Nonparametric and Semiparametric Analysis of a Dynamic Discrete Game

Patrick Bajari,Victor Chernozhukov, Han Hong and Denis Nekipelov[1]

This version: April 2009

Abstract

In this paper, we study the identification and estimation of a dynamic discrete game allowing for discrete or continuous state variables. We first provide a general nonparametric identification result under the imposition of an exclusion restrictions on agents payoffs. Next we analyze large sample statistical properties of nonparametric and semiparametric estimators for the econometric dynamic game model. Numerical simulations are used to demonstrate the finite sample properties of the dynamic game estimators.

Key words: Incomplete information, Static and dynamic games, Nonparametric identification, Semiparametric estimation.

[1]Departments of Economics, University of Minnesota, MIT, Stanford University and UC Berkeley, USA.
The authors acknowledge generous research supports from the NSF and able research assistance from Timothy Armstrong and Boyu Wang. We would like to thank our colleagues and numerous conference and seminar participants for helpful comments. The usual disclaimer applies.

1 Introduction

In this paper, we study the identification and estimation of a dynamic discrete game. A dynamic discrete game is a generalization of a dynamic discrete choice model as in Rust (1987), Hotz and Miller (1993). As in these earlier papers, agents in the model are assumed to solve a dynamic programing problem. Payoffs in each time period depend on the agent's actions, the state variables and random preference shocks. Given current choices, the state variables evolve according to a law of motion which can depend on an agent's actions. A dynamic game generalizes this single agent model to allow the payoffs of one agent to depend on the actions of other agents. Dynamic game models are applicable in many areas such as industrial organization dynamic oligopoly with collusions, e.g. Fershtman and Pakes (2009). Recently, a number of papers have proposed methods to estimate dynamic games including Aguirregabiria and Mira (2007), Berry, Pakes, and Ostrovsky (2003), Pesendorfer and Schmidt-Dengler (2003), Bajari, Benkard, and Levin (2007) and Jenkins, Liu, McFadden, and Matzkin (2004).

The first goal of this paper is to study nonparametric identification of dynamic discrete games. As in Aguirregabiria and Mira (2007), Berry, Pakes, and Ostrovsky (2003) and Pesendorfer and Schmidt-Dengler (2003), we assume that agents move simultaneously in each time period and that the random preference shocks are private information. However, our framework is more general since we allow the state variables to be discrete or continuous. This is attractive for empirical work since in many applications state variables are naturally modeled as continuous. We show that our model parameters are identified if the researcher is willing to make exclusion restrictions, that is, not all state variables can enter the payoffs of all agents. Such restrictions are commonly imposed in empirical research. For example, cost and demand shifters for one firm are frequently excluded from the payoffs of other firms.

Second, we analyze semiparametric and non-parametric estimation procedures. We begin with the analysis of a semiparametric setup, where we only use non-parametric identification assumptions and parameterize the identifiable payoffs of the players without additional restrictions on the state transition law. We find the semiparametric efficiency bound for the payoff parameters, which is the minimum variance of the parameter estimates without parametric assumptions regarding the state transition. Moreover, we show that obtaining the semiparametrically efficient estimates does not require solving for equilibria of the game and computing the corresponding likelihood function. We demonstrate that the treatment of the player's decision problem as a moment equation, generated by her first-order condition allows use to estimate the payoff parameters in one step. We also show that the estimation procedure that allows one to achieve the semiparametric efficiency bound belongs to our class of one-step estimation methods. This is a new approach to the analysis of dynamic games and it generalizes the existing two-step estimation techniques such as those proposed by Aguirregabiria and Mira (2007), Berry, Pakes, and Ostrovsky (2003) and Pesendorfer and Schmidt-Dengler (2003).

An additional advantage of our approach is that it does not rely on the discreteness of the state space which is in particular achieved by using an estimation approach that does not need preliminary estimation of the continuation values of the players. In applied work, many researchers choose to discretize a continuous state variable. Increasing the number of grid points in a two step estimator reduces the bias of the first stage. However, this comes at the cost of increasing the variance of the first stage estimates. In fact, when there are four or more continuous state variables, it can be shown that it is not possible to obtain through discretization $\sqrt{T}$ consistent and asymptotically normal parameter estimates in the second stage, where T is the sample size. Therefore, discretizing the state space does not provide a solution to continuous state variables. The estimation approach of Bajari, Benkard, and Levin (2007) allows for continuous state variables. However, it requires a parametric first stage and the resulting estimates will be biased if the first stage is misspecified.

Third, we find that the reduction of the estimation procedure to one stage allows us to estimate payoffs of players fully non-parametrically. The structure of the non-parametric estimator is based on the player's first-order condition similarly to the semiparametric case. The estimates of the payoff function have a slower than parametric convergence rate. This rate depends on the smoothness of the distribution of the state transition as well as on the support condition on the policy functions of the players. We analyze the non-parametric estimator from the perspective of the mean-square optimality and offer a choice of trimming for the sieve representation of the payoff functions as well as the value functions that provides the procedure with the minimum mean squared error while converging at an optimal non-parametric rate. For the non-parametric estimator we develop a unified large sample theory that nests both continuous and discrete state variables as special cases.

Section 2 discusses identification in a static discrete game model. Section 3 extends the identification analysis to a dynamic game. Section 4 develops nonparametric and semi-parametric estimation methods which follow the lines of the identification conditions to construct estimates for the payoffs based on the non-parametric estimates of the conditional choice probabilities. Section 5 demonstrates how the multi-stage estimation strategy can be improved by representing the decision problem of a player in a dynamic game as a conditional moment equation. Moreover, it demonstrates how to obtain the estimates for the payoff parameters with the minimum variance over the class of models without parametric assumptions regarding the choice probabilities. It also derives the asymptotic distribution for the estimates of the payoff parameters. Section 8 concludes.

2 Nonparametric identification of static games

We begin by describing the model for the case of static games. This serves two purposes. First, this will allow us to discuss our modeling assumptions in a simpler setting. Second, we prove the identification for the static model. This will highlight some key ideas in our identification of the full dynamic model and also will be used as a step in the identification of the more general dynamic model.

In the model, there are a finite number of players $i = 1, ..., n$. Each player simultaneously chooses an action $a_i \in \{0, 1, \ldots, K\}$ out of a finite set. We assume that the set of actions are identical across players. This assumption is for notational convenience only and could easily be relaxed. Let $A = \{0, 1, \ldots, K\}^n$ denote the set of possible actions for all players and $a = (a_1, ..., a_n)$ denote a generic element of A. Also, let $a_{-i} = (a_1, ...a_{i-1}, a_{i+1}, ..., a_n)$ denote a vector of strategies for all players excluding i. The vector $s_i \in S_i$ denotes the state variable for player i. The set S_i can be discrete, continuous or both. Also, define $S = \Pi_i S_i$ and let $s = (s_1, ..., s_n) \in S$ denote a vector of state variables for all n players. We assume that s is common knowledge to all players in the game and is observable to the econometrician.

For each agent, there are $K+1$ private shocks $\epsilon_i(a_i)$ indexed by the actions a_i. Let $\varepsilon_i = (\varepsilon_i(0), ..., \varepsilon_i(K))$ have a density $f(\epsilon_i)$ and assume that the shocks ϵ_i are i.i.d across agents and actions a_i. We shall assume that $\epsilon_i(a_i)$ is distributed extreme value.

Assumption 1 *The error terms $\epsilon_i(a_i)$ are distributed i.i.d. across actions and agents. Furthermore, the error term has an extreme value distribution with density*

$$f(\epsilon_i) = \exp(-\epsilon_i)\exp(-\exp(-\epsilon_i)).$$

We could easily weaken this assumption. However, it is commonly used in the applied literature and will allow us to write a number of formulas in closed form which will simply our study of identification. The vNM utility function for player i is:

$$u_i(a, s, \epsilon_i) = \Pi_i(a_i, a_{-i}, s) + \epsilon_i(a_i).$$

In the above, $\Pi_i(a_i, a_{-i}, s)$ is a scalar which depends on i's own actions, the actions of all other agents a_{-i} and the entire vector of state variables s. We assume that the iid preference shocks $\epsilon_i(a_i)$ are private information for player i. The assumption that the error term $\epsilon_i(a_i)$ is private information is not universal in the literature. For example, Bresnahan and Reiss (1991) assume that the error terms are common knowledge. However, this model requires quite different econometric methods which account for the presence of multiple equilibrium and the possibility of mixed strategies.

A strategy for agent i is a function $a_i = \delta_i(s, \epsilon_i)$ which maps the state s and agent i's private information ϵ_i to an action a_i. Note that agent i's strategy does not depend on ϵ_{-i} since this is assumed to be private information to the other agents in the game. Define

$$\sigma_i(a_i = k|s) = \int 1\{\delta_i(s, \epsilon_i) = k\} f(\epsilon_i) d\epsilon_i.$$

This is the probability that agent i will play strategy k after we margin out ϵ_i.

In equilibrium, player i's belief is that j will play strategy k is $\sigma_j(a_j = k|s)$. Therefore, i's expected utility from choosing the strategy a_i is $\sum_{a_{-i}} \Pi_i(a_i, a_{-i}, s)\sigma_{-i}(a_{-i}|s) + \epsilon_i(a_i)$.

Moving forward, it will be useful to define the *choice specific value function* as

$$\Pi_i(a_i, s) = \sum_{a_{-i}} \Pi_i(a_i, a_{-i}, s)\sigma_{-i}(a_{-i}|s). \tag{1}$$

Note that we can write the expected utility from choosing a_i as

$$\Pi_i(a_i, s) + \epsilon_i(a_i). \tag{2}$$

Recall that the error terms are distributed extreme value. Standard results about the logit model plus the definition of the choice specific value function imply that

$$\sigma_i(a_i = k|s) = \frac{\exp(\Pi_i(a_i, s))}{\sum_{a_i' \in A_i} \exp(\Pi_i(a_i', s))}.$$

Definition 1 *Fix the state s. A Bayes-Nash equilibrium is a collection of probabilities, $\sigma_i^*(a_i = k|s)$ for $i = 1, ..., n$ and $k = 0, ..., K$ such that for all i and all k*

$$\sigma_i^*(a_i = k|s) = \frac{\exp(\Pi_i(a_i, s))}{\sum_{a_i' \in A_i} \exp(\Pi_i(a_i', s))} \quad \textit{and}$$
$$\Pi_i(a_i, s) = \sum_{a_{-i}} \Pi_i(a_i, a_{-i}, s)\sigma_{-i}^*(a_{-i}|s).$$

An equilibrium requires the actions of all players to be a best response to the actions of all other players. Moving forward, it is convenient to define an equilibrium in terms of $\sigma_i(a_i = k|s)$ instead of $\delta_i(s, \epsilon_i)$.

2.1 Identification of the static model

An important question is whether it is possible for us to identify the parameters of our model. One approach to identification is to impose parametric restrictions on $\Pi_i(a, s)$. In what follows, we allow $\Pi_i(a, s)$ to be a general function of s and a do not specify the payoffs $\Pi_i(a_i, s)$ parametrically. We identify $\Pi_i(a_i, a_{-i}, s_i)$ nonparametrically by imposing exclusion restrictions on this function.

Definition 2 *Let $\Pi_i(a_i, a_{-i}, s_i)$ and $\widetilde{\Pi}_i(a_i, a_{-i}, s_i)$ be two different specifications of the payoffs that are not identical, i.e. $\Pi_i(a_i, a_{-i}, s_i) \neq \widetilde{\Pi}_i(a_i, a_{-i}, s_i)$. Also let $\sigma_i(a_i = k|s)$ and $\widetilde{\sigma}_i(a_i = k|s)$ be the corresponding equilibrium choice probabilities for $i = 1, ..., n$. We say that our model is identified if $\Pi_i(a_i, a_{-i}, s_i) \neq \widetilde{\Pi}_i(a_i, a_{-i}, s_i)$ for a subset of s_i with positive probability implies that $\sigma_i(a_i = k|s) \neq \widetilde{\sigma}_i(a_i = k|s)$ for a subset of s with positive probability.*

Our proof of identification is constructive. Assuming that the population probabilities $\sigma_i(a_i = k|s)$ for all k, i and s are known, we reverse engineer the $\Pi_i(a_i, a_{-i}, s_i)$ that rationalize the data. Simple algebra implies that

$$\sigma_i(a_i = k|s) = \frac{\exp(\Pi_i(a_i, s))}{\sum_{a_i' \in A_i} \exp(\Pi_i(a_i', s))} \tag{3}$$

$$\log(\sigma_i(a_i = k|s)) - \log(\sigma_i(a_i = 0|s)) = \Pi_i(a_i = k, s) - \Pi_i(a_i = 0, s) \tag{4}$$

Equation (4) is the well known Hotz-Miller inversion. This equation implies that it is possible to learn the choice specific payoff functions, $\Pi_i(a_i = k, s)$ up to a first difference from knowledge of the choice probabilities $\sigma_i(a_i = k|s)$. Since these choice-specific payoff functions can only be learned up to a first difference, we need to impose the normalization that an "outside good" action always yields zero utility:

$$\Pi_i(a_i = 0, a_{-i}, s) = 0. \tag{5}$$

Assumption 2 *For all i, all a_{-i} and all s,* $\Pi_i(a_i = 0, a_{-i}, s) = 0$.

Having identified the choice specific payoff functions $\Pi_i(a, s)$, we next turn to the problem of identifying primitive mean utilities $\Pi_i(a_i, a_{-i}, s)$. The definition of the choice specific payoff function implies that these two objects are related by the following equation:

$$\Pi_i(a_i, s) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s)\Pi_i(a_i, a_{-i}, s), \qquad \forall i = 1, \ldots, n, a_i = 1, \ldots, K. \tag{6}$$

Given s, this is a system of $n \times K$ equations, because there are n agents and for each agent, there are $K+1$ choices. Then $\Pi_i(a_i, a_{-i}, s)$ are $n \times K \times (K+1)^{n-1}$ free parameters in equation (6). Recall that for each agent, we have normalized the utility for the action $a_i = 0$ to zero regardless of the actions of the other players. Therefore, for each agent i, there are $K \times (K+1)^{n-1}$ free parameters corresponding to the actions K actions available to i which yield nonzero utility and the $(K+1)^{n-1}$ actions of the other agents. Clearly, $n \times K \times (K+1)^{n-1} > n \times K$, which implies that the model is underidentified.

In order to identify the model, we will impose exclusion restrictions on i's payoffs. Partition $s = (s_i, s_{-i})$, and assume that

$$\Pi_i(a_i, a_{-i}, s) = \Pi_i(a_i, a_{-i}, s_i) \tag{7}$$

depends only on the subvector s_i. In other words, we are excluding some component of s from i's payoffs. Such assumptions are commonly used in applied work. For example, many oligopoly models predict that after we control for $-i$'s strategies, i's profits are not influenced by certain cost or demand shifters for $-i$.

If we impose these exclusion restrictions, we can rewrite (6) as

$$\Pi_i(a_i, s_{-i}, s_i) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s_{-i}, s_i)\Pi_i(a_i, a_{-i}, s_i), \tag{8}$$

If there are $(K+1)^{n-1}$ points in the support of the conditional distribution of s_{-i} given s_i, we will have more equations than unknowns.

Theorem 1 *Suppose that Assumptions 1 and 2 hold. Also suppose that for each s_i, there exist $(K+1)^{n-1}$ points in the support of the conditional distribution of s_{-i} given s_i. Assume that the $(K+1)^{n-1}$ equations defined by (8) are linearly independent almost surely. Then the latent utilities $\Pi_i(a_i, s_{-i}, s_i)$ are identified for almost every s_i and s_{-i}.*

We can alternatively state a rank condition, similar to the linear least squares regression model, that is sufficient for identification. This rank condition requires that given each s_i, the second moment matrix of the "regressors" $\sigma_{-i}(a_{-i}|s_{-i}, s_i)$,

$$E\sigma_{-i}(a_{-i}|s_{-i}, s_i)\sigma_{-i}(a_{-i}|s_{-i}, s_i)' \tag{9}$$

is nonsingular. Intuitively, we interpret $\Pi_i(a_i, s_{-i}, s_i)$ as the dependent variable in an ols regression and $\sigma_{-i}(a_{-i}|s_{-i}, s_i)$ as a regressor.

