t, \mathbf{a}^t$ is independent from all past variables $(\mathbf{a}^{t'}, \mathbf{s}^{t'})$ for $t' < t$. The matrix of conditional choice probabilities is denoted by σ , which has dimension $|\mathbf{S}| \times (|\mathbf{A}| \cdot |\mathbf{S}|)$. It consists of conditional probabilities $\sigma(\mathbf{a}|\mathbf{s})$ in row \mathbf{s} , column (\mathbf{a}, \mathbf{s}) , and zeros in row \mathbf{s} , column $(\mathbf{a}, \mathbf{s}')$ with $\mathbf{s}' \neq \mathbf{s}$. We denote the marginal probability of action a_i with $\sigma_i(a_i|\mathbf{s}) = \sum_{\mathbf{a}_{-i}} \sigma(a_i, \mathbf{a}_{-i}|\mathbf{s})$.

State-action transition matrix. Let $g(\mathbf{s}'|\mathbf{a}, \mathbf{s}) = \Pr\{\mathbf{s}^{t+1} = \mathbf{s}' | \mathbf{a}^t = \mathbf{a}, \mathbf{s}^t = \mathbf{s}\}$ denote the state-action transition probability that a state \mathbf{s}' is reached when the current action profile and state are given by (\mathbf{a}, \mathbf{s}) . We also assume that g is time invariant and that given $(\mathbf{a}^t, \mathbf{s}^t), \mathbf{s}^{t+1}$ is independent from all past variables $(\mathbf{a}^{t'}, \mathbf{s}^{t'})$ for $t' < t$. We use the symbol \mathbf{G} to denote the $(|\mathbf{A}| \cdot |\mathbf{S}|) \times |\mathbf{S}|$ dimensional state-action transition matrix in which column $\mathbf{s}' \in \mathbf{S}$ consists of the vector of probabilities $[g(\mathbf{s}'|\mathbf{a}, \mathbf{s})]_{\mathbf{a} \in \mathbf{A}, \mathbf{s} \in \mathbf{S}}$.

State transition matrix. Under the above assumptions on σ and \mathbf{G} , the state variables \mathbf{s}^t obey a (first-order) Markov chain with the (stationary) state transition matrix $\mathbf{P} = \sigma \mathbf{G}$ whose dimension is $|\mathbf{S}| \times |\mathbf{S}|$. A typical element $p(\mathbf{s}'|\mathbf{s}) = \sum_{\mathbf{a} \in \mathbf{A}} \sigma(\mathbf{a}|\mathbf{s})g(\mathbf{s}'|\mathbf{a}, \mathbf{s})$ of the $|\mathbf{S}| \times |\mathbf{S}|$ matrix \mathbf{P} equals the probability that state \mathbf{s}' is reached when the current state is given by \mathbf{s} . Hereafter we focus on the first-order Markov chain.¹

Period payoff. After actions are observed, players collect their period payoffs $\pi_i(\mathbf{a}^t, \mathbf{s}^t, \varepsilon_i^t) \in \mathbb{R}$, and $\pi = (\pi_1, \dots, \pi_N)$. We assume that the payoff is additively separable for the a_i^t -th component of ε_i^t , i.e., $\pi_i(\mathbf{a}^t, \mathbf{s}^t, \varepsilon_i^t) = \bar{\pi}_i(\mathbf{a}^t, \mathbf{s}^t) + \varepsilon_i^t(a_i^t)$. That is, when player i chooses action $a_i^t = k$, then only the k -th element of the vector ε_i^t enters the period payoff. We denote with $\bar{\Pi}_i$ the $|\mathbf{A}| \cdot |\mathbf{S}| \times 1$ dimensional payoff vector $[\bar{\pi}_i(\mathbf{a}, \mathbf{s})]_{\mathbf{a} \in \mathbf{A}, \mathbf{s} \in \mathbf{S}}$ and with Π_i the $|\mathbf{A}| \cdot |\mathbf{S}| \times 1$ dimensional vector of expected payoffs $[\bar{\pi}_i(\mathbf{a}, \mathbf{s}) + E\varepsilon_i(a_i)]_{\mathbf{a} \in \mathbf{A}, \mathbf{s} \in \mathbf{S}}$.

¹ To illustrate how the matrices σ , \mathbf{G} , and \mathbf{P} are defined, consider binary decisions in a single player setting. Every period the player decides whether to be active or not, $a^t \in \{0, 1\}$, and the agent's state variable, $s^t \in \{0, 1\}$, evolves deterministically with $s^t = a^{t-1}$ indicating whether the player was active in the last period or not. In this case, the matrices take the following forms:

$$\sigma = \begin{bmatrix} \sigma(0|0) & \sigma(1|0) & 0 & 0 \\ 0 & 0 & \sigma(0|1) & \sigma(1|1) \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{P} = \begin{bmatrix} \sigma(0|0) & \sigma(1|0) \\ \sigma(0|1) & \sigma(1|1) \end{bmatrix}.$$

Game payoff. Players discount future payoffs with discount factor $\beta \in (0, 1)$. The expected discounted game payoff is $E \left[\sum_{t=1}^{\infty} \beta^{t-1} \pi_i(\mathbf{a}^t, \mathbf{s}^t, \varepsilon_i^t) \right]$. We sometimes denote game payoffs using the $|\mathbf{S}| \times 1$ dimensional ex ante value function $V_i(\sigma) = \sigma \Pi_i + \beta \sigma \mathbf{G} V_i(\sigma)$, or $V_i(\sigma) = (I_{|\mathbf{S}|} - \beta \sigma \mathbf{G})^{-1} \sigma \Pi_i$, where $I_{|\mathbf{S}|}$ denotes the $|\mathbf{S}| \times |\mathbf{S}|$ identity matrix.

Markovian strategies. Player i 's choice of an action is a function $a_i(\mathbf{s}', \varepsilon_i^t) \in \{0, 1, \dots, K\}$ that depends on the publicly observed state vector \mathbf{s}' and the player-specific private information ε_i^t .

Markovian perfect equilibrium (MPE). Player i 's Markovian strategy must be optimally taken given the strategy profile of rivals, and beliefs must be consistent with the strategies.

2.2. Well-known properties

We next describe some important well-known properties for our setup. They concern the equilibrium fixed point mapping, the nonuniqueness of equilibria property, and the existence of a steady-state distribution.

Equilibrium fixed point mapping. Let

$$u_i(a_i, \mathbf{s}, \sigma, \varepsilon_i) = \sum_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} \sigma_{-i}(\mathbf{a}_{-i} | \mathbf{s}) \left\{ \pi_i(\mathbf{a}, \mathbf{s}, \varepsilon_i) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{s}' | a_i, \mathbf{a}_{-i}, \mathbf{s}) V_i(\mathbf{s}', \sigma) \right\} \quad (2.1)$$

denote the continuation value of action a_i , where $\sigma_{-i}(\mathbf{a}_{-i} | \mathbf{s}) = \sum_{a_i} \sigma(a_i, \mathbf{a}_{-i} | \mathbf{s})$ and $V_i(\mathbf{s}', \sigma)$ is the element of $V_i(\sigma)$ corresponding to \mathbf{s}' . Let $1\{\cdot\}$ be the indicator function. The optimality condition of choice a_i implies that $\sigma_i(a_i | \mathbf{s}) = \int_{\varepsilon} 1\{u_i(a_i, \mathbf{s}, \sigma, \varepsilon) \geq u_i(k, \mathbf{s}, \sigma, \varepsilon)\} dF$. Let $\mathbf{p} = (\sigma_1, \dots, \sigma_N)$ be the $(N \cdot K \cdot |\mathbf{S}|) \times 1$ vector that stacks choice probabilities and let $\Psi = (\Psi_1, \dots, \Psi_N)$ be the corresponding $(N \cdot K \cdot |\mathbf{S}|) \times 1$ vector of stacked right hand side elements. We can write the optimality condition in matrix notation as

$$\mathbf{p} = \Psi(\mathbf{p} | \pi, F, \beta, \mathbf{G}). \quad (2.2)$$

The representation lemma in Aguirregabiria and Mira (2007) together with proposition 1 and theorem 1 in Pesendorfer and Schmidt-Dengler (2008) establish the following characterization and existence results.

LEMMA 2.1. *Consider the setup and assumptions of Section 2.1 for the finite action space. Equation (2.2) is a necessary and sufficient condition for choice probabilities to be consistent with an MPE. An MPE exists.*

In the continuous action setting, a set of necessary equilibrium conditions are given by the first order condition $\frac{\partial}{\partial a_i} u_i(a_i, \mathbf{s}, \mathbf{p}, \varepsilon_i) = 0$ where $u_i(a_i, \mathbf{s}, \sigma, \varepsilon_i) = \int_{\mathbf{a}_{-i} \in \mathbf{A}_{-i}} \{\pi_i(\mathbf{a}, \mathbf{s}, \varepsilon_i) + \beta \sum_{\mathbf{s}' \in \mathbf{S}} g(\mathbf{s}' | a_i, \mathbf{a}_{-i}, \mathbf{s}) V_i(\mathbf{s}', \sigma)\} d\sigma_{-i}(\mathbf{a}_{-i} | \mathbf{s})$ and σ_{-i} denotes the distribution function of rival players' actions. The first order condition can be written in the form of equation (2.2) with elements $\sigma_i(a_i | \mathbf{s}) = \sigma_i(a_i | \mathbf{s}) + \frac{\partial}{\partial a_i} u_i(a_i, \mathbf{s}, \sigma, \varepsilon)$.

Nonuniqueness. The fixed point mapping (2.2) can have multiple solutions. Pesendorfer and Schmidt-Dengler (2008) and Doraszelski and Satterthwaite (2010) provide examples of dynamic Markovian games with multiple equilibria.

Limiting steady-state distribution. When the limit exists, let $Q(\mathbf{s}', \mathbf{s}) = \lim_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T 1\{\mathbf{s}^t = \mathbf{s}', \mathbf{s}^0 = \mathbf{s}\}$ denote the long run proportion of time that the Markov chain \mathbf{P} spends in state \mathbf{s}' when starting at the initial state $\mathbf{s}^0 = \mathbf{s}$. Suppose the unconditional long run proportion of time $Q(\mathbf{s}') = \lim_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T 1\{\mathbf{s}^t = \mathbf{s}'\}$ that the Markov chain \mathbf{P} spends in state \mathbf{s}' satisfies $Q(\cdot) = Q(\cdot, \mathbf{s})$ for all initial states \mathbf{s} . Then the $|\mathbf{S}|$ dimensional row vector of

probabilities $\mathbf{Q} = \{Q(\mathbf{s})\}_{\mathbf{s} \in \mathbf{S}}$ is called the steady-state distribution of the Markov chain. Since the state space is finite, \mathbf{Q} describes a multinomial distribution. The properties of Markov chains are well known. We next describe some property useful for our purpose. To do so, we introduce the concept of communicating states.

Communicating states. We say that a state \mathbf{s}' is reachable from \mathbf{s} if there exists an integer T so that the chain \mathbf{P} will be at state \mathbf{s}' after T periods with positive probability. If \mathbf{s}' is reachable from \mathbf{s} , and \mathbf{s} is reachable from \mathbf{s}' , then the states \mathbf{s} and \mathbf{s}' are said to communicate. The random component ε_i^t having full support in the real numbers implies that all actions arise with strictly positive probability for any state $\mathbf{s} \in \mathbf{S}$. Thus, states will communicate if the state-action transition matrix allows that state \mathbf{s}' (or \mathbf{s}) can in principle be reached when starting from state \mathbf{s} (or \mathbf{s}') for any pair of states $\mathbf{s}, \mathbf{s}' \in \mathbf{S}$.

LEMMA 2.2. *Suppose all states of the Markov chain \mathbf{P} communicate (or sometimes called ergodic or irreducible). Then the steady-state distribution \mathbf{Q} exists and is unique. It satisfies $Q(\mathbf{s}) > 0$ for all $\mathbf{s} \in \mathbf{S}$ and $\mathbf{Q} = \mathbf{Q}\mathbf{P}$.*

This lemma guarantees existence and uniqueness of the steady-state distribution and states that the long run proportion of time that the Markov chain \mathbf{P} spends in state \mathbf{s} is strictly positive for any state $\mathbf{s} \in \mathbf{S}$ and the equation $\mathbf{Q} = \mathbf{Q}\mathbf{P}$ must hold. A proof of the above properties is given in proposition 1.14 and corollary 1.17 in Levin et al. (2009).

2.3. Traditional estimation methods

This section reviews the traditional inference approach.

While early work, for example Jofre-Bonet and Pesendorfer (2003), envisioned estimation for a single industry with time series data, subsequent methodological contributions extended the ideas to a cross section of markets with (possibly) short time horizon. The facilitating assumption invoked on the data generating process (DGP) is that the data consist of an identical and independently distributed (iid) sample generated from a single MPE.

For each market j , a sequence of action-state profiles $\mathbf{d}_j = (\mathbf{a}_j^t, \mathbf{s}_j^t)_{t=1, \dots, T}$ is observed, where T is the length of time periods in the data set. Let $n = M \cdot N \cdot T$ denote the total number of observations.

Assumption (DGP): The observed data $(\mathbf{d}_j)_{j=1, \dots, M}$ consist of an iid sample drawn from a single MPE.

Identification. The model is identified if there exists a unique set of model primitives (π, F, β, g) that can be inferred from a sufficiently rich data set describing choices and state transitions from a single equilibrium. Rust (1994), and Magnac and Thesmar (2002) show that exclusion restrictions or functional form assumptions on model primitives are required, and Pesendorfer and Schmidt-Dengler (2008) characterize necessary and sufficient conditions for point identification of parameters. Bajari et al. (2007) consider moment inequality models that do not require point identification.

Estimation. Let $\theta = (\theta^\pi, \theta^F, \beta, \theta^g) \in \Theta \subset \mathbb{R}^q$ denote a vector of parameters to parametrize the model primitives (π, F, β, g) . The estimation methods follow the two-step approach originally developed in Hotz and Miller (1993) for single agent settings. First, consistent estimators of choice probabilities $\hat{\mathbf{p}}_0$ are obtained from the data on actions and states. Second, the equilibrium conditions in (2.2) are invoked to estimate the parameters θ .

The first stage choice probability estimates ideally would involve a nonparametric estimation. Frequently parametric assumptions are invoked to facilitate estimation in practice. Jofre-Bonet

and Pesendorfer (2003) consider a continuous action framework in which choice probabilities are specified by a Weibull distribution function of actions, while Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008) consider a discrete action framework using the empirical frequency estimator. Bajari et al. (2007) allow for both continuous and discrete action variables.

In the second stage, a commonly used approach for discrete actions is to estimate θ based on the minimum distance (MD) problem: $\min_{\theta} \rho_n(\hat{\mathbf{p}}_0, \theta)$, where

$$\rho_n(\hat{\mathbf{p}}_0, \theta) = [\hat{\mathbf{p}}_0 - \Psi(\hat{\mathbf{p}}_0, \theta)]' W^{-1} [\hat{\mathbf{p}}_0 - \Psi(\hat{\mathbf{p}}_0, \theta)] \quad (2.3)$$

for a suitable weight matrix W .

The one-step pseudo-maximum likelihood (PML) estimator, proposed in Aguirregabiria and Mira (2007), is asymptotically equivalent to a MD estimator with the weight matrix $W = \Sigma_{\hat{\mathbf{p}}_0}$, where $\Sigma_{\hat{\mathbf{p}}_0}$ denotes the variance matrix of $\hat{\mathbf{p}}_0$. Pesendorfer and Schmidt-Dengler (2008) propose a general class of MD estimators and show that the efficient weight matrix is given by $W = [I - \nabla_{\mathbf{p}} \Psi] \Sigma_{\hat{\mathbf{p}}_0} [I - \nabla_{\mathbf{p}} \Psi]'$, where $\nabla_{\mathbf{p}} \Psi$ is the Jacobian matrix of Ψ with respect to \mathbf{p} . This efficient estimator is asymptotically equivalent to the maximum likelihood estimator. The generalized method of moments (GMM) estimator proposed in Pakes et al. (2007) is also asymptotically equivalent to a MD estimator.