3 Nonparametric identification of dynamic games

3.1 Dynamic game of incomplete information

In this section, we extend our model to allow for non-trivial dynamics. Our model is similar to the framework proposed by Aguirregabiria and Mira (2007), Berry, Pakes, and Ostrovsky (2003) and Pesendorfer and Schmidt-Dengler (2003). Period returns are defined using a static logit model. The current actions a and state influence the future evolution of the state variable. We shall restrict attention to Markov perfect equilibrium. The methods that we propose here could be applied to other dynamic games, such as a finite horizon where payoffs and the low of motion are time dependent. These extensions require considerable additional notational complexity.

3.2 The Environment

3.2.1 Payoffs

In the model, there are $t = 1, ..., \infty$ time periods. At each time t, we let $a_{it} \in \{0, 1, \ldots, K\}$ denote the choice for agent i. We shall assume that the choice set is identical for all agents and does not depend on the state variable. Both assumptions could be dropped at the cost of notational complexity. Let $s_{i,t} \in S_i$ denote the state variable for agent i at time t. As in the previous section, S_i is a collection of real valued vectors and the state can either be continuous or discrete.

Let $\varepsilon_{it} = (\varepsilon_{it}(0), ..., \varepsilon_{it}(K))$ denote a vector of iid shocks to agent i's payoffs at time t. As in the previous section, we shall assume that the error terms are distributed extreme value. Player i's period utility function is

$$u_i(a_{it}, a_{-it}, s_t, \epsilon_{it}) = \Pi_i(a_{it}, a_{-it}, s_t) + \epsilon_{it}(a_{it}).$$

As in the previous section, we shall develop the model assuming that $\Pi_i(a_{it}, a_{-it}, s_t)$ is a general function of the state variables rather than a member of a particular parametric family. Let $\sigma_i(a_i|s)$ denote the probability that i plays a_i given that the state is s. As in the previous section, we define $\Pi_i(a_{it}, s)$ as $\Pi_i(a_i, s) = \sum_{a_{-i}} \Pi_i(a_i, a_{-i}, s)\sigma_{-i}(a_{-i}|s)$.

3.3 Value Functions

In the model, the evolution of the state variable depends on the current state and the actions of all players. We assume that the state variable evolves according to a first order Markov process $\overline{g}(s'|s, a_i, a_{-i})$. As before, s is perfectly observed by the agent and the econometrician. Player i maximizes expected discounted utility using a discount factor β.

Let $W_i(s, \epsilon_i; \sigma)$ be player i's value function given s and ϵ_i. The value function holds fixed the strategies of the other agents σ_{-i}. The value function then satisfies the following recursive relationship:

$$W_i(s,\epsilon_i;\sigma_{-i}) = \max_{a_i\in A_i}\Big\{\Pi_i(a_i,s)+\epsilon_i(a_i) \qquad (10)$$
$$+\beta\int\sum_{a_{-i}} W_i(s',\epsilon_i';\sigma_{-i})\overline{g}(s'|s,a_i,a_{-i})\sigma_{-i}(a_{-i}|s)f(\epsilon_i')d\epsilon_i'ds'\Big\}.$$

At each state, agents choose an action $a_i \in A_i$ to maximize expected discounted utility. The term $\Pi_i(a_i,s)+\epsilon_i(a_i)$ is the current period return from choosing a_i. The second term captures i's utility from future time periods. In our model, agents choose their actions simultaneously. Therefore, agent i's beliefs about the evolution of the state given his current information will be $\sum_{a_{-i}}\overline{g}(s'|s,a_i,a_{-i})\sigma_{-i}(a_{-i}|s)$. This integrates out agent i's uncertain about the actions of $-i$. The agent also needs to take into account expectations about next periods preference shocks, ϵ_i', by integrating out their distribution using the density $f(\epsilon_i')$.

Definition 3 *A Markov perfect equilibrium is a collection of policy functions, $\delta_i(s,\epsilon_i)$ and corresponding conditional choice probabilities $\sigma_i(a_i|s)$ such that for all i, all s and all ϵ_i, $\delta_i(s,\epsilon_i)$ maximizes the value function $W_i(s,\epsilon_i;\sigma_{-i})$*

$$W_i(s,\epsilon_i;\sigma_{-i}) = \max_{a_i\in A_i}\Big\{\Pi_i(a_i,s;\sigma_{-i})+\epsilon_i(a_i)$$
$$+\beta\int\sum_{a_{-i}} W_i(s',\epsilon_i';\sigma_{-i})\overline{g}(s'|s,a_i,a_{-i})\sigma_{-i}(a_{-i}|s)f(\epsilon_i')d\epsilon_i'ds'\Big\}.$$

In a Markov perfect equilibrium, an agent's strategy $\delta_i(s,\epsilon_i)$ is restricted to be a function of the state (s,ϵ_i). This solution concepts restricts equilibrium behavior by not allowing for time dependent punishment strategies, such as trigger strategies or tit-for-tat which do not depend on payoff relevant state variables. While the Markov perfect equilibrium assumption restricts behavior considerably, it has the advantage that equilibrium behavior can be expressed using familiar techniques from dynamic programming. Since the focus of this paper is on nonparametric identification and estimation, existence of equilibrium will be taken as given in the following analysis.

3.4 Nonparametric identification

Next, we turn to the problem of identification of the model. The strategy for identifying the model will be similar to the static model. We begin with some preliminaries by first defining the choice specific value function and deriving some key equations that must hold in our dynamic model.

The starting point of our analysis is to define the choice specific value function

$$V_i(a_i, s) = \Pi_i(a_i, s) + \beta \int \sum_{a_{-i}} W_i(s', \epsilon_i'; \sigma)\overline{g}(s'|s, a_i, a_{-i})\sigma_{-i}(a_{-i}|s) f(\epsilon_i') d\epsilon_i' ds'. \tag{11}$$

Similar to (1), the choice specific value function is the expected utility from choosing the action a_i, excluding the current period error term $\epsilon_i(a_i)$. As in the static setting, the term $\Pi_i(a_i, s)$ integrates out player i's expectations about the actions of the other players. In a dynamic setting, however, we have to include the utility from future time periods. We do this by integrating out the value function $W_i(s', \epsilon_i'; \sigma)$ with respect to next periods private information, ϵ_i', and state s'. In words, we can interpret the choice specific value function as the returns, excluding $\epsilon_i(a_i)$, from choosing a_i today and then reverting to the solution to the dynamic programming problem (10) in all future time periods. Next, we define the ex ante value function, or social surplus function, as

$$V_i(s) = \int W_i(s, \epsilon_i; \sigma) f(\epsilon_i) d\epsilon_i \tag{12}$$

The ex ante value function is the expected value of W_i tomorrow given that the state today is s. In order to compute this expectation, we integrate over the distribution of s and ϵ_i given that the current state is s.

Using equations (11) and (12), the ex ante and choice specific value functions are related to each other through the following equation

$$V_i(a_i, s) = \Pi_i(a_i, s) + \beta E\left[V_i(s')|s, a_i\right]. \tag{13}$$

Analogous to (2), in the dynamic model, if the state is equal to s, the ex ante value function is related to the choice specific value function by:

$$V_i(s) = E_{\epsilon_i} \max_{a_i} \left[V_i(a_i, s) + \epsilon_i(a_i)\right]. \tag{14}$$

That is, the utility maximizing action maximizes the sum of the choice specific value function plus the private information $\epsilon_i(a_i)$. As in the static model, the equilibrium probabilities and the choice specific value functions are relate through the following equation

$$\sigma_i(a_i|s) = \frac{\exp(V_i(a_i, s))}{\sum_{a_i'} \exp(V_i(a_i', s))}. \tag{15}$$

3.5 Constructive Proof of Identification

As in the static model, we prove the identification of our model constructively. Our strategy is to assume that the econometrician has knowledge of the population choice probabilities $\sigma_i(a_i|s)$. We then show that it is possible to uniquely recover $\Pi_i(a_i, a_{-i}, s)$ after making appropriate normalizations and checking a rank condition.

As in the static model, we begin by taking the log of both sides of (15). Straightforward algebra implies that

$$\log(\sigma_i(a_i = k|s)) - \log(\sigma_i(a_i = 0|s)) = V_i(a_i = k, s) - V_i(a_i = 0, s) \tag{16}$$

This equation demonstrates that it is possible to recover the choice specific value functions up to a first difference, if we know the population choice probabilities.

Next, it follows from (14) and the properties of the extreme value distribution that:

$$\begin{aligned} V_i(s) &= E_{\epsilon_i} \max_{a_i} V_i(a_i, s) + \epsilon_i(a_i) = \log \sum_{k=0}^{K} \exp(V_i(k, s)) \\ &= \log \sum_{k=0}^{K} \exp(V_i(k, s) - V_i(0, s)) + V_i(0, s). \end{aligned} \tag{17}$$

We now combine (17) with equation (13) to yield:

$$\begin{aligned} V_i(0, s) &= \Pi_i(a_i = 0, s) + \beta E\left[V_i(s')|s, a_i = 0\right] \\ &= \Pi_i(a_i = 0, s) + \beta E\left[\log\left(\textstyle\sum_{k=0}^{K} \exp\left(V_i(k, s') - V_i(0, s')\right)\right) + V_i(0, s)|s, a_i = 0\right] \\ &= \Pi_i(a_i = 0, s) + \beta E\left[\log\left(\textstyle\sum_{k=0}^{K} \exp\left(V_i(k, s') - V_i(0, s')\right)\right)|s, a_i = 0\right] \\ &\qquad + \beta E\left[V_i(0, s')|s, a_i = 0\right] \end{aligned} \tag{18}$$

Next, suppose that we are willing to make the "outside good" assumption as in equation (5). Then equation (16) implies that:

$$V_i(0,s) = \beta E\left[\log\left(\sum_{k=0}^{K}\exp\left(V_i(k,s') - V_i(0,s')\right)\right) + V_i(0,s)|s, a_i = 0\right]$$
$$= \beta E\left[\log\left(\sum_{k=0}^{K}\exp\left(\log(\sigma_i(a_i = k|s') - \log(\sigma_i(a_i = 0|s'))\right)\right)|s, a_i = 0\right]$$
$$+ \beta E\left[V_i(0,s)|s, a_i = 0\right]$$

Since the population probabilities $\sigma_i(a_i = k|s)$ are assumed to be known for the purposes of our identification argument, the term

$$\beta E\left[\log\sum_{k=0}^{K}\exp\left(\log(\sigma_i(a_i = k|s')) - \log(\sigma_i(a_i = 0|s'))\right)|s, a_i = 0\right]$$

can be treated as a known constant. Then, equation (18) is a functional equation involving the unknown function $V_i(0,s)$. Blackwell's sufficient conditions imply that for fixed $\sigma_i(a_i|s)$, (18) is a contraction mapping and therefore there is a unique solution for $V_i(0,s)$. As a result, we have shown that $V_i(0,s)$ is identified. Moreover, $V_i(k,s)$ is identified for all k by substituting $V_i(0,s)$ into (16). Finally, we note that the ex ante value functions can be identified by (17) given that we have identified the $V_i(k,s)$.

Next, note that (13) implies that

$$\Pi_i(a_i = k, s) = V_i(a_i = k, s) - \beta E\left[V_i(s')|s, a_i = k\right]. \tag{19}$$

Our identification arguments imply that both terms on the right hand side of (19) are known. This implies that $\Pi_i(a_i = k, s)$ is identified. The rest of identification proof can then follow exactly as in equations (6)-(8). We simply need to construction the $\Pi_i(a_i, a_{-i}, s_i)$ from the static choice specific value functions $\Pi_i(a_i, s)$ by imposing exclusion restrictions.

Theorem 2 *Suppose that Assumptions 1-2 hold. Also suppose that for each s_i, there exist $(K+1)^{n-1}$ points in the support of the conditional distribution of s_{-i} given s_i. Assume that the $(K+1)^{n-1}$ equations defined by (8) are linearly independent. Then the latent utilities $\Pi_i(a_i, a_{-i}, s_i)$ are identified.*

4 Identification-based estimation procedures

In this section, we describe a set of nonparametric and semiparametric estimators for our dynamic game of incomplete information. We begin with the analysis of the two-step estimator, which generalizes the two-step estimation procedures that have been used in the literature for estimating the player payoffs in discrete dynamic games. The two-step estimator is constructed by using the empirical analogue of our identification strategy. The translation from identification arguments to the nonparametric estimator essentially only requires replacing the appropriate conditional expectations with analog sample projections. While there are many possible local and global nonparametric smoothing techniques to estimate conditional expectations, for the clarity of presentation we describe the nonparametric procedure using series expansions as in most of the recent literature (e.g. Newey (1994) and Chen, Linton, and Van Keilegom (2003)).

4.1 A two-step estimator

In the rest of paper we will maintain the assumption that one has access to a data set from a collection of independent markets $m = 1, \ldots, M$ with at least two periods of observations each. Players $i = 1, ..., n$ are observed to play the game in each of these markets. There are $t = 1, ..., T(m)$ plays of the game in market m. We use $a_{i,m,t}$ to denote agent i's actions in market m at time period t and $s_{i,m,t}$ the state variable. The set up of our estimator can be changed to allow for different data structures, such as different players in different markets or varying numbers of players across markets. However, this would come at the cost of greater notational complexity. Our goal will be to estimate $\Pi_i(a_i, a_{-i}, s)$, the nonparametric mean utility parameters at a single point s. While the nonparametric procedure we propose is extremely flexible, it suffers from the standard problem of the curse of dimensionality. However, this method is very useful to exposit how estimation can be constructed analogously to our nonparametric identification arguments. We shall propose a somewhat more practical semiparametric estimator in the next subsection. For expositional reasons, it is useful to exposit the nonparametric case first so that the principals of the semiparametric estimator will be clearer to the reader. The nonparametric estimator is implemented in four steps.

Step 1: Estimate $\widehat{V}_i(k,s)-\widehat{V}_i(0,s)$ using (16) Suppose that we "flexibly" construct an estimator $\widehat{\sigma}_i(a_i|s)$ of the equilibrium choice probabilities $\sigma_i(a_i|s)$. Then equation (16) shows that we can estimate $V_i(k,s)-V_i(0,s)$ as

$$\hat{V}_i(k,s)-\hat{V}_i(0,s)=\log(\widehat{\sigma}_i(k|s))-\log(\widehat{\sigma}_i(0|s)).$$

One method for estimating the choice probabilities flexibly is by using a "sieve logit". Let $q_l(s)$ for $l=1,2,...$ denote a sequence of known basis functions that can approximate any square-integrable function of the state variable s arbitrarily well. It is well known in series estimation that the number of basis terms must go to infinity at a rate that is appropriately slower than the sample size, otherwise the estimator will be inconsistent. How to choose the number of basis terms depends on the data configuration. If the number of observations $T(m)$ increases to infinity at least as fast as the number of total markets M, then the number of basis terms can be a function k of the number of obervations $T(m)$ in market m. In this case the model can be estimated market by market to allow for substantial unobserved heterogeneity across markets. On the other hand, if $T(m)$ is bounded from above, then the number of basis terms should depend on the total number of markets M.