For continuous actions, Jofre-Bonet and Pesendorfer (2003) use the first order conditions for optimal actions $\frac{\partial}{\partial a_i} u_i(a_i, \mathbf{s}, \hat{\mathbf{p}}_0, \varepsilon_i; \theta) = 0$ to conduct inference. The privately known ε_i enters the first order condition of a dynamic bidding game in additive separable form that enables nonparametric identification and estimation of the distribution of ε_i . Srisuma (2013) extends MD problem (2.3) to the case of continuous actions.

Bajari et al. (2007) relax the requirement of point identification of the parameter vector θ . Instead they consider set identified models and propose a two-step estimator for the identified set of parameters. In particular, based on the characterization of an MPE σ , $V_i(\mathbf{s}, \sigma_i, \sigma_{-i}; \theta) \geq V_i(\mathbf{s}, \sigma'_i, \sigma_{-i}; \theta)$ for any alternative σ'_i , they consider the population criterion function

$$\int [\min\{V_i(\mathbf{s}, \sigma_i, \sigma_{-i}; \theta) - V_i(\mathbf{s}, \sigma'_i, \sigma_{-i}; \theta), 0\}]^2 dH(i, \mathbf{s}, \sigma'_i),$$

where H is some distribution over a set for $(i, \mathbf{s}, \sigma'_i)$. Bajari et al. (2007) suggest constructing its empirical counterpart $\rho_n^Y(\theta)$ by estimating the value function $V_i(\cdot; \theta)$ for given θ in the first step and evaluating the integral by simulation. Their set estimator is

$$\hat{\Theta} = \left\{ \theta \in \Theta : \rho_n^Y(\theta) \leq \min_{\theta' \in \Theta} \rho_n^Y(\theta') + \mu_n \right\},$$

for some suitably chosen positive sequence μ_n that decays to zero and may depend on the data; see, e.g., Chernozhukov et al. (2007). Furthermore, their moment inequality approach allows action spaces to be continuous.

Some authors have considered iterating the fixed point equation (2.2). Unfortunately, this can lead to poor statistical properties due to multiplicity of solutions in the fixed point equation inherent to games. Consider the iterated (nested) PML estimator, proposed in Aguirregabiria and Mira (2007). It entails solving the above MD problem iteratively by replacing $\hat{\mathbf{p}}_0$ in Ψ in the ℓ -th iteration with $\hat{\mathbf{p}}_{\ell}$ obtained from $\hat{\mathbf{p}}_{\ell} = \Psi(\theta_{\ell}, \hat{\mathbf{p}}_{\ell-1})$. Pesendorfer and Schmidt-Dengler (2010) show analytically in a dynamic game framework with a unique symmetric equilibrium that this type of iteration can lead almost surely to inconsistent (Lyapunov stable) estimates. For a fixed number of iterations, Bugni and Bunting (2020) study asymptotic properties of the iterated PML and

MD estimators. They show that the above described optimally-weighted one-step MD estimator is asymptotically more efficient than iterated MD and iterated PML estimators. Kasahara and Shimotsu (2012) propose modified algorithms with the aim to alleviate the convergence issues of iterated methods. The computation time of the K-PML and K-MD estimators, which may differ in applications, does not appear to be driven by differences in the computational complexities. This can be seen from the fact that both the K-PML and K-MD problems can be viewed as MD problems which differ only in the weight matrices used.

As indicated in Pesendorfer and Schmidt-Dengler (2008) and reconfirmed in Bugni and Bunting (2020), the asymptotic theory of the MD estimator of θ becomes analogous to that of the standard GMM theory (e.g., Newey and McFadden, 1994) under standard conditions. GMM theory directly implies that the optimally-weighted one-step MD estimator is asymptotically optimal. There is no gain achievable by iteration. Other estimation methods, such as the (iterated) PML, are typically inefficient for estimating dynamic games.

It is insightful to note that in the single agent setup, the fixed point mapping for \mathbf{p} typically satisfies the zero Jacobian with respect to \mathbf{p} at the fixed point so that additional policy mapping iterations have no first-order effect on the asymptotic distribution (Aguirregabiria and Mira, 2002). In contrast, the mapping Ψ for dynamic games typically fails to satisfy such a zero Jacobian property. To recover this zero Jacobian property, Dearing and Blevins (2021) develop an efficient version of the iterative PML estimator by considering an alternative mapping (say, $\tilde{\Psi}$) which approximates the full Newton step of the original constrained maximum likelihood problem. They show that the iterative maximum likelihood estimator using $\tilde{\Psi}$ is asymptotically efficient and the first-order asymptotic properties of the estimator do not vary across iterations.

Unobserved heterogeneity. The benchmark setup discussed above has been extended in various directions. In particular, several papers develop estimation methods to accommodate unobserved market heterogeneity. Aguirregabiria and Mira (2007) extend their PML estimator to allow discrete and finitely supported unobserved heterogeneity in which the pseudo-likelihood function is specified by the finite mixture form. Arcidiacono and Miller (2011) and Connault (2015) extend the estimation approach to account for unobserved heterogeneity in state variables which refers to an element that is known to players and may evolve over time in a Markovian way, but is not observed by the econometrician. Importantly, the initial unobserved heterogeneity for a specific market is an iid draw from a known probability distribution.

Other extensions. Srisuma and Linton (2012) extend two-step games estimators to allow for a continuous state space. Doraszelski and Judd (2012) and Arcidiacono et al. (2016) consider random (poisson) arrival of decision nodes which reduces the computational complexity of the ex ante value function calculation. Egedal et al. (2015) propose a constrained maximum likelihood estimator. Miessi Sanches et al. (2016) propose conditions under which the minimization problem in (2.3) simplifies to standard OLS/GLS. Abbring and Campbell (2010) study last-in first-out dynamics, which typically implies a unique MPE. Abbring et al. (2018) propose a simple dynamic model of entry/exit which has an essentially unique symmetric MPE.²

² Although this paper focuses on the MPE, several papers consider alternative equilibrium concepts. Weintraub et al. (2008) and Ifrach and Weintraub (2016) weaken the dependence of firms' strategies on rival firms' state variables. Aguirregabiria and Magesan (2019) relax the equilibrium consistency condition of firms' beliefs. Fershtman and Pakes (2012) and Asker et al. (2020) relax the equilibrium conditions and introduce the experience-based-equilibrium concept.

3. MULTIPLICITY: TESTING

A maintained assumption in the conventional estimation approaches is that the DGP is identical across markets. Researchers assume that action-state profiles are generated from an identical equilibrium in all markets, and policy functions can be estimated by pooling those data across markets. If this assumption does not hold, then policy functions will be inconsistently estimated and the resulting inference for structural parameters will be erroneous. We next review the literature that explores statistical tests of validity of this pooling assumption and also presence of multiple equilibria.