To simplify notation in the following we shall focus on the case when $T(m)$ is finite and denote the number of basis terms as $k(M)$. Our asymptotic theory section shall derive how k depends on the sample size. We will denote the column vector of $k(M)$ basis terms by

$$q^{k(M)}(s)=\left(q_1(s),...,q_{k(M)}(s)\right)'.$$

The sieve logit estimator is simply the standard multinomial logit where the independent variables are $q^{k(M)}(s)$. We will estimate the choice probabilities $\widehat{\sigma}_i(k|s)$ separately for each agent. Obviously, pooling observations across agents is possible if we are willing to assume that agents will play the same strategies if they have the same state variables. We opt for a specification in which strategies vary across agents since this approach is more general. We will let $\gamma_{i,k}$ denote the parameters for agent i for a particular choice $k=0,1,...,K$ and $\gamma_i=(\gamma_{i,0},\gamma_{i,1},...,\gamma_{i,K})$ a vector which collects all the $\gamma_{i,k}$.We let $\gamma_{i,k}^{k(M)}$ denote the first $k(M)$ parameters corresponding to the basis vector $q^{k(M)}(s)'$. We estimate our model parameters as

$$\widehat{\gamma}_i = \arg\max_{\gamma_i} \sum_{m=1}^{M} \sum_{k=0}^{K} \sum_{t=1}^{T(m)} 1(a_{i,m,t} = k) \log \frac{\exp(q^{k(M)}\left(s_{i,m,t}\right)' \gamma_{i,k}^{k(M)})}{\sum_{k'=0}^{K} \exp(q^{k(M)}\left(s_{i,m,t}\right)' \gamma_{i,k'}^{k(M)})}.$$

Our estimate of $\widehat{\sigma}_i(k|s)$ and $\widehat{V}_i(k,s) - \widehat{V}_i(0,s)$ are then

$$\widehat{\sigma}_i(k|s) = \frac{\exp(q^{k(M)}\left(s_{i,m,t}\right)' \widehat{\gamma}_{i,k}^{k(M)})}{\sum_{k'=0}^{K} \exp(q^{k(M)}\left(s_{i,m,t}\right)' \widehat{\gamma}_{i,k'}^{k(M)})},$$

$$\widehat{V}_i(k,s) - \widehat{V}_i(0,s) = \log\left(\widehat{\sigma}_i(k|s)\right) - \log\left(\widehat{\sigma}_i(0|s)\right).$$

We also need to construct an estimate of $g(s'|s, a_i, a_{-i})$. The details of estimating $g(s'|s, a_i, a_{-i})$ will vary with the application. In many problems, the law of motion for the state variable is deterministic and therefore does not need to be directly estimated. Another common case is when $g(s'|s, a_i, a_{-i})$ is defined by a density. Let $g(s'|s, a_i, a_{-i}, \alpha)$ be a flexible parametric density with parameter α. In this case, one could use maximum likelihood or other appropriate methods to form an estimate $\widehat{\alpha}$ of α. Making a parametric distributional assumption about g is used here only for expositional convenience. One can also nonparametrically estimate g, in which case we only need to make a slight modification to the estimator.

Step 2: Estimate the choice specific value function for k=0, $\hat{V}_i(0,s)$. Step 1 only identifies the choice specific value functions up to a first difference. As in our identification arguments, we next construct an estimate of $\hat{V}_i(0,s)$ by iterating on the empirical analogue of equation (18). In order to do this, we first need to construct an estimate the density of next periods state s' given that the current periods state is s and the action chosen by player i is $a_i = 0$. We will denote this density as $\widehat{g}(s'|s, a_i = 0)$. Using the results from step 1, we can construct this density as:

$$\widehat{g}(s'|s, a_i = 0) = \sum_{a_{-i}} g(s'|s, a_i = 0, a_{-i}, \widehat{\alpha}) \widehat{\sigma}_{-i}(k|s)$$

The empirical analogue of (18) is then

$$\widehat{V}_i(0,s) = \beta \int_{ds'} \log\left(\sum_{k=0}^{K} \exp(\widehat{V}_i(k,s') - \widehat{V}_i(0,s'))\right) \widehat{g}(s'|s, a_i = 0) + \beta \int_{ds'} \widehat{V}_i(0,s)\widehat{g}(s'|s, a_i = 0).$$

The term $\beta \int \log\left(\sum_{k=0}^{K} \exp(\widehat{V}_i(k,s') - \widehat{V}_i(0,s'))\right) \widehat{g}(s'|s, a_i = 0)ds'$ is known from the previous step. In practice, we imagine computing this intergral using standard methods for numerical integration. Given that we know this term, the above equation can be viewed as a functional equation in $\widehat{V}_i(0,s)$. Define the operation T by:

$$T \circ \widehat{V}_i(0,s) = \beta \int_{ds'} \log\left(\sum_{k=0}^{K} \exp(\widehat{V}_i(k,s') - \widehat{V}_i(0,s'))\right) \widehat{g}(s'|s, a_i = 0) + \beta \int_{ds'} \widehat{V}_i(0,s)\widehat{g}(s'|s, a_i = 0) \tag{20}$$

As in our identification section, it is easy to verify that (20) satisfies Blackwell's sufficient conditions for a contraction and therefore has a unique fixed point. There is a large literature on solving functional equations defined by contraction mappings and in applied work we imagine using standard numerical methods to solve for $\widehat{V}_i(0,s)$.

An alternative method based on series expansion can also be used to estimate $V_i\,(0,s)$ without the need of explicitly estimating $\hat{g}\,(s'|s, a_i = 0)$ nonparametrically and calculating the fixed point to the above contraction mapping. Consider a linear series approximation of the value function $V_i\,(0,s)$: $V_i\,(0,s) = q^{k(M)}(s)'\,\theta_i^{k(M)}$. By the law of iterated expectation we can write

$$\begin{aligned} E\left\{1\,(a_i = 0)\,q^{k(M)}(s)q^{k(M)}(s)'\,\theta_i^{k(M)}\right\} =& \beta E\left\{1\,(a_i = 0)\,q^{k(M)}(s)\log\left(\sum_{k=0}^{K} \exp(\widehat{V}_i(k,s') - \widehat{V}_i(0,s'))\right)\right\} \\ &+ \beta E\left\{1\,(a_i = 0)\,q^{k(M)}(s)q^{k(M)}(s')'\,\theta_i^{k(M)}\right\}. \end{aligned}$$

Therefore $\hat{\theta}_i^{k(M)}$ can be estimated by an empirical analog: $\hat{\theta}_i^{k(M)} = \left(\hat{X} - \beta\hat{Z}\right)^{-1}\hat{Y}$, where

$$\hat{X} = \sum_{m=1}^{M}\sum_{t=1}^{T} 1\,(a_{i,m,t} = 0)\,q^{k(M)}(s_{m,t})q^{k(M)}(s_{m,t})',$$

$$\hat{Z} = \sum_{m=1}^{M}\sum_{t=1}^{T-1} 1\,(a_{i,m,t} = 0)\,q^{k(M)}(s_{m,t})q^{k(M)}(s_{m,t+1})'.$$

and

$$\hat{Y} = \sum_{m=1}^{M}\sum_{t=1}^{T-1} 1\,(a_{i,m,t} = 0)\,q^{k(M)}(s_{m,t})\log\left(\sum_{k=0}^{K} \exp(\widehat{V}_i(k, s_{m,t+1}) - \widehat{V}_i(0, s_{m,t+1}))\right).$$

The baseline choice specific value function is then estimated by the "fitted value" from the "linear regression": $\hat{V}_i\,(0,s) = q^{k(M)}(s)\hat{\theta}_i^{k(M)}$.

Step 3: Estimate the static choice specific payoff function $\hat{\Pi}_i(k,s)$ We evaluate the empirical analogue of (19) to estimate the static choice specific payoff function which we denote as $\hat{\Pi}_i(k,s)$. From the previous step, we have constructed an estimate of $\hat{V}_i(0,s)$ and from step 1 we have constructed an estimate of $\widehat{V}_i(k,s) - \widehat{V}_i(0,s)$. Putting these two steps together implies that we have an estimate of $\widehat{V}_i(k,s)$ for all i,k,s.

The empirical analogue of equation (19) is then

$$\widehat{\Pi}_i(a_i = k, s) = \widehat{V}_i(a_i = k, s) - \beta \int \widehat{V}_i(s')\widehat{g}(s'|s, a_i = k)ds'. \tag{21}$$

As a practical matter, in order to evaluate the above expression, it is useful to use the empirical analogue of equation (17), that is,

$$\widehat{V}_i(s) = \log \sum_{k=0}^{K} \exp(\widehat{V}_i(k,s))) \tag{22}$$

Substituting (22) into (21) yields:

$$\widehat{\Pi}_i(a_i = k, s) = \widehat{V}_i(a_i = k, s) - \beta \int \left(\log \sum_{k=0}^{K} \exp(\widehat{V}_i(k,s))) \right) \widehat{g}(s'|s, a_i = k)ds'. \tag{23}$$

As in the previous steps, $\widehat{\Pi}_i(a_i = k, s)$ can be evaluated using standard methods from numerical integration.

Step 4: Estimate the nonparametric mean utilities $\widehat{\Pi}_i(a_i, a_{-i}, s_i)$. The final step of our analysis is to perform the empirical analogue of inverting the linear system (8). Recall that in order to identify the system we needed to make an exclusion restriction. That is, the state has to be partitioned as $s = (s_i, s_{-i})$ and the variables s_{-i} are assumed not to enter into i's mean utilities. This allows us to write i's utility as $\Pi_i(a_i, a_{-i}, s_i)$.

One approach to inverting this system will be to run a local linear regression (see Fan and Gijbels (1992)). Local linear regression is essentially a weighted least squares regressions where the weights are defined using a kernel distance between the observations. Formally, our estimator for $\Pi_i(a_i, a_{-i}, s_i)$ is defined as the solution to the following minimization problem:

$$\widehat{\Pi}_i(a_i, a_{-i}, s_i) = \underset{\Pi_i(a_i,a_{-i},s_i)}{\arg\min} \sum_{m=1}^{M} \sum_{t=1}^{T(m)} (\hat{\Pi}_i(a_i, s_{m,t}) - \sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s_{m,t})\Pi_i(a_i, a_{-i}, s_i))^2 w(m,t), \tag{24}$$
$$w(m,t) = K(\frac{s_{imt} - s_i}{h}).$$

In the above, s_{mt} is the state variable in market m at time t, and s_{imt} is the component of s_{mt} that enters i's mean utilities. The term $K(\frac{s_{imt}-s_i}{h})$ is the distance, as measured by the kernel function K, between s_{imt} and s_i. Our weighting scheme overweighs observations near s_i and underweighs observations that are farther away. The term h is the bandwidth. The minimization problem (24) can be interpreted as a regression in which the static choice specific value function, $\hat{\Pi}_i(a_i, s_{m,t})$, is the dependent variable and $\hat{\sigma}_{-i}(a_{-i}|s_{m,t})$ are the regressors. The regression coefficients are $\Pi_i(a_i, a_{-i}, s_i)$, the structural mean utility parameters. The exclusion restrictions guarantee that standard the rank condition from the theory of regression is satisfied.

Choosing the bandwidth involves a variance-bias trade off. A smaller h reduces the bias by increasing the weight on nearby observations, but increases the variance of our estimator. In practice, cross validation, rules of thumb or simply "eyeballing" the bandwidth are commonly used in applied work. The theory of local linear regression establishes that if we shrink the bandwidth h at an appropriate rate, we will have a consistent estimate of $\Pi_i(a_i, a_{-i}, s_i)$.

4.2 Semiparametric payoff function

While the nonparametric estimator in the previous section is very flexible, it is not very practical for samples of small and intermediate sizes. Without a sufficiently large sample, nonparametric estimators suffer from a curse of dimensionality and may be poorly estimated. Also, the final estimates may be quite sensitive to ad hoc assumptions about the bandwidth or choice of the kernel. Therefore, it is desirable to have a semiparametric approach to the problem. We will specify $\Pi_i(a, s_i, \theta)$ to depend on a finite number of parameters. Parametric specifications are almost universal in the empirical literature. Frequently, applied researchers will assume that utility is linear in the strucutral parameters.

In what follows, we shall assume that the mean utilities take the form

$$\Pi_i(a, s_i) = \Phi_i(a, s_i)'\theta_{i,a}. \tag{25}$$

Here, $\Phi_i(a, s_i)$ is a known vector valued function and θ is used to weight the elements of the basis function.

In our semiparametric model, steps 1-3 of the nonparametric section are left unchanged. In step one, we estimate the choice probabilities $\widehat{\sigma}_i(k|s)$ flexibly using a sieve multinomial logit. We then apply the Hotz-Miller inversion to learn $\widehat{V}_i(k, s) - \widehat{V}_i(0, s)$. Steps 2 and 3 allow us to estimate $\hat{\Pi}_i(a_i, s_{mt})$, the static choice specific value function given that the action is a_i and the state is s_{mt}. Note that all of these steps are nonparametric and do not impose ad hoc functional form restrictions.

In our semiparametric estimator, we simply modify (24) in step 4 to include the parametric restrictions in (25).

$$\widehat{\theta}_{i,a} = \arg \min_{\Pi_i(a_i,a_{-i},s_i)} \sum_{m=1}^{M} \sum_{t=1}^{T(m)} (\hat{\Pi}_i(a_i, s_{m,t}) - \sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s_{m,t})\Phi_i(a, s_i)'\theta_{i,a})^2$$

An advantage of the semiparametric estimator is that it can be shown that $\widehat{\theta}$ converges to the true parameter value at a rate proportional to the square root of the sample size and has a normal asymptotic distribution. This is a common result in semiparametric estimation. Even though the nonparametric part of our model, $\hat{\sigma}_{-i}(a_{-i}|s_{m,t})$ and $\hat{\Pi}_i(a_i, s_{m,t})$, the payoff parameters θ converge at the slower rates.

4.3 Limit distribution of the semiparametric estimator

Given the ease of the estimation procedure discussed in the previous sections, we suggest that the most practical method of inference is bootstrapping. However, understanding the derivation of the asymptotic distribution is still important for ensuring the theoretical validity of resampling method such as bootstrap or subsampling. This section gives high level conditions and describes the form of the asymptotic variance using a powerful set of results developed in Newey (1994).

We will not study functional dependence among different $\theta_{i,a}$ components of the parameter vector θ. To simplify notation in the following we will use θ to denote a particular component $\theta_{i,a}$ of equation (25) in section 4.2, where $a_i = k$. In particular, θ (short-handed for $\theta_{i,a}$ where $a_i = k$) is identified by the relation $\sigma_i(k|s)\Pi_i(k,s) = E\left[d_i^k \Phi(a, s_i)|s\right]' \theta$, where

$$\Pi_i(k,s) = V_i(k,s) + \beta E\left[\log \sigma_i\left(0|s'\right)|s, a_i = k\right] - \beta E\left[V_i\left(0,s'\right)|s, a_i = k\right].$$

This suggests that a class of instrument variable estimators for θ, which includes the least square estimator in section 4.2 as a special case, takes the form of an empirical analog of

$$Ed_i^k A(s)\Phi(a,s_i)'\theta = Ed_i^k A(s)V_i(k,s) + \beta Ed_i^k A(s)\log\sigma_i\left(0|s'\right) - \beta Ed_i^k A(s)V_i\left(0,s'\right).$$

Then we can write $\hat{\theta} = \hat{X}^{-1}\hat{Y}$, where ($\sum_\tau$ is used to denote $\sum_{t=1}^{T-1}\sum_{m=1}^{M}$),

$$\hat{X} = \sum_\tau d_{i\tau}^k A(s_\tau)\Phi(a_\tau, s_{i\tau})/M(T-1),$$

and

$$\hat{Y} = \frac{1}{M(T-1)}\sum_\tau d_{i\tau}^k A(s_\tau)\left[\log\frac{\hat{\sigma}_i(k|s_\tau)}{\hat{\sigma}_i(0|s_\tau)} + \beta\log\hat{\sigma}_i(0|s_{m,t+1}) + \hat{V}_i(0,s_\tau) - \beta\hat{V}_i(0,s_{m,t+1})\right].$$

The instrument matrix itself can be estimated nonparametrically as $\hat{A}(s)$. The least square estimator in section 4.2 effectively uses $\hat{E}[\phi(a,s_i)|s,a_i=k]$ as the instrument matrix $\hat{A}(s)$. It is well known, however, that estimation of the instruments has no effect on the asymptotic distribution under suitable regularity conditions. Therefore with no loss of generality we treat the instruments $A(s)$ as known in deriving the form of the asymptotic distribution. Furthermore, by a standard law of large number $\hat{X}$ can be replaced by its population limit

$$X = Ed_{i\tau}^k A(s_\tau)\Phi(a_\tau, s_{i\tau}).$$

Therefore the asymptotic distribution of $\hat{\theta}$ will solely be determined by the convergence of $\sqrt{M(T-1)}\left(\hat{Y} - Y\right)$, where

$$Y = \frac{1}{M(T-1)}\sum_\tau d_{i\tau}^k A(s_\tau)\left[\log\frac{\sigma_i(k|s_\tau)}{\sigma_i(0|s_\tau)} + \beta\log\sigma_i(0|s_{m,t+1}) + V_i(0,s_\tau) - \beta V_i(0,s_{m,t+1})\right].$$

This asymptotic distribution is given in the following theorem.

Theorem 3 *Under suitable regularity conditions,*

$$\sqrt{M(T-1)}\left(\hat{Y}-Y\right)=\frac{1}{\sqrt{M(T-1)}}\sum_{\tau}\sum_{j=1}^{4}\phi_j\left(a_{i\tau},s_\tau\right)+o_p\left(1\right).$$

where

$$\begin{aligned}
\phi_1\left(a_{i\tau},s_\tau\right)= & \quad A\left(s_\tau\right)\left(d_{i\tau}^k-\sigma_i\left(k|s_\tau\right)\right)-\tfrac{\sigma_i(k|s_\tau)}{\sigma_i(0|s_\tau)}A\left(s_\tau\right)\left(d_{i\tau}^0-\sigma_i\left(0|s_\tau\right)\right),\\
\phi_2\left(a_{i\tau},s_\tau\right)= & \quad E\left[d_{i(m,t-1)}^k A\left(s_{m,t-1}\right)|s_\tau\right]\tfrac{1}{\sigma_i(0|s_\tau)}\left(d_i^0-\sigma_i\left(0|s_\tau\right)\right),\\
\phi_3\left(a_{i\tau},s_\tau\right)= & \quad \beta\delta\left(s_\tau\right)\Big[-\left(d_{i\tau}^0\log\sigma_i\left(0|s_{m,t+1}\right)-E\left(d_{i\tau}^0\log\sigma_i\left(0|s_{m,t+1}\right)|s_{m,t}\right)\right)\\
& +\left(d_{i\tau}^0V_i\left(0,s_{m,t+1}\right)-E\left(d_{i\tau}^0V_i\left(0,s_{m,t+1}\right)|s_{m,t}\right)\right)-\tfrac{1}{\beta}V_i\left(0|s_{m,t}\right)\left(d_{it}^0-\sigma_i\left(0|s_{m,t}\right)\right)\\
& +\tfrac{d_{i\tau}^0d_{i,m,t+1}^0}{\sigma_i(0|s_{m,t+1})}-E\left(\tfrac{d_{i\tau}^0d_{i,m,t+1}^0}{\sigma_i(0|s_{m,t+1})}|s_{m,t}\right)+d_{i\tau}^0-\sigma_i\left(0|s_\tau\right)\Big]
\end{aligned}$$

In the above, $\delta\left(s_\tau\right)$ is defined as the solution to the following functional relation:

$$\delta\left(s_t\right)\sigma_i\left(0|s_t\right)-\beta E\left(\delta\left(s_{t-1}\right)d_{i,t-1}^0|s_t\right)=\sigma_i\left(k|s_t\right)A\left(s_t\right).$$

$\phi_4\left(a_{i\tau},s_\tau\right)$ is defined similarly to $\phi_3\left(a_{i\tau},s_\tau\right)$, with $\beta\delta\left(s_\tau\right)$ replaced by $-\beta^2\bar{\delta}\left(s_\tau\right)$, where $\bar{\delta}\left(s_\tau\right)$ is now defined as the solution to the functional relation of

$$\delta\left(s_t\right)\sigma_i\left(0|s_t\right)-\beta E\left(\delta\left(s_{t-1}\right)d_{i,t-1}^0|s_t\right)=E\left(d_{i,t-1}^kA\left(s_{t-1}\right)|s_t\right).$$

An immediate consequence of theorem 3 is the asymptotic normality of the semiparametric estimator $\hat{\theta}$ under suitable regularity conditions

$$\sqrt{M(T-1)}\left(\hat{\theta}-\theta\right)\xrightarrow{d}N\left(0,X^{-1}Var\left(\sum_{j=1}^{4}\phi_j\left(a_{i\tau},s_\tau\right)\right)X^{-1}\right).$$

A few remarks are in order. First, because the second stage moment equations are exact identities when evaluated at the true parameter values, if the first stage nonparametric functions are exactly known and do not need to be estimated, the second stage parameters will have zero asymptotic variance. In other words, all the variations in the asymptotic variance are generated from the first stage nonparametric estimation of the conditional choice probabilities and the transition probabilities. Secondly, as shown in Newey (1994), the form of the asymptotic variance given in theorem 3 is obtained from a pathwise derivative calculation and does not depend on the exact form of the nonparametric methods that are being used to estimated, as long as suitable regularity conditions are met. Therefore, Theorem 3 is stated without regard to the particular nonparametric method and its regularity conditions used in the first stage analysis. Both of these will be discussed in the following two sections. While resampling method is a clear preferable method of inference, it is in principle possible to estimate nonparametrically each component of the asymptotic variance in theorem 3. Alternatively, Ackerberg, Chen, and Hahn (2009) proposes a recent approach that can also be used to estimate the asymptotic variance consistently.

4.4 Unobserved heterogeneity

Unobserved heterogeneity is an important concern for dynamic discrete choice models. A recent insight from this literature is that it is sufficient to estimate a reduced form model of conditional choice probabilities and transition probabilities that account for the presence of the unobserved heterogeneity. A variety of such methods are available in the recent literature, some allowing for a fixed number of support points in the distribution of unobserved state variables while others allowing for a continuous of unobserved state variables. For each of the discrete and continuous support cases of the unobserved state variables, some methods limited to only non time varying unobserved state variables while other methods might allow for serially correlated unobserved state variables.

In the following, we will take as given the ability of estimate a first stage model of conditional choice probabilities and conditional transition probabilities that incorporate the presence of general (discrete and continuous, time invariant and serially correlated) state variables. Therefore, we will assume that it is possible to use one of the methods available in the existing literature to estimate a reduced form model of $\hat{\sigma}_i(k|s), \forall i, k$ and $\hat{g}(s'|s, a)$, where now s' and s include both observed and unobserved state variables that can be either discrete or continuous, either time-invariant or serially correlated.

We now note that the entire nonparametric identification process in section 3.4 and the entire estimation procedure, both nonparametric and semiparametric, described in section 4, depend only on the first stage $\hat{\sigma}_i(k|s), \forall i, k$ and $\hat{g}(s'|s, a_i)$. Therefore, as long as the state transition process is assumed to be common across individuals, we can follow exactly the same procedures outlined in sections 3.4 and 4 to estimate the primitive mean utility functions $\Pi_i(a, s_i)$ and $\Phi_i(a, s_i)'\theta$. Perhaps the best way to understand this argument is through simulations. Given knowledge of $\hat{\sigma}_i(k|s), \forall i, k$ and $\hat{g}(s'|s, a_i)$, a research can generate a data set with as many markets and as many time periods as desired, and apply the estimation procedures described in the previous subsections of section 4 to the simulated data set.

5 Semiparametric efficient estimation

In the previous section we described a multi-stage semiparametric procedure which allows us to estimate the finite-dimensional parameters of the profit function. This procedure is very intuitive because it follows directly from the identification argument. The asymptotic distribution of this estimator also has an explicit analytic structure. However, this approach inherits the disadvantages of many multi-stage estimation techniques. First of all, the standard errors are hard to compute because of propagation of errors from the previous steps of the procedure which will depend on the degree of smoothness of the unknown functions of the model. Second, this multistage estimation procedure is not semiparametrically efficient. it is well known that it is difficult to design multistage estimation procedures that can achieve semiparametric efficiency bounds, because at each subsequent step has to compensate the estimation errors that will arise from previous estimation errors.

In this section we will propose an efficient one step estimation procedure using the framework of conditional moment models. It has the advantage that given the choice of instrument functions and the weighting matrix, practical inference can be performed using standard parametric methods as if a finite dimensional linear parametric model of $V_i(k, s)$ and $\Pi(a, s_i)'\theta$ is estimated, as long as the estimation noise in the estimation of $\sigma_i(k|s)$ is appropriately accounted for.

By formulating the model in a conditional moment framework and making use of the stationary controlled Markov process structure, we can avoid direct estimation of the transition density of the state variable. This simplifies the derivation of the semiparametric efficiency bound of the model and the statement of the regularity conditions for the semiparametric efficient estimator.

The conditional moment formulation is derived from the Bellman equations for individual players. Recall the Bellman equations of interest:

$$V_i(k,s) = \Pi_i(k,s;\gamma) + \beta \int \sum_{a_{-i}\in A_{-i}} \sigma_{-i}(a_{-i}|s) \log\left[\sum_{l=0}^{K} \exp(V_i(l,s'))\right] g(s'|s, a_i = k, a_{-i})\, ds',$$

where

$$\sigma_i(a_i = k|s) = \frac{\exp(V_i(k,s))}{\sum_{l=0}^{K} \exp(V_i(l,s))},$$

for $i = 1,\dots,n$ and $k = 0,\dots,K$. Denote $d^{i,l}$ the dummy for choice l by player i. We can use the second equation to substitute it into the first one, which leads to $n \times K$ conditional moment equations for each $(T-1)\times M$ observations:

$$\begin{aligned} E\Big[& d_{m,t}^{a_i,k}\left(V_i(0,s_{m,t}) - \beta V_i(0,s_{m,t+1}) + \beta \log \sigma_i(0|s_{m,t+1})\right) \\ & -d_{m,t}^{a_i,k}\left(1 - d_{m,t}^{a_i,0}\right)\left[\Pi_i(a_i,a_{-i},s_{m,t};\gamma) + \log\,\sigma_i(0|s_{m,t}) - \log\,\sigma_i(a_i|s_{m,t})\right] \Big|\, s_{m,t}\Big] = 0. \end{aligned} \tag{26}$$

Together with the following $n \times K$ moment conditions for each $T \times M$ observations,

$$E\left(d_{m,t}^{a_i,k}|s_{m,t}\right) = \sigma_i(k|s_{m,t}), \tag{27}$$

(26) and (27) form a system of conditional moment restrictions that fully characterize the implications from the structural dynamic discrete choice model. This system of conditional moment restrictions can be used to obtain asymptotically normal semiparametric estimators that can achieve the semiparametric efficiency bound by adapting the recipe prescribed in Ai and Chen (2003). In their notation of $E[\rho(w_{m,t},\gamma,V(\cdot),\sigma(\cdot))|s_{m,t}] = 0$, where $w_{m,t}$ are all the random variables the model, γ are the finite dimensional parameters, $V(\cdot)$ and $\sigma(\cdot)$ are the infinite dimensional unknown parameters, we can write, for $h(\cdot) = (V(\cdot),\sigma(\cdot))$:

$$\rho(w_{m,t},\gamma,h(\cdot)) = \left(\rho_1(w_{m,t},\gamma,V(\cdot),\sigma(\cdot))', \rho_2(w_{m,t},\gamma,V(\cdot),\sigma(\cdot))'\right)',$$

where ρ_1 is the $T \times m \times n \times K$ dimensional collection of $d_{m,t}^{a_i,k} - \sigma_i(k|s_{m,t})$, and

$$\begin{aligned} \rho_2(w_{m,t},\gamma,V(\cdot),\sigma(\cdot)) = {} & d_{m,t}^{a_i,k}\left(V_i(0,s_{m,t}) - \beta V_i(0,s_{m,t+1}) + \beta\log\sigma_i(0|s_{m,t+1})\right) \\ & -d_{m,t}^{a_i,k}\left(1 - d_{m,t}^{a_i,0}\right)\left[\Pi_i(a_i,a_{-i},s_{m,t};\gamma) + \log\,\sigma_i(0|s_{m,t}) - \log\,\sigma_i(a_i|s_{m,t})\right]. \end{aligned}$$

The conditional moment restrictions in (26) and (27) can be transformed into unconditional moments by forming an instrument matrix $z_{m,t}$ using the state variables $s_{m,t}$, its lags $s_{m,t-\tau}$ and polynomial powers $s_{m,t}$ and its lags, such that the number of instruments in $z_{m,t}$ increases at appropriate rates as the sample size increases to infinity. Equations (26) and (27) implies the following moment vectors with elements

$$E\Big[d_{m,t}^{a_i,p} z_{m,t}\Big(V_i\left(0,s_{m,t}\right)-\beta V\left(0,s_{m,t+1}\right)+\log\tfrac{\sigma_i(a_i|s_{m,t})}{\sigma_i(0|s_{m,t})}-\beta\log\tfrac{\sigma_i(a_i|s_{m,t+1})}{\sigma_i(0|s_{m,t+1})}$$
$$-\left(1-d_{m,t}^{a_i,0}\right)\left[\Pi_i\left(a_i,a_{-i},s_{m,t};\gamma\right)-\beta\log\sigma_i\left(a_i|s_{m,t+1}\right)\right]$$
$$+d_t^{a_i,0}\beta\log\sigma_i\left(0|s_{m,t+1}\right)\Big)\Big]=0,$$

and $Ez_{m,t}\left(d_{m,t}^{a_i,p}-\sigma_i\left(p|s_{m,t}\right)\right)=0$. To estimate γ we can follow two steps.

Step 1

We approximate the conditional choice probabilities using orthogonal series:

$$\sigma_i\left(a_i=p\,|\,s\right)=q^{k_1(MT)\prime}\left(s\right)b_{i,p}^1+\Delta_{k_1(MT)},$$

and approximate the value function similarly

$$V_i\left(0,\ s\right)=q^{k_2(MT)\prime}\left(s\right)\ b_i^2+\Delta_{k_2(MT)},$$

where $\Delta_{k_1(M)}$ and $\Delta_{k_2(M)}$ are numerical approximation errors that decrease to zero as $k_1\left(M\right)$ and $k_2\left(M\right)$ increase to infinity with M at appropriate rates.

Step 2

Next we form an instrument $z_{m,t}$ by stacking an orthogonal series of functions of the state variables $s_{m,t}$, $\left(q_0\left(s_{m,t}\right),\ldots,q_{k_3(MT)}\left(s_{m,t}\right)\right)$. This produces an over-identified empirical moment vector with the elements, for $b=\left(b_1^{i,p},b_2^i,\forall i,p\right)$,

$$\widehat{\varphi}\left(\gamma,b\right)=\sum_{m,t}\varphi_{m,t}\left(\gamma,b\right)\quad\text{where}\quad\varphi_{m,t}\left(\gamma,b\right)=\rho\left(w_{m,t},\gamma,b\right)\otimes z_{m,t}.$$

Then we introduce a weighting matrix $\mathcal{W}$ with both row and column dimensions $\dim\left(z_{m,t}\right)\times\dim\left(\rho\right)$. In the simplest case we can use the identity matrix in lieu of $\mathcal{W}$. Using a given weighting matrix we form a GMM objective and minimize it with respect to parameters of interest γ as well as the parameters of the expansion of the value function

$\min_{\gamma, b} \widehat{\varphi}(\gamma, b)' \mathcal{W} \widehat{\varphi}(\gamma, b)$. In particular, if we let $Z \equiv (z_{m,t}, \forall, m, t)'$ denote the data matrix for the instruments, the following choice of the weighting matrix

$$\mathcal{W} = I \otimes (Z'Z)^{-1} \left(\sum_{m,t} \Omega_{m,t}^{-1} \otimes z_{m,t} z_{m,t}' \right) I \otimes (Z'Z)^{-1}.$$

yields the semiparametric minimum distance estimators of Ai and Chen (2003). In the above $\Omega_{m,t}$ is a candidate estimate of the conditional variance covariance matrix of $\rho(w_{m,t}, \gamma, h(\cdot))$ given $s_{m,t}$. When $\Omega_{m,t} \equiv I$ an identity matrix, the estimator becomes a nonlinear two stage least square estimator. When $\Omega_{m,t} = \Omega$ is homoscedastic across observations, this becomes a nonlinear three stage least square estimators. Semiparametric efficiency bound is achieved when $\Omega_{m,t}$ is a consistent estimate of $Var(\rho(w_{m,t}, \gamma, h(\cdot)) | s_{m,t})$, in which case it becomes a heteroscedasticity weighted nonlinear three stage least square estimator. When $\rho(\cdot)$ is a scalar, the semiparametric efficient minimum distance estimator is a weighted nonlinear two stage least square estimator.