Based on the setup described in the previous section, the DGP of the profile $(\mathbf{a}_j^t, \mathbf{s}_j^t)_{t=1, \dots, T}$ of the j -th market is characterized by the $|\mathbf{S}| \times (|\mathbf{A}| \cdot |\mathbf{S}|)$ dimensional conditional choice probability matrix \mathbf{p}_j , which consists of conditional probabilities $\Pr\{\mathbf{a}_j^t = \mathbf{a} | \mathbf{s}_j^t = \mathbf{s}\}$ in row \mathbf{s} , column (\mathbf{a}, \mathbf{s}) , and zeros in row \mathbf{s} , column $(\mathbf{a}, \mathbf{s}')$ with $\mathbf{s}' \neq \mathbf{s}$, and the $(|\mathbf{A}| \cdot |\mathbf{S}|) \times |\mathbf{S}|$ dimensional state-action transition matrix \mathbf{G}_j for $\Pr\{\mathbf{s}_j^{t+1} = \mathbf{s}' | \mathbf{a}_j^t = \mathbf{a}, \mathbf{s}_j^t = \mathbf{s}\}$. Note that the transition matrix of states is written as $\mathbf{p}_j \mathbf{G}_j$. If all states of the Markov chain $\mathbf{p}_j \mathbf{G}_j$ communicate, then by Lemma 2.2, there exists a unique steady-state distribution \mathbf{Q}_j and the identical distribution hypothesis across markets $j = 1, \dots, M$ may be tested by homogeneity of the steady-state distribution

$$H_0^Q : \mathbf{Q}_1 = \dots = \mathbf{Q}_M. \quad (3.1)$$

The $|\mathbf{S}| \times 1$ dimensional vector of relative frequencies $\widehat{\mathbf{Q}}_j = [T^{-1} \sum_{t=1}^T 1\{\mathbf{s}_j^t = \mathbf{s}\}]_{\mathbf{s} \in \mathbf{S}}$ is a nonparametric estimator of the steady-state distribution \mathbf{Q}_j . Otsu et al. (2016) show that a test statistic for H_0^Q is

$$\mathcal{T}_Q = T \sum_{j=1}^M (\widehat{\mathbf{Q}}_j - \widehat{\mathbf{Q}}) \widehat{\mathbf{V}}^{-1} (\widehat{\mathbf{Q}}_j - \widehat{\mathbf{Q}}) \xrightarrow{d} \chi^2((M-1)(|\mathbf{S}|-1)), \quad (3.2)$$

under H_0^Q , where $\widehat{\mathbf{Q}} = M^{-1} \sum_{j=1}^M \widehat{\mathbf{Q}}_j$ and $\widehat{\mathbf{V}}^{-1}$ means a generalized inverse of $\widehat{\mathbf{V}}$, the asymptotic variance estimator of $T^{1/2}(\widehat{\mathbf{Q}}_j - \mathbf{Q}_j)$ by using, e.g., the Newey and West (1987) estimator. The term $\widehat{\mathbf{V}}$ in (3.2) may be replaced with the identity matrix and a researcher can employ a bootstrap critical value. For example, for every bootstrap iteration, a set of states can be randomly chosen from the steady-state distribution $\widehat{\mathbf{Q}}$ for every market j .³

The above test requires that the steady-state distributions exist and that the Markov chains are in the steady-state; see Lemma 2.2. That is, regardless of H_0^Q , it is assumed that all states in the chain $\mathbf{p}_j \mathbf{G}_j$ communicate for each j .

Otsu et al. (2016) relax this assumption and propose an alternative test based on the conditional state distribution given the initial state. This situation is relevant in new industries in which the steady-state has not been reached yet, or when there are some absorbing states. Using the state transition matrix $\mathbf{P} = \mathbf{pG}$, the conditional distribution $\mathbf{s}^t | \mathbf{s}^1 = \mathbf{s}$ is described by $\iota_s' (\mathbf{pG})^t$, where ι_s takes one at the element corresponding to \mathbf{s} and zero otherwise. The null hypothesis of identical DGP can be considered in the form of

$$H_0^s : \left[\frac{1}{T-1} \sum_{t=2}^T \Pr\{\mathbf{s}^t = \mathbf{s}' | \mathbf{s}^1 = \mathbf{s}\} \right]_{\mathbf{s}' \in \mathbf{S}} = \frac{1}{T-1} \sum_{t=2}^T \iota_s' \mathbf{P}^t. \quad (3.3)$$

³ Otsu et al. (2016) also propose testing procedures for homogeneity of the conditional choice probability matrices (i.e., $H_0^p : \mathbf{p}_1 = \dots = \mathbf{p}_M$) and state-action transition matrices (i.e., $H_0^g : \mathbf{G}_1 = \dots = \mathbf{G}_M$) for a given M . These tests are also asymptotically valid as $T \rightarrow \infty$.

The left hand side is a vector of model-free conditional probabilities, and the right hand side is the model-based prediction for those probabilities. The hypothesis H_0^s is implied from two assumptions: (i) the state variables $(s_j^t)_{t=1, \dots, T}$ for $j = 1, \dots, M$ are iid over j , which allows us to express the hypothesis H_0^s without using a market index j ; and (ii) the Markov chain is first-order and time-homogeneous. The state transition matrix $\mathbf{P} = [\Pr\{s^{t+1} = s' | s^t = s\}]_{s, s' \in \mathcal{S}}$ on the right hand side of (3.3) can be estimated by the frequency estimator $\hat{\mathbf{P}} = \left[\frac{\sum_{j=1}^M \sum_{t=1}^{T-1} 1\{s_j^{t+1} = s', s_j^t = s\}}{\sum_{j=1}^M \sum_{t=1}^{T-1} 1\{s_j^t = s\}} \right]_{s, s' \in \mathcal{S}}$.

Let $\hat{\mathbf{Q}}_s^t = \left[\frac{\sum_{j=1}^M 1\{s_j^t = s', s_j^{t-1} = s\}}{\sum_{j=1}^M 1\{s_j^{t-1} = s\}} \right]_{s' \in \mathcal{S}}$ be the relative frequency estimator for the vector of conditional probabilities $[\Pr\{s^t = s' | s^{t-1} = s\}]_{s' \in \mathcal{S}}$ for $t = 2, \dots, T$ for a given initial state s . If the model parametrized by \mathbf{pG} is correct, then a test statistic can be defined based on the contrast $\mathbf{C}'_s = (T - 1)^{-1} (\sum_{t=2}^T \hat{\mathbf{Q}}_s^t - t'_s \sum_{t=2}^T \hat{\mathbf{P}}^t)$, which satisfies

$$\mathcal{T}_s = M \mathbf{C}'_s \hat{\mathbf{V}}_s^- \mathbf{C}_s \xrightarrow{d} \chi^2(|\mathcal{S}| - 1), \tag{3.4}$$

as $M \rightarrow \infty$ with fixed T under H_0^s , where $\hat{\mathbf{V}}_s^-$ is a generalized inverse of an estimator of the asymptotic variance of $\sqrt{M} \mathbf{C}_s$ under H_0^s . The test based on \mathcal{T}_s requires $T \geq 3$. As in (3.2), the term $\hat{\mathbf{V}}_s$ in (3.4) may be replaced with the identity matrix and a bootstrap critical value can be employed.

In contrast to the asymptotic tests by Otsu et al. (2016), Bugni et al. (2021) propose an approximation to the exact randomization test (see chapter 15.2 of Lehmann and Romano, 2005) of the homogeneity hypothesis across markets and time periods

$$H_0^{\text{hom}} : \Pr\{\mathbf{a}_j^t = \cdot | s_j^t = \cdot\} \text{ and } \Pr\{s_j^{t+1} = \cdot | \mathbf{a}_j^t = \cdot, s_j^t = \cdot\} \text{ are identical for all } j, t,$$

by employing a Markov chain Monte Carlo (MCMC) algorithm. Under the setup in Section 2.1 and H_0^{hom} , the sufficient statistic of the data $\mathbf{d} = (\mathbf{a}_j^t, s_j^t)_{j=1, \dots, M, t=1, \dots, T}$ is given by

$$U(\mathbf{d}) = \left(\begin{array}{c} \{s_j^1\}_{j=1}^M, \left\{ \sum_{j=1}^M 1\{\mathbf{a}_j^T = \mathbf{a}, s_j^T = s\} \right\}_{\mathbf{a} \in \mathbf{A}, s \in \mathcal{S}}, \\ \left\{ \sum_{j=1}^M \sum_{t=1}^{T-1} 1\{\mathbf{a}_j^t = \mathbf{a}, s_j^t = s, s_j^{t+1} = s'\} \right\}_{\mathbf{a} \in \mathbf{A}, s \in \mathcal{S}, s' \in \mathcal{S}} \end{array} \right),$$

and thus in principle we can consider the exact randomization test.