Remark 1:

By appropriate choices of the instrument functions and the weighting matrix, the conditional moment framework also incorporates the multistage procedure in the previous section as special cases. If the same orthogonal series is used in approximating $V_i(0, s_{m,t})$, $\sigma_i(p|s_{m,t})$ and in obtaining the instruments, and if $k_1(MT) = k_2(MT) = k_3(MT)$, the instrumented moment conditions (27) are exactly identifying, and $\sigma_i(p|s)$ are computed from least square regressions:

$$\widehat{\sigma}_i(a_i = p \,|\, s) = q^{k(M)\prime}(s) \left(\sum_{m,t} q^{k(M)}(s_{m,t}) \, q^{k(M)\prime}(s_{m,t}) \right)^{-1} \sum_{m,t} q^{k(M)}(s_{m,t}) \, d_t^{a_i,p}.$$

Given the estimate of $\widehat{\sigma}_i(a_i = p \,|\, s)$, the component of the instrumented moment condition $\rho_1(w_{m,t}, \cdot)$ that corresponds to $k = 0$ is also exactly identifying and depends only on $V_i(0, s)$. Hence $V_i(0, s)$ can be estimated by a single equation two stage least square regression with dependent variables $\beta \log \widehat{\sigma}_i(0|s_{m,t+1})$, independent variables $q^{k(M)}(s_{m,t}) - \beta q^{k(M)}(s_{m,t+1})$ and instrument matrix $q^{k_{MT}}(s_{m,t})$. Subsequently, given estimates of $\widehat{V}_i(0, s)$ and $\widehat{\sigma}_i(a_i = p \,|\, s)$, the parameters

$$\gamma_i^p = \left(\gamma_{i,a_{-i}}^p, \forall a_{-i} \right),$$

for $i = 1, \ldots, n$, $p = 1, \ldots, K$ in a linear profit function specification $\Pi_i(p, a_{-i}; \gamma) = \Phi_i(p, a_{-i})' \gamma_i^p$, can be estimated by single equation linear two stage least square regression methods when $\Omega(x_i) \equiv I$, with dependent variables

$$Y_{i,p,m,t} = \widehat{V}_i(0, s_{m,t}) - \beta \widehat{V}(0, s_{m,t+1}) + \log \frac{\widehat{\sigma}_i(a_i|s_{m,t})}{\widehat{\sigma}_i(0|s_{m,t})} - \beta \log \frac{1}{\widehat{\sigma}_i(0|s_{m,t+1})}$$

and the vector of independent variables $\mathbf{X}_t$ with elements

$$X_{i,p,m,t} = -\left(1 - d_{m,t}^{a_i,0}\right) \Phi_i(p, a_{-i}),$$

and instrument matrix $Z = (z_{m,t}, \forall m, t)'$. Efficiency can be improved by weighted 2SLS or weighted 3SLS by choosing $\hat{\Omega}(x_i)$ appropriately.

Remark 2:

The semiparametric efficient minimum distance estimator of Ai and Chen (2003) can be interpreted both in light of weighted nonlinear three stage least square estimator and the efficient instrument method of Newey (1990) for finite dimensional parameters. The semiparametric minimum distance objective function can be equivalently rewritten as

$$\sum_{m,t} \hat{\rho}(s_{m,t}, b, \gamma)' \hat{\Omega}_{m,t}^{-1} \hat{\rho}(s_{m,t}, b, \gamma),$$

where $\hat{\rho}(s_{m,t}, b, \gamma)$ is an estimate of $\hat{E}(\rho(w_{m,t}, b, \gamma) | s_{m,t})$,

$$\hat{\rho}(s, b, \gamma) = z\left(Z'Z\right)^{-1} \sum_{m,t} z_{m,t} \rho(w_{m,t}, b, \gamma)'.$$

Its first order condition resembles the efficient instrument estimator of Newey (1990):

$$\sum_{m,t} \frac{\partial}{\partial(b, \gamma)} \hat{\rho}(s_{m,t}, b, \gamma)' \hat{\Omega}_{m,t}^{-1} \hat{\rho}(s_{m,t}, b, \gamma).$$

The efficient instrument estimator of Newey (1990) only differs in using $\rho(s_{m,t}, b, \gamma)$ in place of the second $\hat{\rho}(s_{m,t}, b, \gamma)$ in light of the law of iterated expectation, and instead uses the first order condition of

$$\sum_{m,t} \frac{\partial}{\partial(b, \gamma)} \hat{\rho}(s_{m,t}, b, \gamma)' \hat{\Omega}_{m,t}^{-1} \rho(s_{m,t}, b, \gamma).$$

The following theorem adapts the semiparametric efficiency bound in Ai and Chen (2003) to our model. In our model the unknown function $V(\cdot)$ enters non-linearly as a function of the state variable in the period t and in the period $t+1$.

Using the results from Ai and Chen (2003) we can provide the semiparametric efficiency bound for estimating the parameter γ of the payoff function. Denote

$$\Sigma_0\left(s_{m,t}\right) = \mathrm{Var}\left(\rho\left(w_{m,t}, \gamma_0, h_0\left(\cdot\right)\right)|s_{m,t}\right).$$

The semiparametric efficiency bound expressed in theorem 4 will depend on the functional derivatives of the moment conditions ρ_1 in (26) and ρ_2 in (27) on the unknown functions $h_1^i\left(\cdot\right) = V_i\left(0,\cdot\right)$ and $h_2^{i,k}\left(\cdot\right) = \sigma_i\left(k|\cdot\right)$. The functional derivative of the conditional moment functions with respect to these unknown functions can be expressed using the linear expectation operator

$$\mathcal{P}_i^k \circ f = E\left[f\left(s_{m,t+1}\right) \mid s_{m,t} = s, a_{m,t}^i = k\right],$$

where expectation is defined for the conditional density

$$\sum_{a_{-i}} g\left(s_{m,t+1}|s_{m,t} = s, a_i = k, a_{-i}\right)\sigma_{-i}\left(a_{-i}|s_{m,t} = s\right).$$

The operator $\mathcal{P}_i^k \circ f$ is assumed to have a discrete spectrum with eigenfunctions $\left\{\Theta_j^{i,k}\left(s\right)\right\}_{j=0}^{\infty}$ and eigenvalues $\left\{\lambda_j^{i,k}\right\}_{j=0}^{\infty}$ different from zero. Then we can find that

$$\frac{d\,E\left[\rho_1^{i,k}\left(w_{m,t},\gamma_0,h_0\left(\cdot\right)\right)|s_{m,t}\right]}{d\,h_1^i}\left[\psi\right] = \sigma_i\left(k|s\right)\sum_{j=0}^{\infty}\psi_j\left(1-\beta\lambda_j^{i,k}\right)\Theta_j^{i,k}\left(s\right),$$

for all sequences of real numbers ψ which belong to $\mathcal{H} = \left\{\psi|\sum\limits_{j=0}^{\infty}|\psi_j|\,\left\|\Theta_j^{i,k}\left(s\right)\right\| < \infty\right\}$, Furthermore, we also calculate that

$$\frac{d\,E\left[\rho_1^{i,k}\left(w_{m,t},\gamma_0,h_0\left(\cdot\right)\right)|s_{m,t}\right]}{d\,h_2^{i,0}}\left[\psi\right] = \beta E\left[d_{m,t}^{a_i,k}\frac{1}{\sigma_i(0|s_{m,t+1})}h_2^{i,0}\left(s_{m,t+1}\right)|s_{m,t}\right].$$

and for $k \neq 0$ the linear derivative of,

$$\frac{d\,E\left[\rho_1^{i,k}\left(w_{m,t},\gamma_0,h_0\left(\cdot\right)\right)|s_{m,t}\right]}{d\,h_2^{i,k}}\left[\psi\right]=h_2^{i,k}\left(s_{m,t}\right).$$

Finally, for all k,

$$\frac{d\,E\left[\rho_2^{i,k}\left(w_{m,t},\gamma_0,h_0\left(\cdot\right)\right)|s_{m,t}\right]}{d\,h_2^{i,k}}\left[\psi\right]=-h_2^{i,k}\left(s_{m,t}\right).$$

The functional derivatives in the direction of the unknown functions $\frac{d\,E[\rho(w_{m,t},\gamma_0,h_0(\cdot))|s_{m,t}]}{d\,h}\left[\psi\right]$ are formed by stacking the above individual components together.

Then for each component of γ solve the minimization problem

$$\min_{\psi^{(j,0)}\in\mathcal{H}} E\left\{\left(\frac{d\,E[\rho(w_{m,t},\gamma_0,h_0(\cdot))|s_{m,t}]}{d\,\gamma_j}-\frac{d\,E[\rho(w_{m,t},\gamma_0,h_0(\cdot))|s_{m,t}]}{d\,h}\left[\psi^{(j,0)}\right]\right)\Sigma_0\left(s_{m,t}\right)^{-1}\right.$$
$$\left.\times\left(\frac{d\,E[\rho(w_{m,t},\gamma_0,h_0(\cdot))|s_{m,t}]}{d\,\gamma_j}-\frac{d\,E[\rho(w_{m,t},\gamma_0,h_0(\cdot))|s_{m,t}]}{d\,h}\left[\psi^{(j,0)}\right]\right)\right\}.$$

Form the vector

$$D_{\psi^{(0)}}\left(s_{m,t}\right)=\frac{d\,E\left[\rho\left(w_{m,t},\gamma_0,h_0\left(\cdot\right)\right)|s_{m,t}\right]}{d\,\gamma'}-\frac{d\,E\left[\rho\left(w_{m,t},\gamma_0,h_0\left(\cdot\right)\right)|s_{m,t}\right]}{d\,h}\left[\psi^{(0)}\right].$$

The following theorem follow directly from the result provided in Ai and Chen (2003):

Theorem 4 *The semiparametric efficiency bound for estimation of γ in equation (26) can be found as*

$$V\left(\gamma\right)=E\left[D_{\psi^{(0)}}\left(s_{m,t}\right)'\,\Sigma_0\left(s_{m,t}\right)^{-1}D_{\psi^{(0)}}\left(s_{m,t}\right)\right]^{-1}.$$

5.1 Asymptotic distribution for semiparametric estimator

We impose the following regularity assumptions on the functions in the model to assure that the two-stage conditional moment-based estimation method delivers consistent estimates for the Euclidean parameter in the per period payoff function as well as the non-parametric estimate of the continuation value of players.

Assumption 3

1. *Parameter space Γ is a convex compact set. Profit function $\Pi_i(a_i, a_{-i}, s; \gamma)$ is continuous in γ for each $(a_i, a_{-i}) \in \mathcal{A}$. Moreover, for each $\gamma \in \Gamma$ profit function is bounded:*

$$\sup_{a \in \mathcal{A},\, s \in \mathcal{S}} |\Pi_i(a_i, a_{-i}, s; \gamma)| < \infty.$$

2. *The data $\left\{\{a_{1t}, \ldots, a_{nt}, s_t, s_{t+1}\}_{t=1}^{T-1}\right\}_{m=1}^{M}$ are i.i.d. generated by the stationary distribution determined by Markov transition kernel for the state variable.*
3. *The approximating series expansion $\{q^{k(m)}\}$ forms a basis in $\mathcal{C}^{k(m)}(\mathcal{S})$, such that the eigenvalues of $E\left[q^{k(m)}(s_{t+1})\, q^{k(m)\prime}(s_{t+1})\,|s_t = s\right]$ are bounded away from zero for all $s \in \mathcal{S}$. The operator*

$$\mathcal{P}_i \circ f = E\left[f(s_{m,t+1}) \mid s_{m,t} = s,\, a_i\right],$$

where expectation is defined for the conditional density

$$\sum_{a_{-i}} g(s_{m,t+1}|s_{m,t} = s, a_i = k, a_{-i})\, \sigma_{-i}(a_{-i}|s_{m,t} = s),$$

which has a discrete spectrum with eigenfunctions $\left\{\Theta_j^{i,k}(s)\right\}_{j=0}^{\infty}$ such that for each j we can find $j' \leq j$ for which $\langle q_j^{k(m)},\, \Theta_j^{i,k}\rangle \neq 0$. In addition, $\limsup_{m\to\infty} E\left[m^{-1/2}\left(1 + \beta\Lambda_j^{i,k}\right) q_j^{k(m)}\right] < \infty$.

4. *The value function $V_i(s_t)$ is piece-wise continuous on $\mathcal{S}$ and bounded. Moreover, for each $V_i(\cdot) \in \mathcal{V}$ there exists a vector $\mu \in \mathbb{R}^{k(n)}$ such that $E\left[\left(V(s_t) - \mu'\, q^k\right)^2\right] = o(1)$.*
5. *For a given $V(\cdot)$ and transition density, there exists a unique solution $\gamma \in \Gamma$ to the system of equations*

$$E\left[\varphi_i(s_t, s_{t+1}, a; V_i, \gamma) \mid s_t\right] = 0,$$

for $i = 1, \ldots, n$.

These assumptions allow us to apply the results from Newey and Powell (2003) for each order of approximation $k(m)$. By the appropriate choice of basis we can guarantee that the approximation error is negligible as compared to the estimation error. The estimation problem is linear in parameters: expansion coefficients for $V(\cdot)$ and the Euclidean parameter

γ. For each finite approximation order $k(m)$ we can assure that the estimated parameters are consistent estimates for the functions given the order of approximation. When $m \to \infty$ approximation error approaches zero and the estimated coefficients will be consistent for the true coefficients. Given that by assumption the value function admits consecutive approximations in the basis $\{q^K(s)\}$ for each $K \in \mathbb{N}$, the fitted values $\widehat{b}^{K\prime} q^K(s)$ will be consistent for the true value function in the limit.

We can provide a similar set of assumptions that will assure the asymptotic normality of the estimates.

Assumption 4 *1. There exists a metric $\|\cdot\|_s$ such that the product space $\mathcal{V} \times \Gamma$ is compact. Moreover, the space $\{q^\infty(s)\} \times \Gamma$ is dense in $\mathcal{V} \times \Gamma$ for the chosen metric.*

2. For the covering number in the family of the moment functions defined by consecutive series approximations

$$\log N\left(\varepsilon,\, \{q^{k(m)}(s)\} \times \Gamma,\, \|\cdot\|_s\right) \leq Ck(m)\log\left(\frac{k(m)}{\epsilon}\right).$$

3. The weighting matrix can be estimated consistently such that

$$\left\|\widehat{\mathcal{A}}(s,\, d^a) - \mathcal{A}(s,\, d^a)\right\| = o_p\left(m^{-1/4}\right).$$

Moreover for each $\|\mu - \mu_0^{k(m)}\| < Cm^{-1/4}$ and each $\|\gamma - \gamma_0\| < Cm^{-1/4}$

$$\left\|\left(\widehat{\mathcal{A}}(s,\, d^a) - \mathcal{A}(s,\, d^a)\right)\varphi\left(s', s, a; \mu' q^{k(m)}(s), \gamma\right)\right\| = o_p\left(m^{-1/4}\right)$$

4. The variance of the moment function $Var\left(\varphi(s', s, a; V_0, \gamma_0) \,\middle|\, s\right)$ is positive definite for all $s \in \mathcal{S}$.

5. For each direction $h \in \mathcal{C}^{k(m)}(\mathcal{S})$ we define the directional derivative of the moment function as a vector $\partial_h \varphi = \left(\frac{\partial \varphi}{\partial \gamma}, \left(\frac{\partial \varphi}{\partial V}\right)_h\right)$, where

$$\frac{\partial \varphi_i}{\partial \gamma} = \frac{\partial \Pi_i(a_i, a_{-i}, s; \gamma)}{\partial \gamma}, \quad \text{and} \quad \left(\frac{\partial \varphi_i}{\partial V_i}\right)_h = \sum_{j=0}^{\infty} h_j \left(1 - \beta \lambda_j^{i,k}\right) \Theta_j^{i,k}(s).$$

We assume that in the ball of radius $Cm^{-1/4}$ around the true value (V_0, γ_0) in $\mathcal{V} \times \Gamma$ the directional derivative $\partial_h \varphi$ is Hölder-continuous with respect to norm $\|\cdot\|_s$ and bounded above by a linear functional of h, $F[h]$ such that $E[F[h]] < \infty$. Choose h^* such that

$$E\left\{\partial_h \varphi (V_0, \gamma_0)' E\left[\mathcal{A}(s_{t+1}, d^a)\, \mathcal{A}(s_{t+1}, d^a)' \,\middle|\, s_{t=s}\right] \partial_h \varphi (V_0, \gamma_0)\right\}$$

is minimized with respect to h. Then uniformly in the chosen ball

$$E\left(\|\partial_{h^*} \varphi (V_0, \gamma_0) - \partial_{h^*} \varphi (V, \gamma)\|^2\right) = o\left(m^{-1/4}\right),$$

where we use a standard Euclidean norm.

The following theorem is an immediate consequence of Ai and Chen (2003), which we state without proof.