Let \mathbf{D} be support of \mathbf{d} , and define the set Γ of mappings from \mathbf{D} to \mathbf{D} such that each element γ of Γ satisfies $U(\tilde{\mathbf{d}}) = U(\gamma(\tilde{\mathbf{d}}))$ for all $\tilde{\mathbf{d}} \in \mathbf{D}$. Then the randomization test with significance level α is given by

$$\phi_{\text{rd}}(\mathbf{d}) = 1 \left\{ \left[\frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} 1\{\tau(\gamma(\mathbf{d})) \geq \tau(\mathbf{d})\} \right] \leq \alpha \right\},$$

where $|\Gamma|$ is the number of elements of Γ and $\tau(\cdot)$ is any scalar-valued function to define the test statistic. This test is exactly valid; i.e., under H_0^{hom} , the rejection probability satisfies $E[\phi_{\text{rd}}(\mathbf{d})] \leq \alpha$ for any finite N, T , and M (Lehmann and Romano, 2005). However, since Γ is usually difficult to enumerate, this test is practically infeasible. To address this practical issue, Bugni et al. (2021) propose an MCMC-based approximation for $\phi_{\text{rd}}(\mathbf{d})$. In particular, they construct an iterative

algorithm which randomly picks an element $\gamma^{(k)}(\mathbf{d})$ from Γ , develop the approximate test

$$\phi_{\text{rd}}^{(K)}(\mathbf{d}) = 1 \left\{ \left[\frac{1}{K} \sum_{k=1}^K 1\{\tau(\gamma^{(k)}(\mathbf{d})) \geq \tau(\mathbf{d})\} \right] \leq \alpha \right\},$$

and show its validity as the number of iterations increases; i.e., $\limsup_{K \rightarrow \infty} E[\phi_{\text{rd}}^{(K)}(\mathbf{d})] \leq \alpha$ for any finite N , T , and M . The main idea of their iterative algorithm is to draw the action and state profiles $(\mathbf{a}^{(k)}, \mathbf{s}^{(k)}) = \gamma^{(k)}(\mathbf{d})$ separately from configurations that yield the same value of the sufficient statistic $U(\mathbf{d})$.

de Paula and Tang (2021) study testable implications of multiple equilibria in static and dynamic discrete games, where players' private signals are correlated but the econometrician can split the sample into clusters within which equilibrium selections under multiplicity are correlated. In the context of dynamic games, correlation among players' private signals naturally arises through time-varying serially correlated unobserved heterogeneity, say η^t , and the clusters can be constructed by matching players across different periods in the same dynamic game.

To be specific, consider the setup in Section 2.1, where the publicly observable state variables are augmented as (\mathbf{s}^t, η^t) . Although \mathbf{s}^t is observable to the econometrician, η^t is not. We partition disjointly the set of players, \mathcal{I}_A and \mathcal{I}_B , and then define $\mathbf{a}_A^t = \{a_i^t : i \in \mathcal{I}_A\}$, $\mathbf{a}_B^t = \{a_i^t : i \in \mathcal{I}_B\}$, and matrices \mathbf{P}_{AB} and $\tilde{\mathbf{P}}_{AB}$ for joint probability masses of $(\mathbf{a}_A^t, \mathbf{a}_B^t)$ and $(\mathbf{a}_A^{t-1}, \mathbf{a}_B^{t+1})$, respectively. Under the assumption that a single equilibrium is played in each market and some rank conditions on the matrices defined by conditional probability masses of $P_\omega\{\mathbf{a}_A^t, \mathbf{a}_B^t | \eta^t = \eta\}$ for each MPE $\omega = 1, \dots, d_\omega$ conditional on η , de Paula and Tang (2021) explore mixture representations of \mathbf{P}_{AB} and $\tilde{\mathbf{P}}_{AB}$ by unobserved heterogeneity and equilibrium selection, and show that the number of MPE is identified as $d_\omega = \{\text{rank}(\mathbf{P}_{AB})\}^2 / \text{rank}(\tilde{\mathbf{P}}_{AB})$. Therefore, inference on the number of MPE including uniqueness can be conducted by applying the rank test in Kleibergen and Paap (2006).

4. MULTIPLICITY: INFERENCE

What can we do if the tests described in the last section reject the null of homogeneous DGP across markets or unique equilibrium? This section reviews some existing approaches and recent advances for inference on structural parameters under equilibrium multiplicity.

First, a researcher may search for a subsample where the null of homogeneity is accepted. Second, unobserved-heterogeneity robust inference methods, such as Aguirregabiria and Mira (2007), Arcidiacono and Miller (2011), and Connault (2015), can be adopted under some further assumptions to estimate structural parameters under the presence of equilibrium multiplicity. However the following caveats arise: unobserved heterogeneity typically affects choice probabilities and payoffs jointly, while equilibrium multiplicity affects choice probabilities only. Hence, the channel in which unobserved heterogeneity enters has to exclude payoff effects and focus on choice probability effects. Second, the unobserved-heterogeneity robust inference methods postulate a unique outcome arising with a given probability. As the pioneering work by Tamer (2003) shows, postulating a particular probability distribution over equilibrium outcomes is an 'ad hoc' assumption that changes the model and can lead to inconsistent estimates if specified incorrectly. It requires knowledge of the 'equilibrium selection rule'. Third, parametric specifications employed for the distribution of unobserved heterogeneity are typically restrictive and may cause inconsistency. For instance, a recent paper that accounts for unobserved heterogeneity in

dynamic games is Arcidiacono et al. (2016). It postulates a parametric functional form in which unobserved heterogeneity affects choice probabilities. Yet, if this approach were adopted to account for multiplicity, then this would translate into a particular functional form how multiplicity may affect choice probabilities. If this functional form is misspecified, then this may lead to inconsistent parameter estimates.

For large T , Bajari et al. (2007) propose to estimate choice probabilities for each market separately. To conduct inference that accounts for equilibrium multiplicity, the second step estimator can be based on pooling all equilibrium conditions from all markets. This approach permits consistent inference that accommodates different equilibria in different markets. Yet, when T is small the market-specific choice probabilities can no longer be consistently estimated.

4.1. Set inference

An inference method for the identified set of parameters which does not require specification of the equilibrium selection rule is proposed in Otsu et al. (2021). They consider the asymptotic setting where T fixed and the number of markets M diverges. In this subsection we describe this approach.

The econometrician observes action-state profiles $\mathbf{d}_j = (\mathbf{a}_j^t, \mathbf{s}_j^t)_{t=1, \dots, T}$ generated by the equilibrium play of some underlying economic model in market $j = 1, \dots, M$. The market-specific choice probabilities \mathbf{p}_j are usually not observed in the data, and treated as random variables. It is assumed that $(\mathbf{d}_j, \mathbf{p}_j)_{j=1, \dots, M}$ is an iid sequence. Then Otsu et al. (2021) consider the $d = 2 \cdot K \cdot |\mathbf{S}| \cdot N$ dimensional moment function for the observable \mathbf{d} and unobservable \mathbf{p} :

$$g(\mathbf{d}, \mathbf{p}; \theta) = \begin{bmatrix} [\mathbf{p}^{(a_i, \mathbf{s}, i)} - \Psi^{(a_i, \mathbf{s}, i)}(\mathbf{p}; \theta)]_{a_i \in \mathbf{A}_i, \mathbf{s} \in \mathbf{S}, i \in \mathbf{N}} \\ [f^{(\mathbf{s})} \mathbf{p}^{(a_i, \mathbf{s}, i)} - f^{(a_i, \mathbf{s}, i)}]_{a_i \in \mathbf{A}_i, \mathbf{s} \in \mathbf{S}, i \in \mathbf{N}} \end{bmatrix}, \quad (4.1)$$

where the upper script ' (a_i, \mathbf{s}, i) ' means the element corresponding to player i , action a_i , and state \mathbf{s} . The upper block in (4.1) consists of the equilibrium restrictions (2.2) and the lower block contains moment restrictions placed on the choice probabilities by the data. Here $f^{(a_i, \mathbf{s}, i)} = \sum_{t=1}^T 1\{a_i^t = a_i, \mathbf{s}^t = \mathbf{s}\}$ is the frequency of action-state profile (a_i, \mathbf{s}) , and $f^{(\mathbf{s})} = \sum_{t=1}^T 1\{\mathbf{s}^t = \mathbf{s}\}$ is the frequency of state \mathbf{s} . Although $E[g(\mathbf{d}, \mathbf{p}; \theta)] = 0$ holds, GMM cannot be applied to estimate θ due to the latent variables \mathbf{p} .