Theorem 5 *Under assumptions 3 and 4, for $\hat{\gamma}$ defined in steps 1 and 2 of the previous section, $\hat{\gamma} \xrightarrow{p} \gamma_0$, and for $V(\gamma)$ given in theorem 4*

$$\sqrt{MT}\,(\hat{\gamma} - \gamma_0) \xrightarrow{d} N(0, V(\gamma)).$$

6 Non-parametric two-stage estimation

The moment equation (26) in general does not depend on the dimensionality of the payoff parameter γ. Making γ infinite-dimensional will cost us losing the parametric convergence rate. However, given the identification assumptions we will be able to provide a fully non-parametric estimate of the per-period payoff function. We can suggest an estimation procedure which is equivalent to the efficient estimation procedure in the semiparametric case.

Step 1 Estimate conditional choice probabilities non-parametrically using the orthogonal series representation:

$$\widehat{\sigma}_i (a_i = p \mid s) = q^{k(M)\prime}(s) \left(\sum_{m,t} q^{k(M)}(s_{m,t})\, q^{k(M)\prime}(s_{m,t})\right)^{-1} \sum_{m,t} q^{k(M)}(s_{m,t})\, d_t^{a_i,p}.$$

Step 2

Consider a series approximation for the value function

$$V_i\left(a_i = p,\, s\right) = q^{k(M)\prime}\left(s\right)\, b^{i,p} + \Delta_{k(M)},$$

where $\Delta_{k(M)}$ is a numerical approximation error, and a similar expansion for the payoff function

$$\Pi_i\left(a_i = p,\, a_{-i}\, s\right) = q^{k(M)\prime}\left(s\right)\, \gamma^{i,p,a_{-i}} + \Delta'_{k(M)},$$

For implementability of the procedure at this step we need the payoff function to be continuous (or, at least, has a finite set of points of first-order discontinuity). Next we form an instrument $z_{m,t}$ by stacking the state variables $s_{m,t}$ across the markets forming vectors s_t, and then choosing the linearly independent subset of vectors from the collection

$$\left(q_0\left(s_{m,t-\tau}\right), \ldots, q_{k(M)}\left(s_{m,t-\tau}\right)\right),$$

for all $0 \leq \tau \leq t-1$. Additional instruments come from other functions of $a_{m,t}$ and and the estimated choice probabilities $\widehat{\sigma}_i\left(j|s_{m,t+1}\right)$. This produces an empirical moment vector with $2k(M)$ unknown expansion coefficients with the elements

$$\begin{aligned}\widehat{\varphi}_{i,p}\left(\gamma,\, b\right) = \tfrac{1}{T}\sum_{t=1}^{T-1} d_t^{a_i,p} z_t \Bigg(& b^{i,p\prime}\left(q^{k(M)}\left(s_{m,t}\right) - \beta q^{k(M)}\left(s_{m,t+1}\right)\right) \\ & - \left(1 - d_t^{a_i,0}\right)\left[q^{k(M)\prime}\left(s\right)\, \gamma^{i,p,a_{-i}} - \beta \log \widehat{\sigma}_i\left(a_i|s_{m,t+1}\right)\right] \\ & + d_t^{a_i,0} \beta \log\left(1 - \sum_{j=1}^{K} \widehat{\sigma}_i\left(j|s_{m,t+1}\right)\right)\Bigg).\end{aligned}$$

Then we introduce a weighting matrix $\mathcal{W}$ with dimensions $n\,K\,m \dim\left(z_t\right) \times n\,K\,m \dim\left(z_t\right)$. In the simplest case we can use the identity matrix in lieu of $\mathcal{W}$. For this weighting matrix we form a GMM objective and minimize it with respect to parameters of interest γ as well as the parameters of the expansion of the value function

$$\min_{\gamma,\, b} \widehat{\varphi}\left(\gamma,\, b\right)' \, \mathcal{W}\, \widehat{\varphi}\left(\gamma,\, b\right).$$

In this estimation procedure the object of interest is the entire surface of the profit function, which can be computed as

$$\widehat{\Pi}_i\left(a_i = p,\, a_{-i}\, s\right) = q^{k(M)\prime}\left(s\right)\widehat{\gamma}^{i,p,a_{-i}}.$$

We need to determine the conditions that assure consistency and non-degeneracy of the asymptotic distribution of the pointwise estimate of the payoff function as well as find the rate of convergence of the estimator. Previous we imposed conditions that assure convergence of the semiparametric estimator. We can supplement them with additional assumptions which will provide consistency and asymptotic normality in the non-parametric case.

Assumption 5 *1. The payoff function $\Pi_i\left(a_i, a_{-i}, \cdot\right)$ belongs to the functional space$\mathcal{C}^p\left(S\right)$ for $p > 1$. Moreover, the orthocomplement of projecting the payoff function onto some Hilbert space $\mathcal{H}$, defined by the set of basis functions $\left\{q_t\left(\cdot\right)\right\}_{t=0}^{p}$ with the scalar product $\langle\cdot,\cdot\rangle$ has a norm in $\mathcal{C}^\infty\left(S\right)$ decreasing in p. Moreover its projection on the first p basis vectors converge absolutely, uniformly in the argument as $p \to \infty$.*

2. *For a truncation sequence $k(m) < m^{2r}$ the error of approximation of $\Pi_i(\cdot)$ and $V\left(\cdot\right)$by the basis function $\left\{q_t\left(\cdot\right)\right\}_{t=0}^{k(m)}$ is $o\left(m^{-2r}\right)$ with respect to the norm implied by the scalar product in $\mathcal{H}$.*

3. *$\widehat{\sigma}_i(a_i|a_i,\cdot)$ is asymptotically normal pointwise in Ω and converges at rate q. The truncation sequence $k(m)'$ giving the convergence rate m^q is $o\left(k(m)'\right)$. The approximation error of $h_i(\cdot)$ with respect to the norm in $\mathcal{H}$ is of order smaller than m^q*

This set of assumptions allows us to formulate the following theorem, which is proven in the appendix.

Theorem 6 *Given assumption 3, 4 and 5,*

$$m^{\min\{q,r\}}\left(\widehat{V}_i^{k(m)}\left(s\right) - V_i\left(s\right)\right) \xrightarrow{d} N\left(0,\, \omega_v^2\right),$$

and

$$m^{\min\{q,r\}}\left(\widehat{\Pi}_i^{k(m)}\left(a_i, a_{-i}, s\right) - \Pi_i\left(a_i, a_{-i}, s\right)\right) \xrightarrow{d} N\left(0,\, \omega_\pi^2\right),$$

7 Simulations

To demonstrate the performance of proposed estimators in finite samples, we conduct two sets of numerical simulations. The first set is a simple two by two entry game with discrete state variables and the second set is a single agent dynamic discrete choice model with continuous state variables.

In the first set of numerical simulations, each of the two players has one state variable that takes two possible values. Each player simultaneously decides whether to enter a market. The payoff to not entering into the market is normalized to zero regardless of the action of the competing player. The payoff for entering the market, which is also a function of the action of the competing player, is a draw from the uniform distribution between -2 and 2. The distributions of the payoff are independent across both the combination of the states and across the actions of the competing players. Therefore, we do not impose restrictions on how the action of the competing player affects the payoff to entering into the market, and allow the actions of both players to be either substitutes or complements. The transition probability matrices for a new state condition on the previous state and the actions of both players are also randomly generated from uniform distributions between 0 and 1. They are normalized so that the transition probability matrix is a proper stochastic matrix. The discount rate is set to 0.9.

Once generated, the payoff matrix and the transition probability matrices are held constant across the simulation runs. Following the recipe described in the estimation section, we first estimate the entry probabilities from independently generated data on the entry indicators, and then invert out the choice specific continuation value function and the choice specific static expected utility function. Finally, the primitive payoffs are recovered from the choice specific static expected utility functions.

The following tables report the results across 1000 simulation runs. The number of markets (nmarket), reported in the following tables refer to the number of observations (markets) generated for each combination of the state variables.

These tables show that the estimator performs well in finite sample, and that the amount of estimation error decreases monotonically as the sample size increases.

In the second set of numerical simulations for a single agent dynamic discrete choice

Table 1: Simulation summary for entry utilities, nmarket=100

i	a_{-i}	state	1st quartile	median	mean	3rd quartile	std	true Pi
1	1	1	-0.238	-0.006885	-0.0028	0.225	0.36	-0.52
1	1	2	-0.22	0.00004	0.0071	0.2304	0.343	0.749
1	2	1	-0.33	-0.011	-0.0349	0.28	0.48	-1.023
1	2	2	-0.21	0.0038	0.017	0.24	0.34	0.81
2	1	1	-0.25	0.027	0.023	0.31	0.44	0.53
2	1	2	-0.36	-0.021	-0.005	0.35	0.57	-1.005
2	2	1	-0.31	-0.022	0.0032	0.31	0.482	1.15
2	2	2	-0.38	0.013	-0.021	0.36	0.619	-1.600

Table 2: Simulation summary for entry utilities, nmarket=500

i	a_{-i}	state	1st quartile	median	mean	3rd quartile	std	true Pi
1	1	1	-0.109	-0.006	-0.002	0.102	0.161	-0.52
1	1	2	-0.095	0.0017	0.0011	0.097	0.146	0.749
1	2	1	-0.15	-0.0044	-0.0068	0.13	0.211	-1.02
1	2	2	-0.092	-0.0059	0.00004	0.1	0.146	0.812
2	2	1	-0.109	0.013	0.009	0.12	0.18	0.53
2	2	2	-0.164	0.0085	-0.0013	0.15	0.23	-1.005
2	2	1	-0.13	0.0017	0.0043	0.14	0.203	1.15
2	2	2	-0.15	0.0001	0.002	0.16	0.24	-1.60

Table 3: Simulation summary for entry utilities, nmarket=1000

i	a_{-i}	state	1st quartile	median	mean	3rd quartile	std	true Pi
1	1	1	-0.073	-0.00002	0.001	0.077	0.109	-0.52
1	1	2	-0.075	-0.0049	-0.0032	0.072	0.106	0.749
1	2	1	-0.107	-0.003	-0.005	0.092	0.14	-1.023
1	2	2	-0.066	0.0053	0.004	0.075	0.108	0.812
2	2	1	-0.078	0.0023	0.0046	0.086	0.127	0.537
2	2	2	-0.11	0.0024	-0.0033	0.107	0.166	-1.005
2	2	1	-0.098	-0.0021	-0.001	0.091	0.14	-1.60

Table 4: Simulation summary for entry utilities, nmarket=2000

i	a_{-i}	state	1st quartile	median	mean	3rd quartile	std	true Pi
1	1	1	-0.05	0.005	0.0038	0.059	0.0772	-0.52
1	1	2	-0.053	-0.00017	-0.00037	0.051	0.074	0.749
1	2	1	-0.078	-0.0066	-0.0062	0.0603	0.1007	-1.023
1	2	2	-0.045	0.0018	0.0017	0.05	0.075	0.812
2	2	1	-0.055	0.0017	0.0039	0.064	0.089	0.537
2	2	2	-0.088	-0.005	-0.0049	0.07	0.119	-1.005
2	2	1	-0.066	-0.004	-0.0021	0.059	0.097	1.150
2	2	2	-0.079	0.011	0.0051	0.086	0.124	-1.600

model, the state variable follows a continuous distribution and evolves continuously according to a normal AR(1) process:

$$s_t = \varphi\left(a_i\right)\, s_{t-1} + \sigma\, \varepsilon_t,$$

where ε_t is a standard normal random variable and $\varphi\left(a_i\right) = 0.81(a_i = 0) + 0.31(a_i = 1)$. The probability of choosing action 1 is assumed to take the following flexible functional form:

$$\sigma_i\left(a_i = k \,|\, a_{-i}, s_t\right) = \alpha_{0ik}\left(a_{-i}\right) + \alpha_{1ik} s_t + \alpha_{2ik} s_t^2 + \sum_{j=1}^{J}\left[\beta_{0j} + \beta_{1j}\cos\left(p_j s_t\right) - \beta_{2j}\sin\left(p_j s_t\right)\right],$$

where parameters α are fixed. The goal of this empirical exercise is to compare the payoff function estimated from the sample, generated by the state variable and the policy function using our one-stage estimation method and the utility function that we can recover by numerically solving the first-order condition for the player. We begin with describing the numerical computation algorithm. For each player i the value function associated with choice 0 can be expressed as

$$V_{i,0}\left(s\right) = \beta \int_{-\infty}^{+\infty} \log\left[\sum_{r=0}^{K} \exp\left(V_{i,r}\left(s'\right)\right)\right] \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\left(s' - \varphi(0)\, s\right)^2}{2\sigma^2}}\, ds'.$$

Using the relation $\sigma_i\left(k \,|\, s\right) = \frac{\exp\left(V_{i,k}(s)\right)}{\sum_{r=0}^{K}\exp\left(V_{i,r}(s)\right)}$, this expression can be written as a functional relation to solve for the continuation value function:

$$V_{i,0}\left(s\right) = \beta \int_{-\infty}^{+\infty} \left[V_{i,0}\left(s'\right) - \log\, \sigma_i\left(0 \,|\, s'\right)\right] \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\left(s' - \varphi(0)\, s\right)^2}{2\sigma^2}}\, ds'.$$

The value function will be approximated on a discrete uniform grid using linear extrapolation and the integral will be approximated by a Gauss-Hermite Gaussian quadrature method. The value function for the grid points will be solved from a system of linear equations. In particular, by a change of variables

$$V_{i,0}(s) = \frac{\beta}{\sqrt{\pi}} \int_{-\infty}^{+\infty} \left[V_{i,0}\left(\sqrt{2}\sigma\, x + \varphi\,(0)\; s\right) - \log\, \sigma_i \left(0 \,\middle|\, \sqrt{2}\sigma\, x + \varphi\,(0)\; s\right)\right] e^{-x^2}\, dx$$

$$\approx \frac{\beta}{\sqrt{\pi}} \sum_{n=1}^{N} \omega_n \left[V_{i,0}\left(\pm\sqrt{2}\sigma\, x_n + \varphi\,(0)\; s\right) - \log\, \sigma_i \left(0 \,\middle|\, \pm\sqrt{2}\sigma\, x_n + \varphi\,(0)\; s\right)\right],$$

where ω_n are the weights and x_n are the points of $2N$-point Gauss-Hermite quadrature approximation for the integral of interest. We aim to solve for the value function at a uniform grid $\mathbf{S}_G = \{s_1,\, s_2\, \ldots,\, s_G\}$ for the state variable: $V_{i,0}\,(s_g) = V_{i,0,g}$. For numerical computations we will use linear interpolation The intermediate values of the value function will be approximated by linear interpolation: for instance, if $s \in [s_g,\, s_{g+1}]$ then $V_{i,0}\,(s) \approx V_{i,0,g} + \frac{V_{i,0,g+1} - V_{i,0,g}}{s_{g+1} - s_g}\,(s - s_g)$. Let $\xi_{g,n,p}$ correspond to the index of the grid point that is not further from the point $(-1)^p\sqrt{2}\sigma\, x_n + \varphi\,(0)\; s_g$ than the cell length and has the smallest absolute value. Then the discretized Bellman equation can be written as G linear equations for the grid function:

$$V_{i,0,g} - \frac{\beta}{\sqrt{\pi}} \sum_{n=1}^{N} \sum_{p=0}^{1} \left[a_{g,n,p} V_{1,0,\xi_{g,n,p}} + b_{g,n,p} V_{i,0,\xi_{g,n,p}+1}\right]$$
$$= -\tfrac{\beta}{\sqrt{\pi}} \sum_{n=1}^{N} \sum_{p=0}^{1} \omega_n\, \log\, \sigma_i \left(0 \,\middle|\, (-1)^p\sqrt{2}\sigma\, x_n + \varphi\,(0)\; s_g\right)$$

Denote Δ the step of the grid. Then we can express the above coefficients as

$$a_{g,n,p} = \tfrac{\beta}{\Delta\sqrt{\pi}}\omega_n \left[s_{\xi_{g,n,p}+1} - (-1)^p\, \sqrt{2}\sigma\, x_n - \varphi\,(0)\; s_g \right],$$
$$b_{g,n,p} = \tfrac{\beta}{\Delta\sqrt{\pi}}\omega_n \left[(-1)^p\, \sqrt{2}\sigma\, x_n + \varphi\,(0)\; s_g - s_{\xi_{g,n,p}} \right].$$

We compare the utility function that we obtain from a numerical solution of the Bellman equation with the estimated payoff that we obtain using our method. The following table tabulates the integrated difference between the utility function that is numerically computed and the utility function that is estimated from a randomly generated sample. We use the stationary density of the state variable for the comparison. Specifically, if $\widehat{u}_T(\cdot)$ is the estimated utility from sample of size T and $u(\cdot)$ is the numerical solution, the reported criterion is

$$Q_T = \sqrt{T} \int\limits_{\mathcal{S}} \left(\widehat{u}_T(s) - u(s)\right)\, \pi(s)\, ds,$$

where $\pi(\cdot)$ is the stationary density of the state variable. We obtain this integral using the Monte-Carlo integration technique. To do so we make the joint draws from the state variable transition and the decision rules using a preliminary draw of the state variable. We generate the state variable as well as the policy rule as a Markov chain until it reaches the stationary distribution (we determine that by the behavior of the distribution mean across the blocks of consecutive draws). Then if N_s is the number of draws from the stationary distribution, we compute the approximate criterion

$$Q_T^{N_s} = \frac{\sqrt{T}}{N_s} \sum_{t=1}^{N_s} \left(\widehat{u}_T(s_t) - u(s_t)\right).$$

This object converges to the integral of interest as the number of draws increases. For our purposes we use 2.5 million draws.