In this setup, following Schennach (2014), the identified region for the structural parameters θ can be defined as

$$\Theta_0 = \left\{ \theta \in \Theta : \inf_{\lambda \in \Lambda} |E_{\lambda \times \mu}[g(\mathbf{d}, \mathbf{p}; \theta)]| = 0 \right\}, \quad (4.2)$$

where μ is the probability measure of the observables \mathbf{d} , λ is the conditional probability measure of \mathbf{p} given \mathbf{d} , and Λ is the set of all conditional probability measures supported on the set of choice probabilities. In other words, the identified region Θ_0 is the set of parameters in which some conditional measure λ can rationalize the moment conditions. The expectation $E_{\lambda \times \mu}[\cdot]$ is infeasible to compute because the true distribution λ of the equilibrium choice probabilities is unknown and the econometrician does not want to specify the distributional form. As described in Section 2, the fixed point problem in (2.2) can have multiple solutions, and the set inference approach in Otsu et al. (2021) explicitly allows such a scenario. In particular, based on the duality approach in Schennach (2014), they treat the market-specific choice probabilities \mathbf{p} as latent variables with their conditional probability measure λ treated as an infinite-dimensional nuisance

parameter, and show by entropic latent variable integration that

$$\theta \in \Theta_0 \text{ if and only if } \inf_{\gamma \in \mathbb{R}^d} |E_\mu[\tilde{g}(\mathbf{d}, \theta, \gamma)]| = 0, \quad (4.3)$$

where

$$\tilde{g}(\mathbf{d}, \theta, \gamma) = \frac{\int g(\mathbf{d}, \mathbf{p}; \theta) \cdot \exp(\gamma' g(\mathbf{d}, \mathbf{p}; \theta)) dF(\mathbf{p})}{\int \exp(\gamma' g(\mathbf{d}, \mathbf{p}; \theta)) dF(\mathbf{p})},$$

with a user-specified probability measure F . In practice, F may be set as the uniform distribution on the strategy set \mathcal{Q} . A similar result can be obtained for the case where \mathbf{p} is parametrized by $\boldsymbol{\tau} \in \bar{\mathcal{Q}}$ (say, $\mathbf{p}(\boldsymbol{\tau})$), as far as $\bar{\mathcal{Q}}$ is compact and $\Psi(\mathbf{p}(\cdot); \theta)$ is continuous at each $\theta \in \Theta$. The dual problem (4.3) is equivalent to the primal problem (4.2). The moment condition in the dual involves the function $\tilde{g}(\mathbf{d}, \theta, \gamma)$ that is an integral of the original moment function $g(\mathbf{d}, \mathbf{p}; \theta)$ under an exponential distribution of the unobservables. Schennach (2014) shows this selected exponential distribution has the property of being the ‘least favorable’ distribution of the unobservables.

Note that, while the intermediate problem (4.2) requires optimization over the infinite-dimensional set Λ , the equivalent problem (4.3) entails optimization over only a finite-dimensional Euclidean space for γ . This property allows for a practically feasible characterization of the identified region Θ_0 for the structural parameters. Based on the data $(\mathbf{d}_j)_{j=1, \dots, M}$, we can replace the population moment in (4.3) with the sample analogue $\bar{g}(\theta, \gamma) = M^{-1} \sum_{j=1}^M \tilde{g}(\mathbf{d}_j, \theta, \gamma)$ and define a GMM-type criterion as

$$Q(\theta) = \sup_{\gamma \in \mathbb{R}^d} -\frac{1}{2} \bar{g}(\theta, \gamma)' \hat{V}(\theta, \gamma)^{-1} \bar{g}(\theta, \gamma), \quad (4.4)$$

where $\hat{V}(\theta, \gamma)$ is some estimator of $Var(\tilde{g}(\mathbf{d}, \theta, \gamma))$. A simple but conservative confidence set suggested by Schennach (2014) is

$$\hat{C} = \{\theta \in \Theta : -2MQ(\theta) \leq \chi_{d, \alpha}^2\}, \quad (4.5)$$

where $\chi_{d, \alpha}^2$ is the $(1 - \alpha)$ -th quantile of the χ^2 distribution with degree of freedom d . Schennach (2014) shows its asymptotic validity in the sense that $\lim_{n \rightarrow \infty} \Pr\{\theta \notin \hat{C}\} \leq \alpha$ for all $\theta \in \Theta_0$.

A drawback of the confidence set \hat{C} is: when the dimension d of the moment function \tilde{g} (or γ) is high, the critical value tends to be large. Recall that $d = m_s \cdot N \cdot K + 1$. Thus, if the number of states, players, or actions is large, the confidence set may be too large to obtain a meaningful conclusion.

Otsu et al. (2021) propose an adapted version of Kleibergen (2005) statistic to the moment function \tilde{g} defined by the entropic latent variable integration. For each θ , let $\tilde{\gamma}(\theta)$ be an estimator of the solution in (4.4) satisfying certain regularity conditions. Then the test statistic is written as

$$K(\theta) = M \bar{g}(\theta, \tilde{\gamma}(\theta))' \hat{V}(\theta, \tilde{\gamma}(\theta))^{-1} \hat{D}(\theta) [\hat{D}(\theta)' \hat{V}(\theta, \tilde{\gamma}(\theta))^{-1} \hat{D}(\theta)]^{-1} \\ \times \hat{D}(\theta)' \hat{V}(\theta, \tilde{\gamma}(\theta))^{-1} \bar{g}(\theta, \tilde{\gamma}(\theta)), \quad (4.6)$$

where $\hat{D}(\theta)$ is a $d \times q$ matrix with the l -th column

$$\frac{\partial \bar{g}(\theta, \tilde{\gamma}(\theta))}{\partial \theta_l} - \hat{G}_l(\theta, \tilde{\gamma}(\theta))' \hat{V}(\theta, \tilde{\gamma}(\theta))^{-1} \bar{g}(\theta, \tilde{\gamma}(\theta)),$$

and $\hat{G}_l(\theta, \gamma) = M^{-1} \sum_{j=1}^M \tilde{g}(\mathbf{d}_j, \theta, \gamma) \partial \tilde{g}(\mathbf{d}_j, \theta, \gamma) / \partial \theta_l$. In the limit, the distribution of this statistic satisfies $K(\theta_0) \xrightarrow{d} \chi_q^2$ for each $\theta_0 \in \Theta_0$. The confidence set based on this statistic is obtained as

$$\tilde{C} = \{\theta \in \Theta : K(\theta) \leq \chi_{q,\alpha}^2\}, \quad (4.7)$$

where $\chi_{q,\alpha}^2$ is the $(1 - \alpha)$ -th quantile of the χ^2 distribution with degree of freedom q . Note that the critical value $\chi_{q,\alpha}^2$ depends only on the dimension of structural parameters θ and is robust to the dimension of the moment function \tilde{g} .

In practice, the dimensionality of the function g can be high, which may pose computational difficulties, or may not be practical for the data at hand. We briefly outline three ways in which the dimensionality is typically reduced in applied work. First, if the data are not sufficiently rich to consider a frequency distribution of the vector \mathbf{p} for every realization of the market-state-action-player set, then the second element in (4.1) can be replaced with moment conditions of action variables possibly aggregated across states. Second, to facilitate computation of the d -dimensional strategies \mathbf{p} , the strategy function may be parametrized with a low dimensional parameter vector $\boldsymbol{\tau} \in \tilde{\mathcal{Q}} \subseteq \mathbb{R}^l$, by assuming $\mathbf{p}^{(a_i, s, i)} = \mathbf{p}(a_i, \mathbf{s}; \boldsymbol{\tau})$. This approach is commonly used for empirical dynamic models and referred to as a policy function approximation. Equation (2.2) is then replaced with the corresponding fixed point equation for $\boldsymbol{\tau}$ defined on $\tilde{\mathcal{Q}}$. Finally, for computational reasons it can be useful to aggregate the first element in equation (4.1) across action-state-players into a single equation. The equilibrium condition can be equivalently formulated as a problem of finding zeros of the following one-dimensional equation: $\sum_{i \in \mathbf{N}} \sum_{\mathbf{s} \in \mathbf{S}} \sum_{a_i \in \mathbf{A}_i} \{\mathbf{p}^{(a_i, s, i)} - \Psi^{(a_i, s, i)}(\mathbf{p}; \theta)\}^2 = 0$.

Observe that this estimation approach readily accommodates unobserved payoff elements as nuisance parameters as well. Suppose the period payoff has additionally an additive payoff shock v_{ij} , which is time invariant and market specific. The equilibrium equation system (2.2) becomes $\mathbf{p} = \Psi(\mathbf{p}, \mathbf{v}; \theta)$ and the function g is redefined accordingly. With these modifications in place, the structural parameters can be partially identified in the presence of unobserved heterogeneity as before. This is a notable difference to the approach of Arcidiacono and Miller (2011), which requires the researcher to use a multinomial distribution for unobserved heterogeneity.

4.2. Point identification

Luo et al. (2022) investigate nonparametric identification of the model in Section 2.1, but allowing both multiple equilibria and unobservable market heterogeneity which may vary with time. Let η^t be a state variable with finite support that is publicly observable to players but is unobservable to econometricians, and e^* be the index for equilibria conditional choice probabilities with finite support (i.e., $\{\mathbf{p} : \mathbf{p} = \Psi(\mathbf{p}; \pi, F, \beta, \mathbf{G})\} = \{\mathbf{p}^{e^*} : e^* = 1, \dots, |e^*|\}$). Then define the unobserved ‘type’ variable $\tau^t = \tau(\eta^t, e^*)$ with finite support $\{\tau_1, \dots, \tau_{|\tau|}\}$ for unknown $|\tau|$. In the first stage, Luo et al. (2022) identify the conditional choice probabilities of $\mathbf{a}^t | \mathbf{s}^t, \tau^t$ and transition probabilities $\mathbf{s}^{t+1} | \mathbf{a}^t, \mathbf{s}^t, \tau^t$ by applying the eigenvalue-eigenvector decomposition technique to identify mixture models (e.g., Hu and Shum, 2012) based on four periods of data. In the second stage, they employ the conventional identification argument as in Pesendorfer and Schmidt-Dengler (2008) to identify payoff primitives based on the identified conditional choice and transition probabilities. Furthermore, by comparing the identified payoff primitives at different values of τ^t , one can also distinguish the unobservable heterogeneity and multiple equilibria.

If the payoff functions are same (or different) at different values of τ^t , they represent multiple equilibria (or unobserved heterogeneity).

Let $\mathbf{d}^t = (\mathbf{a}^t, \mathbf{s}^t)$ for $t = 1, 2, 3, 4$ be observables for four periods. The key step in the identification analysis by Luo et al. (2022) is the first stage, where the conditional probabilities for $(\mathbf{d}^3, \tau^3)|(\mathbf{d}^2, \tau^2)$ and $\mathbf{d}^4|(\mathbf{d}^3, \tau^3)$ are identified so that the conditional choice probabilities for $\mathbf{a}^3|(\mathbf{s}^3, \tau^3)$ and transition probabilities for $\mathbf{s}^4|\mathbf{a}^3, \mathbf{s}^3, \tau^3$ are implied. First of all, since (\mathbf{d}^t, τ^t) follows a first-order Markov process, the joint distribution of three periods of observables can be written in a multiplicatively separable way (given (\mathbf{d}^2, τ^2))

$$\Pr\{\mathbf{d}^3, \mathbf{d}^2, \mathbf{d}^1\} = \sum_{\tau^2} \Pr\{\mathbf{d}^3|\mathbf{d}^2, \tau^2\} \Pr\{\mathbf{d}^2, \tau^2, \mathbf{d}^1\},$$

i.e., correlation between \mathbf{d}^3 and \mathbf{d}^1 given \mathbf{d}^2 is associated with unobserved heterogeneity and/or multiple equilibria. Indeed Luo et al. (2022) show that under certain rank conditions, the number of types $|\tau|$ is identified as the maximum rank of the probability matrix $\Pr\{\mathbf{d}^3 = \cdot, \bar{\mathbf{d}}^2, \mathbf{d}^1 = \cdot\}$ over $\bar{\mathbf{d}}^2$. Second, by assuming ‘limited feedback’ from the current unobserved state variable η^t onto the next-period observable state \mathbf{s}^t in the sense that $\Pr\{\mathbf{s}^t, \eta^t|\mathbf{s}^t, \eta^t\} = \Pr\{\mathbf{s}^t|\eta^t, \mathbf{s}^t\} \Pr\{\eta^t|\mathbf{s}^t, \eta^t\}$, the joint distribution of four periods of observables admits the following mixture representations

$$\begin{aligned} \Pr\{\mathbf{d}^4, \mathbf{d}^3, \mathbf{d}^2, \mathbf{d}^1\} &= \sum_{\tau^3} \Pr\{\mathbf{d}^4|\mathbf{d}^3, \tau^3\} \Pr\{\mathbf{d}^3|\tau^3, \mathbf{d}^2\} \Pr\{\tau^3, \mathbf{d}^2, \mathbf{d}^1\} \\ &= \sum_{\tau^3} \Pr\{\mathbf{d}^4|\mathbf{d}^3, \tau^3\} \left[\sum_{\tau^2} \Pr\{\mathbf{d}^3, \tau^3|\mathbf{d}^2, \tau^2\} \Pr\{\mathbf{d}^2, \mathbf{d}^1, \tau^2\} \right], \end{aligned}$$

which can be utilized to identify the distributions of $\mathbf{d}^4|(\mathbf{d}^3, \tau^3)$ and $(\mathbf{d}^3, \tau^3)|(\mathbf{d}^2, \tau^2)$, respectively, via the eigenvalue-eigenvector decomposition technique under some rank and variable linkage conditions to deal with the type-matching issue induced by up-to-label-swapping identification of mixtures.

5. CONCLUSION

There are several interesting directions of future research. First, while Maskin and Tirole (2001) conceptualize Markovian strategies as a way to capture the *simplest form of behaviour that is consistent with rationality*, which leads to the notion of payoff relevant state variables, applied work has taken a richer interpretation of strategies’ dependence on state variables. Collard-Wexler (2013) allow strategies to depend on plant size, past plant size for each firm, plus fifty market-level demand levels, which leads to a state space consisting of 1.4 million elements. Arcidiacono et al. (2016) allow strategies to be a function of the number of stores operated by each firm, population, and five time-invariant unobserved variables equal to discretized values of a standard normal. A rich strategy space potentially allows a large set of equilibria that should be balanced against the original aim of Markovian strategies. An open research question concerns the minimal size of the state space required to capture rational behaviours consistent with the data.

Second, if the required state space is indeed large as suggested by applied work, then applications of machine learning tools to conduct inference on dynamic games will be a promising direction. For example, Semenova (2018) extends the two-step set inference approach in Bajari et al. (2007), where the state space (and thus the dimension of \mathbf{p}) is high-dimensional, by applying the Neyman orthogonalized moment function and cross fitting to deal with the bias in

the first stage estimation of \mathbf{p} using machine learning methods (e.g., Chernozhukov et al., 2018). This methodology may be extended to other estimation strategies and accommodate unobserved heterogeneity and multiple equilibria.