Table 5: Simulation summary for entry utilities

sample size	mean	variance	median	90% quantile/10% quantile ratio
50	-0.2075	1.0898	-0.2371	0.0007
100	-0.2064	1.1111	-0.1776	0.0012
150	-0.2075	0.9341	-0.1935	0.0011
200	-0.2056	1.0461	-0.1934	0.0010
250	-0.2047	1.0346	-0.1936	0.0009
300	-0.2041	0.9111	-0.1851	0.0010

As the table shows the nonparametric procedure for recovering the primitive utilities works well in finite samples. In particular, the following figure illustrates the median of numerically recovered utility with top and bottom 10% quantiles for 600 Monte-Carlo draws.

Figure 1: Median and Percentils of numerically recovered utility

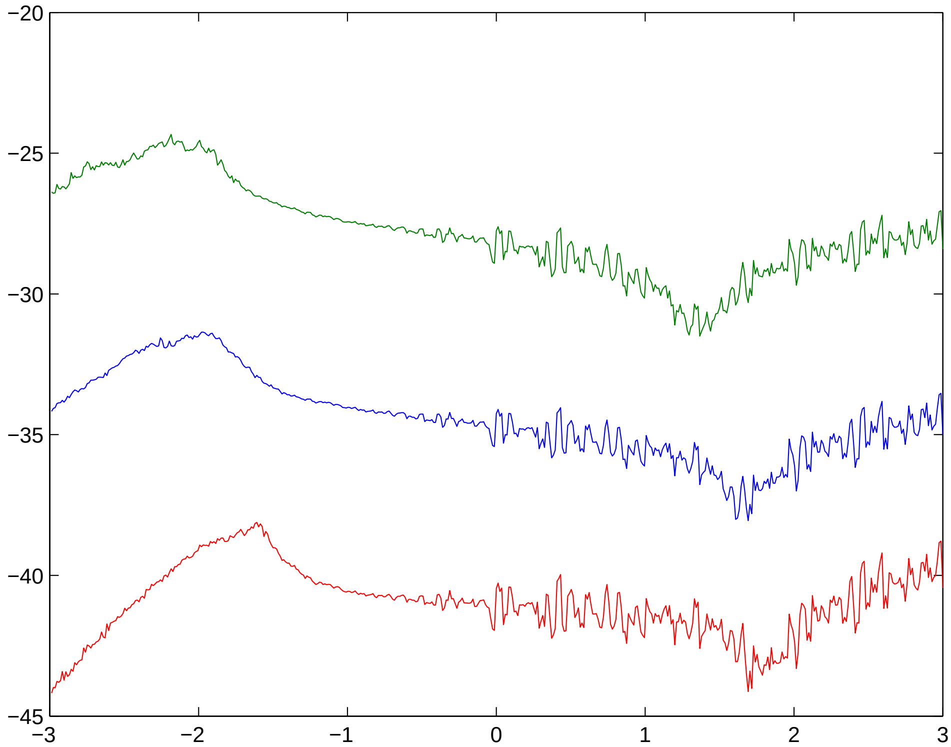

8 Conclusion

We study nonparametric identification of a dynamic discrete game model of incomplete information, and develop nonparametric and semiparametric estimators that have flexible computational properties and desirable statistical properties. Our identification analysis provides a unified framework for both discrete and continuous state variables, and suggests a natural implementation of a nonparametric estimator. In addition, we derive the semiparametric efficiency bound and propose a one-step semiparametric efficient estimator under the assumptions that the transition process is nonparametrically specified while the static payoff functions are parametric. A set of numerical simulations are used to demonstrate the properties of the model.

References

Ackerberg, D., X. Chen, and J. Hahn (2009): "A Practical Asymptotic Variance Estimator for Two-Step Semiparametric Estimators," unpublished manuscript, UCLA and Yale University.

Aguirregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," Forthcoming, *Econometrica*.

Ai, C., and X. Chen (2003): "Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions," *Econometrica*, 71(6), 1795–1843.

Bajari, P., C. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

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

Bresnahan, T., and P. Reiss (1991): "Empirical Models of Discrete Games," *Journal of Econometrics*, 48, 57—81.

Chen, X., O. Linton, and I. Van Keilegom (2003): "Estimation of Semiparametric Models when the Criterion Function Is Not Smooth," *Econometrica*, 71(5), 1591–1608.

Dunford, N., and J. T. Schwarz (1958): *Linear Operators. Part I: General Theory.* Wiley.

Fan, J., and I. Gijbels (1992): "Variable bandwidth and local linear regression smoothers," *The Annals of Statistics*, pp. 2008–2036.

Fershtman, C., and A. Pakes (2009): "Finite state dynamic games with asymmetric information: A framework for applied work," SSRN working paper.

Hotz, J., and R. Miller (1993): "Conditional Choice Probabilties and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529.

Jenkins, M., P. Liu, D. McFadden, and R. Matzkin (2004): "The Browser War: Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects," UC Berkeley, working paper.

Newey, W. (1990): "Semiparametric Efficiency Bounds," *Journal of Applied Econometrics*, 5(2), 99–135.

——— (1994): "The Asymptotic Variance of Semiparametric Estimators," *Econometrica*, 62, 1349–82.

Newey, W., and J. Powell (2003): "Instrumental variable estimation of nonparametric models," *Econometrica*, pp. 1565–1578.

Pesendorfer, M., and P. Schmidt-Dengler (2003): "Identification and Estimation of Dynamic Games," NBER working paper No. w9726.

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.

A Proof of theorem 3

We use the pathwise derivation method of Newey (1994) to calculate the form of the semiparametric asymptotic variance. Newey (1994) shows that for a general moment condition $m\left(z, h\left(\cdot\right)\right)$ that is a function of the nonparametric functional $h(s) = E(y|s)$, if one can find a function $\delta(s)$ such that for all parametric path $h_\eta(\cdot)$ of $h(\cdot)$, $\frac{\partial}{\partial\eta} Em(z, h_\eta(\cdot)) = E\delta(s)\frac{\partial}{\partial\eta} h(s)$, then the asymptotic variance of $\sum_\tau m\left(z_\tau, \hat{h}(\cdot)\right)$ will be given by

$$Var\left(m\left(z_\tau, h_0(\cdot)\right) + \delta(s_\tau)(y - h(s_\tau))\right).$$

In our case $m(z_\tau, h_0(\cdot)) \equiv 0$, therefore the asymptotic variance will be driven only by $\delta(s_\tau)(y - h(s_\tau))$. In the following analysis, we derive the form of y and $\delta(s)$ for our problem. To simplify notations, in the following we often omit the subscript i.

Part 1 corresponds to $Ed_\tau^k A(s_\tau)(\log\sigma_\eta(k|s_\tau) - \log\sigma_\eta(0|s_\tau))$. The pathwise derivative is given by $\frac{\partial}{\partial\eta} E\sigma(k|s_\tau)A(s_\tau)\log\frac{\sigma_\eta(k|s_\tau)}{\sigma_\eta(0|s_\tau)} = E\sigma(k|s_\tau)A(s_\tau)\left(\frac{1}{\sigma(k|s_\tau)}\frac{\partial}{\partial\eta}\sigma_\eta(k|s_\tau) - \frac{1}{\sigma(0|s_\tau)}\frac{\partial}{\partial\eta}\sigma_\eta(0|s_\tau)\right)$. Hence $\phi_1(a_{i\tau}, s_\tau)$.

Part 2 corresponds to $\beta Ed_t^k A(s_t)\log\sigma_\eta(0|s_{t+1})$. Its pathwise derivative is

$$\beta E\left[E\left(d_t^k A(s_t)|s_{t+1}\right)\right]\frac{1}{\sigma(0|s_{t+1})}\frac{\partial}{\partial\eta}\sigma_\eta(0|s_{t+1}).$$

Part 3 corresponds to $Ed_t^k A(s_t)V_\eta(0, s_t)$. First we derive an impression for $\frac{\partial}{\partial\eta}V_\eta(0, s_t)$, denoted $\dot{V}_\eta(0, s_t)$, in terms of a conditional expectation. Note that $V_\eta(0, s_t)$ is defined by

$$V_\eta(0, s_t) - \beta E_\eta\left[V_\eta(0, s_{t+1})|s_t, 0\right] = -\beta E_\eta\left[\log\sigma_\eta(0|s_{t+1})|s_t, 0\right],$$

which can be equivalently written as

$$\sigma_\eta(0|s_t)V_\eta(0, s_t) - \beta E_\eta\left[d_t^0 V_\eta(0, s_{t+1})|s_t\right] = -\beta E_\eta\left[d_t^0\log\sigma_\eta(0|s_{t+1})|s_t\right].$$

Differentiate totally with respect to η:

$$\begin{aligned}\sigma(0|s_t)\dot{V}_\eta(0, s_t) - \beta E\left[d_t^0\dot{V}_\eta(0, s_{t+1})|s_t\right] =& \\ \beta\dot{E}_\eta\left[d_t^0 V(0, s_{t+1})|s_t\right] - \dot{\sigma}_\eta(0|s_t)V(0, s_t) - \beta\dot{E}_\eta\left[d_0^t\log\sigma(0|s_{t+1})|s_t\right]& \\ - \beta E\left[d_t^0\frac{1}{\sigma(0|s_{t+1})}\dot{\sigma}_\eta(0|s_{t+1})|s_t\right].&\end{aligned}$$

The last term on the right hand side can be written as

$$E\left[d_t^0 \frac{1}{\sigma(0|s_{t+1})}\dot{\sigma}_\eta(0|s_{t+1})|s_t\right] = \int_{d_t^0, s_{t+1}} \int_{d_{t+1}^0} \frac{d_t^0 d_{t+1}^0}{\sigma(0|s_{t+1})} \dot{f}_\eta\left(d_{t+1}^0|s_{t+1}\right) f\left(d_t^0, s_{t+1}|s_t\right)$$

$$= \int_{d_t^0, s_{t+1}} \int_{d_{t+1}^0} \frac{d_t^0 d_{t+1}^0}{\sigma(0|s_{t+1})} \dot{f}_\eta\left(d_{t+1}^0|s_{t+1}, d_t^0, s_t\right) f\left(d_t^0, s_{t+1}|s_t\right)$$

$$= \int_{d_t^0, s_{t+1}} \int_{d_{t+1}^0} \frac{d_t^0 d_{t+1}^0}{\sigma(0|s_{t+1})} \left[\dot{f}_\eta\left(d_{t+1}^0, s_{t+1}, d_t^0|s_t\right) - f\left(d_{t+1}^0|s_{t+1}, d_t^0, s_t\right) \dot{f}_\eta\left(d_t^0, s_{t+1}|s_t\right)\right]$$

$$= \dot{E}_\eta\left[\frac{d_t^0 d_{t+1}^0}{\sigma(0|s_{t+1})}|s_t\right] - \dot{E}_\eta\left[d_t^0|s_t\right].$$

The second equality above follows from the Markov property and the conditional independence assumption. Therefore one can write

$$\sigma(0|s_t)\dot{V}_\eta(0,s_t) - \beta E\left[d_t^0 \dot{V}_\eta(0,s_{t+1})|s_t\right] = \dot{E}_\eta(y_t|s_t)$$

where $y_t = \beta d_t^0 V(0,s_{t+1}) - V(0,s_t) d_t^0 - \beta d_t^0 \log\sigma(0|s_{t+1}) - \beta\frac{d_t^0 d_{t+1}^0}{\sigma(0|s_{t+1})} + \beta d_t^0$.

In the next step we verify that the function $\delta(s_t)$ given in theorem 3 satisfies

$$E d_t^k A(s_t)\dot{V}_\eta(0,s_t) = E\delta(s_t)\dot{E}_\eta(y_t|s_t).$$

The left side is $E\sigma(k|s_t)A(s_t)\dot{V}_\eta(0,s_t)$ while the right side can be written as

$$E\delta(s_t)\dot{E}_\eta(y_t|s_t) = E\delta(s_t)\left(\sigma(0|s_t)\dot{V}_\eta(0,s_t) - \beta E\left[d_t^0\dot{V}_\eta(0,s_{t+1})|s_t\right]\right)$$

$$= E\delta(s_t)\sigma(0|s_t)\dot{V}_\eta(0,s_t) - \beta E d_{t-1}^0 \delta(s_{t-1})\dot{V}_\eta(0,s_t)$$

$$= E\left[\delta(s_t)\sigma(0|s_t) - \beta E\left(d_{t-1}^0\delta(s_{t-1})|s_t\right)\right]\dot{V}_\eta(0,s_t) = E\sigma(k|s_t)A(s_t)\dot{V}_\eta(0,s_t).$$

Part 4 is completely analogous to part 3, if we replace $\sigma_k(s_t)A(s_t)$ by $-\beta E\left(d_{t-1}^k A(s_{t-1})|s_t\right)$.