Third, it is also interesting to bridge the gap to the subgame-perfect dynamic games framework. Miller et al. (2021) estimate a novel dynamic price-leadership game, in which one firm proposes super-markups to rivals (over Bertrand prices) in order to maximize its own discounted payoff subject to participation constraints of rivals, which entails punishment with reversion to permanent static Bertrand pricing. The game, which can be solved using techniques of Abreu et al. (1990), is shown to provide a good fit for the observed pricing path. A future research agenda is to build on the recursive methods introduced by Abreu et al. (1990) and to develop an estimable class of subgame-perfect equilibria that relaxes the Markovian strategy framework.

CONFLICTS OF INTEREST

There is no potential conflicts of interest.

REFERENCES

- Abbring, J. H. and J. R. Campbell (2010). Last-in first-out oligopoly dynamics. *Econometrica* 78(5), 1491–527.
- Abbring, J. H., J. R. Campbell, J. Tilly and N. Yang (2018). Very simple Markov-perfect industry dynamics: Theory. *Econometrica* 55(2), 721–35.
- Abreu, D., D. Pearce and E. Stacchetti (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58(5), 1041–63.
- Aguirregabiria, V. and A. Magesan (2019). Identification and estimation of dynamic games when players' beliefs are not in equilibrium. *Review of Economic Studies* 87(2), 582–625.
- Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica* 70(4), 1519–43.
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75(1), 1–53.
- Aguirregabiria, V., A. Collard-Wexler and S. P. Ryan (2021). Dynamic games in empirical industrial organization. NBER Working Paper 29291.
- Aradillas-López, A. (2020). The econometrics of static games. *Annual Review of Economics* 12(1), 135–65.
- Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79(6), 1823–67.
- Arcidiacono, P., P. Bayer, J. Blevins and P. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *Review of Economic Studies* 83, 889–931.
- Asker, J., C. Fershtman, J. Jeon and A. Pakes (2020). A computational framework for analyzing dynamic auctions: The market impact of information sharing. *RAND Journal of Economics* 51(3), 805–39.
- Bajari, P., L. C. Benkard and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75(5), 571–99.
- Bresnahan, T. and P. C. Reiss (1991). Empirical models of discrete games. *Journal of Econometrics* 48(1–2), 57–81.
- Bugni, F. A. and J. Bunting (2020). On the iterated estimation of dynamic discrete choice games. *Review of Economic Studies* 88(3), 1031–73.

- Bugni, F. A., J. Bunting and T. Ura (2021). Testing homogeneity in dynamic discrete games in finite samples. Working Paper. arXiv preprint: arXiv:2010.02297.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* 21, C1–C68.
- Chernozhukov, V., H. Hong and E. Tamer (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica* 75(5), 1243–84.
- Ciliberto, F. and E. Tamer (2009). Market structure and multiple equilibria in airline markets. *Econometrica* 77, 1791–828.
- Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica* 81(3), 1003–37.
- Connault, B. (2015). Chapter 1, in "Essays on Structural Dynamic Discrete Models", PhD dissertation, Princeton University, Princeton N.J.
- Dearing, A. and J. R. Blevins (2021). Efficient and convergent sequential pseudo-likelihood estimation of dynamic discrete games. Working Paper, arXiv:1912.10488.
- de Paula, A. (2013). Econometric analysis of games with multiple equilibria. *Annual Review of Economics* 5(1), 107–31.
- de Paula, A. and X. Tang (2021). Testable implications of multiple equilibria in discrete games with correlated types. Working Paper, arXiv:2012.00787.
- Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics* 3(1), 53–93.
- Doraszelski, U. and M. Satterthwaite (2010). Computable Markov-perfect industry dynamics. *RAND Journal of Economics* 41(2), 215–43.
- Egedal, M., Z. Lai and C.-L. Su (2015). Estimating dynamic discrete-choice games of incomplete information. *Quantitative Economics* 6(3), 567–97.
- Fershtman, C. and A. Pakes (2012). Dynamic games with asymmetric information: A framework for empirical work. *Quarterly Journal of Economics* 127(4), 1611–61.
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60(3), 497–529.
- Hu, Y. and M. Shum (2012). Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics* 171(1), 32–44.
- Ifrach, B. and G. Y. Weintraub (2016). A framework for dynamic oligopoly in concentrated industries. *Review of Economic Studies* 84(3), 1106–50.
- Jofre-Bonet, M. and M. Pesendorfer (2003). Estimation of a dynamic auction game. *Econometrica* 71, 1443–89.
- Jovanovic, B. (1989). Observable implications of models with multiple equilibria. *Econometrica* 57(6), 1431–37.
- Kasahara, H. and K. Shimotsu (2012). Sequential estimation of structural models with a fixed point constraint. *Econometrica* 80(5), 2303–19.
- Kleibergen, F. (2005). Testing parameters in gmm without assuming that they are identified. *Econometrica* 73(4), 1103–23.
- Kleibergen, F. and R. Paap (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics* 133(1), 97–126.
- Lehmann, E. L. and J. P. Romano (2005). *Testing Statistical Hypothesis: Third Edition*. Springer: New York.
- Levin, D. A., Y. Peres and E. L. Wilmer (2009). *Markov Chains and Mixing Times*. Providence, RI: American Mathematical Society.

- Luo, Y., P. Xiao and R. Xiao (2022). Identification of dynamic games with unobserved heterogeneity and multiple equilibria. *Journal of Econometrics* 226(2), 343–67.
- Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica* 70(2), 801–16.
- Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica* 56(3), 571–99.
- Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory* 100(2), 191–219.
- Miessi Sanches, F. A., D. J. Silva and S. Srisuma (2016). Ordinary least squares estimation of a dynamic game model. *International Economic Review* 57(2), 623–34.
- Miller, N. H., G. Sheu and M. C. Weinberg (2021). Oligopolistic price leadership and mergers: The United States beer industry. *American Economic Review* 111(10), 3123–59.
- Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. In R. F. Engle and D. McFadden (Eds.), *Handbook of Econometrics*(1 edn), vol. 4, pp. 2111–245, Elsevier: Amsterdam.
- Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55(3), 703–8.
- Otsu, T., M. Pesendorfer and Y. Takahashi (2016). Pooling data across markets in dynamic Markov games. *Quantitative Economics* 7(2), 523–59.
- Otsu, T., M. Pesendorfer, Y. Takahashi and Y. Sasaki (2021). Estimation of (static or dynamic) games under equilibrium multiplicity. *International Economic Review*, forthcoming.
- Pakes, A., M. Ostrovsky and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *RAND Journal of Economics* 38(2), 373–99.
- Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75, 901–28.
- Pesendorfer, M. and P. Schmidt-Dengler (2010). Sequential estimation of dynamic discrete games. *Econometrica* 78, 833–42.
- Rust, J. (1994). Structural estimation of markov decision processes. In R. F. Engle and D. McFadden (Eds.), *Handbook of Econometrics*(1 edn), vol. 4, pp. 3081–143, Elsevier: Amsterdam.
- Schennach, S. M. (2014). Entropic latent variable integration via simulation. *Econometrica* 82(1), 345–85.
- Semenova, V. (2018). Machine learning for dynamic discrete choice. Working Paper, arXiv:1808.02569.
- Srisuma, S. (2013). Minimum distance estimators for dynamic games. *Quantitative Economics* 4, 549–83.
- Srisuma, S. and O. Linton (2012). Semiparametric estimation of Markov decision processes with continuous state space. *Journal of Econometrics* 166(2), 320–41.
- Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies* 70(1), 147–65.
- Weintraub, G. Y., C. L. Benkard and B. Van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica* 76(6), 1375–411.

Co-editor Jaap Abbring handled this manuscript.