B Proof of theorem 4

First we need to characterize the tangent set of the model. The likelihood of the model will be determined by the choice probabilities and the transition density for the state variable. Given that choices of players are observed by the econometrician, the log-likelihood of the model can be written as

$$\mathcal{L}(s,s',d) = \sum_{i=1}^{n}\sum_{k=0}^{K} d^{i,k}\log\,\sigma_i(a_i = k\,|\,s) + \sum_{a\in\mathcal{A}} d^a \log\,g(s\,|\,s',a) + \log\,p(s'),$$

where $g(\cdot|s', a)$ is the transition density of the state variable, d^a is the indicator of the action profile a, and $p(\cdot)$ is the stationary density of the state variable. We choose a particular parameterization path θ for the model and compute the score by differentiating the model along the path:

$$S_\theta(s, s', d) = \sum_{a \in \mathcal{A}} d^a s_{1\theta}(s \mid s', a) + s_{2\theta}(s') + \sum_{i=1}^{n} \sum_{k=0}^{K-1} \left(\frac{d^{i,k}}{\sigma_i(k|s)} - \frac{d^{i,K}}{\sigma_i(K|s)} \right) \dot{\sigma}_i(k|s),$$

where $E[s_{1\theta}(s \mid s', a) \mid s', a] = 0$, $E[s_{2\theta}(s')] = 0$, $E\left[|s_{1\theta}(s \mid s', a)|^2 \mid s', a\right] < \infty$, $E|s_{2\theta}(s')|^2 < \infty$, and $E|\sigma_i(k|s)|^2 < \infty$. Then we characterize the tangent set as

$$\mathcal{T} = \left\{ \sum_{a \in \mathcal{A}} d^a \eta_1(s \mid s', a) + \eta_2(s') + \sum_{i=1}^{n} \sum_{k=0}^{K-1} \eta_3(s) \left(\frac{d^{i,k}}{\sigma_i(k|s)} - \frac{d^{i,K}}{\sigma_i(K|s)} \right) \right\},$$

with $E[\eta_1(s \mid s', a) \mid s', a] = 0$, $E[\eta_2(s')] = 0$, $E\left[|\eta_1(s \mid s', a)|^2 \mid s', a\right] < \infty$, $E|\eta_2(s')|^2 < \infty$, and $E|\eta_3(s)|^2 < \infty$. We will derive the semiparameric efficiency bound for this model under the absence of parametric restrictions on the state transition density. To derive the bound we find the parametric and the non-parametric parts of the score of the model using a particular parametrization path for the non-parametric component. For the chosen parametric path θ we denote

$$\frac{\partial V_i(k, s)}{\partial \theta} = \zeta_i(k, s) \quad \text{and} \quad \frac{\partial V_i(k, s)}{\partial \gamma'} = \tilde{\zeta}_i(k, s).$$

Also denote $\pi_i(k, s) = \frac{\partial \Pi(k, s; \beta)}{\partial \gamma'}$. We form vectors $V^i = (V_i(1, s), \ldots, V_i(K, s))'$, $V = (V^1, \ldots, V^n)'$, $\sigma^i = (\sigma_i(1|s), \ldots, \sigma_i(K|s))'$ and
$\zeta = (\zeta_1(1, s), \ldots, \zeta_1(K, s), \ldots, \zeta_n(K, s))'$. First of all, we note that we can transform the original moment equation. Consider the operator

$$\mathcal{P}_i \circ f = E[f(s') \mid s, a_i],$$

where expectation is defined for the conditional density $\sum_{a_{-i}} g(s'|s, a_i = k, a_{-i}) \sigma_{-i}(a_{-i}|s)$. This operator has a discrete spectrum with eigenfunctions $\left\{\Theta_j^{i,k}(s)\right\}_{j=0}^{\infty}$ and eigenvalues $\left\{\lambda_j^{i,k}\right\}_{j=0}^{\infty}$ different from zero. This follows directly from the properties of the Hibert-Schmidt operators which can be found in Dunford and Schwarz (1958). Then we can represent the value function as

$$V_i(k, s) = \sum_{j=0}^{\infty} a_j^{i,k} \Theta_j^{i,k}(s).$$

Then we can transform the moment equation to

$$\widetilde{\varphi}(s,\, s',\, a\,;\, \gamma,\, V_i,\, \sigma_i) = \sum_{j=0}^{\infty} a_j^{i,k} \left(1 - \beta\lambda_j^{i,k}\right) \Theta_j^{i,k}(s)$$

$$+ \left(1 - d^{a_i,0}\right) \left[-\Pi_i\left(a_i, a_{-i}, s;\, \gamma\right) + \beta \log\, \sigma_i\left(a_i|s'\right)\right] + d^{a_i,0} \beta \log \left(1 - \sum_{j=1}^{K} \sigma_i\left(j|s'\right)\right).$$

Then we can define a directional derivative of the moment function with respect to V_i in the direction h as

$$\left(\frac{\partial \varphi_i}{\partial V_i}\right)_h = \sum_{j=0}^{\infty} h_j \left(1 - \beta\lambda_j^{i,k}\right) \Theta_j^{i,k}(s),$$

for all h with $\sum_{j=0}^{\infty} |h_j| \left\|\Theta_j^{i,k}(s)\right\| < \infty$. Differentiating the unconditional moment equation with respect to the parametrization path we obtain

$$E\left[\mathcal{A}(s, d^a)\, \pi(s) \left(1 - d^{a,0}\right)\right] \dot{\gamma} - E\left[\mathcal{A}(s, d^a) \left(\tfrac{\partial \varphi}{\partial V}\right)_h\right] \dot{h}$$

$$= \beta E\left[\mathcal{A}(s, d^a) \left(\tfrac{d^{a\neq 0}}{\sigma(a|s')} - \tfrac{d^{a,0}}{\sigma(0|s')}\right)\right] + E\left[\mathcal{A}(s', d^a)\ \varphi\, s_{1\theta}\right].$$

We consider the right-hand side and try to find a function $\widetilde{\Psi}$ such that the expression on the right-hand side can be represented as $\langle \Psi,\, S_\theta \rangle$. This function can be obtained as

$$\widetilde{\Psi} = \mathcal{A}(s,\, d^a) \left\{ \left(\varphi - E\left[\varphi \,|\, s,\, a\right]\right) + \frac{d^{a\neq 0} - \sigma(a|s)}{\sigma(a|s')} - \frac{d^{a,0} - \sigma(0|s)}{\sigma(0|s')} \right\}.$$

We note that conditional moment equation (26) holds and we can differentiate it with respect to the parameterization path. Then we can substitute the expression for the derivative into the expression for the unconditional moment. This allows us to express the directional derivative of γ and, consequently, the efficient influence function for a fixed instrument matrix:

$$\Psi = E\left[\mathcal{A}(s, d^a) \left(\pi(s)\left(1 - d^{a,0}\right) - \left(\tfrac{\partial \varphi}{\partial V}\right)_h\right)\right]^{-1} \widetilde{\Psi}.$$

The semiparametric efficiency bound as a minimum variance of the influence function. Denoting

$$\Omega(s,\, a) = \mathrm{Var}\left(\varphi + \frac{d^{a\neq 0}}{\sigma(a|s')} - \frac{d^{a,0}}{\sigma(0|s')} \,\middle|\, s,\, a\right).$$

Using standard GMM arguments, we can express the efficiency bound for fixed instrument as

$$V_h(\beta) = \left(\left(\pi(s)\left(1 - d^{a,0}\right) - \left(\tfrac{\partial \varphi}{\partial V}\right)_h\right) \zeta\left(d^a, s\right)' \, \Omega(s, a)^{-1} \, \zeta\left(d^a, s\right) \left(\pi(s)\left(1 - d^{a,0}\right) - \left(\tfrac{\partial \varphi}{\partial V}\right)_h\right)\right)^{-1}.$$

The efficiency bound overall can be found as $V_{h^*}(\beta)$ for h^* solving

$$\inf_h \left(\pi(s)\left(1-d^{a,0}\right)-\left(\frac{\partial\varphi}{\partial V}\right)_h\right)\zeta\left(d^a,s\right)'\Omega\left(s,a\right)^{-1}\zeta\left(d^a,s\right)\left(\pi(s)\left(1-d^{a,0}\right)-\left(\frac{\partial\varphi}{\partial V}\right)_h\right).$$

The optimal instrument matrix can be explicitly written as

$$\mathcal{M}(s)=E\left[\left(\pi(s)\left(1-d^{a,0}\right)-\sum_{j=0}^{\infty}h_j^*\left(1-\beta\lambda_j\right)\Theta_j(s)\right)\zeta\left(d^a,s\right)'\Omega\left(s,a\right)^{-1}\,\Bigg|\,s\right].$$

Q.E.D.

Proof of theorem 6

We can use the Bellman equation to express the estimate of the payoff function in terms of the estimate of the value function. We use a series projection estimator to estimate $V_i(k,s)-V_i(0,s)$. To evaluate the elements of the Bellman equation for player i we need to analyze the right hand side function

$$h_i(s)=E\left\{\log\sum_{k=0}^{K}\left(V_i\left(k,s'\right)-V_i\left(0,s'\right)\right)\Bigg|\,s\right\}.$$

Function $h_i(s)$ admits a series representation $h_i(s)=\sum_{j=1}^{k(m)}q_j(s)\lambda_{i,j}^{k(m)}+o\left(\|q_{k(m)}(s)\|\right)$, where we use the standard Sobolev norm. The coefficients for this representation can be obtained from the coefficients for $V_i(k,s)-V_i(0,s)$. This result can be used to find a series representation for $V_i(0,s)$ which needs to be estimated. To do that we proceed by analyzing the nonparametric conditional expectation estimation component of step two, which takes the form of

$$V_i(s,0)=\beta\int V_i\left(s',0\right)g_i\left(s'|s,0\right)ds'+h_i(s)=\left(\mathcal{K}_i\circ V_i\right)(s,0)+h_i(s), \tag{28}$$

where $g_i\left(s'|s,0\right)=\sum_{a_{-i}\in\mathcal{A}_{-i}}g\left(s'|s,0,a_{-i}\right)\sigma\left(a_{-i}|s\right)$.

This is an integral equation for $V_i(\cdot,0)$. We assume that the integral operator $\mathcal{K}_i$ and the term $h_i(\cdot)$ satisfy the standard assumptions assuring the existence of a smooth solution of this equation. In particular $s\in S$, $V_i:S\mapsto\mathbb{R}_+$, both the kernel function $g_i(\cdot)$ and the function $h_i(\cdot)$ have derivatives up to order $p\geq k(m)$, which assures a high degree of smoothness of the value function. Thus $V_i\in C^p(S)$, and $\mathcal{K}_i:C^p(S)\mapsto C^p(S)$. A standard method for solving this equation is to represent solution by a series expansion over a particular basis in $C^p(S)$. We will use the basis

$q^{k(m)}(s) = \left(q_1(s), \ldots, q_{k(m)}(s)\right)'$ for these purposes. Then the approximation for the value function can be written as:

$$V_i\left(\cdot, 0\right) = q^{k(m)}(s)'\,\theta_i^{k(m)}.$$

We endow the space $C^p\left(S\right)$ with an inner product $\langle \cdot, \cdot \rangle$ and introduce matrices

$$\Gamma = \left(\langle q_t(s), q_j(s)\rangle\right)_{t,j=1}^{k(m)} \quad \text{and} \quad G_i = \left(\langle \mathcal{K}\, q_t(s), q_j(s)\rangle\right)_{t,j=1}^{k(m)}.$$

We define the inner product for two functions $f, g \in C^p\left(S\right)$ as:

$$\langle f,\, g\rangle = \int_S f(s)g(s)\pi\left(ds\right) = E\left[f(s)\, g(s)\right],$$

where $\pi(\cdot)$ is a stationary distribution measure for the state space S. In general, this measure is not available. For this reason, we substitute it with the empirical measure $\pi^m\left(\cdot\right)$, which we require to be weekly converging to $\pi(\cdot)$. We call the space associated with the inner product generated by $\pi^m\left(\cdot\right)$ by $C^{pm}\left(S\right)$. This space is only a semi-Hilbert space as the inner product in it might have a non-empty kernel (and, thus, the associated norm is only a seminorm). We will use the same basis in $C^{pm}\left(S\right)$ as before. In general, this unfortunate property will not create additional complications as long as the measure of the kernel of the seminorm associated with the inner product vanishes as the number of available markets m increases. For this reason, in the further discussion we will assume that the stationary measure in S is known, and then extend our results to the case when we are using empirical measure instead.

We can use the expansion for $h_i(\cdot)$ to derive the series approximation for the value function $V_i\left(\cdot, 0\right)$. In this case the vector of coefficients in the series representation of the value function can be found as:

$$\theta_i^{k(m)} = \left(\Gamma - \beta G_i\right)^{-1}\Gamma\,\lambda_i^{k(m)}.$$

This result is obtained from substituting series expansions for $h_i(\cdot)$ and $V_i(\cdot, 0)$ into equation (28) and projecting both sides of this equation on the basis vectors $q^{k(m)}(\cdot)$.

These coefficients allow us to obtain an approximation for the value of the function $V_i\left(\cdot, 0\right)$ which can be expressed as:

$$V_i^{k(m)}\left(s, 0\right) = q^{k(m)}(s)'\left(\Gamma - \beta G_i\right)^{-1}\Gamma\,\lambda_i^{k(m)}$$

For sufficiently smooth coefficients of the original integral equation, this expression will provide an approximation of order $k(m)$ such that the norm of the deviation of the approximation from the true solution will be bounded from above by $\frac{L}{k(m)!}\sup\limits_{s,s'\in\Omega}\|s-s'\|^{k(m)}$, where $\Omega\subset\mathcal{S}$ is a subset of the state space where the value function is approximated by the series expansion. Note that all components of this formula are exactly known, although the matrices are specific to a particular basis. For instance, if $q^{k(m)}(\cdot)$ is a system of Legendre polynomials then $\Gamma=\operatorname{diag}\left\{\sqrt{\frac{2}{2k(m)+1}}\right\}$.

We estimate coefficients in the series representation of the value function from the data. To do so, first, we estimate the state transition probability. We assume that an estimator with the rate $r\in(0,\,1/2]$ is available which produces the estimate that is point-wise asymptotically normal at s' uniformly over s in Ω:

$$n^r\left(\widehat{g}_i\left(s'\mid s,0\right)-g_i\left(s'\mid s,0\right)\right)\xrightarrow{d}N\left(0,\sigma_g^2(s',s)\right).$$

We assume for convenience that this estimate is obtained using an estimation procedure which can be approximated by a series expansion with the order of precision at least $o_p\left(n^{-r}\right)$. To estimate the vector of coefficients $\lambda^{k(m)}$ we use the data from the observed states and values of $h_i(\cdot)$ to estimate it. Note that the values of $h_i(\cdot)$ are obtained from the Hotz-Miller-type inversion and thus contain noise. By the nature of this inversion we can in principle evaluate $\widehat{h}_i(\cdot)$ at any point of Ω. Although the probabilities of actions are estimated non-parametrically, by Delta-method we can assure that for some $q\in(0,\,1/2]$ we obtain a point-wise asymptotically normal estimator of $h_i(\cdot)$ in Ω. In particular we use a spectral representation of $h_i(\cdot)$ to estimate it non-parametrically and obtain the coefficients λ^p. Thus

$$n^q\left(\widehat{h}_i\left(s\right)-h_i\left(s\right)\right)\xrightarrow{d}N\left(0,\sigma_h^2(s)\right).$$

We consider the properties of the pointwise approximation error for the value function:

$$\begin{aligned}\hat{V}_i^{k(m)}\left(s,0\right)-V_i\left(s,0\right)&=q^{k(m)}(s)'\left(\Gamma-\beta G_i\right)^{-1}\langle\widehat{h}_i(s)-h_i(s),q^{k(m)}(s)\rangle\\&+\beta q^{k(m)}(s)'\left(\Gamma-\beta G_i\right)^{-1}\langle\left(\widehat{\mathcal{K}}_i-\mathcal{K}_i\right)q^{k(m)}(s),\,q^{k(m)}(s)'\rangle\left(\Gamma-\beta G_i\right)^{-1\prime}\Gamma\,\lambda^{k(m)}+\Delta^{k(m)}.\end{aligned}$$

In this expression $\Delta^{k(m)}$ is a residual function. In the expression for the error in the estimate of the value function the matrices only play the role of normalization while the asymptotic behavior of the error is governed by the integrated error in the estimated components of the Bellman equation. This normalization does not change the rate of convergence of the estimators, and the order of polynomial expansion is determined only by the degree of smoothness of the function approximation. Assumption 5 restricts the operator $\mathcal{K}$ to be bounded. Consider the transformation $\lambda\mapsto\Gamma^{1/2}\lambda$ and $q^{k(m)}(\cdot)\mapsto\Gamma^{-1/2}q^{k(m)}(\cdot)$. This a rotation of the basis which does not change the asymptotic properties. In fact, indicating the rotated variables by tildes we get:

$$m^q\left(\tilde{q}^{k(m)}(s)'\widetilde{\hat{\lambda}}_i^{k(m)}-\tilde{q}^{k(m)}(s)'\tilde{\lambda}_i^{k(m)}\right)\xrightarrow{d}N\left(0,\,\sigma_\psi^2\right).$$

Specifically, $\sigma_\psi^2 = \lim_{m\to\infty} \text{trace}\left\{m^{2q}\tilde{\Omega}_\lambda \tilde{q}^{k(m)}(s)\tilde{q}^{k(m)}(s)'\right\} = \lim_{m\to\infty} \text{trace}\left\{m^{2q}\Omega_\lambda q^{k(m)}(s)q^{k(m)}(s)'\right\} = \sigma_\psi^2(s)$.

Next, note that $I_{k(m)} \leq \left(I_{k(m)} - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2}\right)^{-1} \leq (1-\beta)^{-1} I_{k(m)}$, where inequality should be treated as the difference between the two matrices is a positive semi-definite matrix. We can show that the last inequality is valid in to steps. First, the matrix $\Gamma - G_i$ is positive semi-definite because the operator $\mathcal{K}$ is defined by a density function. Second, the matrix $(1-\beta)^{-1}\left(I_{k(m)} - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2}\right) - I_{k(m)}$ is positive semi-definite. To see that, consider decomposition

$$I_{k(m)} - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2} = (1-\beta)\, I_{k(m)} + \beta\Gamma^{-1/2}\left(G_i - \Gamma\right)\Gamma^{-1/2} \geq (1-\beta)\, I_{k(m)}.$$

As a result:

$$\begin{aligned}
&\text{trace}\left\{m^{2q}\left(\Gamma - \beta G_i\right)^{-1}\Gamma\Omega_\lambda\Gamma\left(\Gamma - \beta G_i\right)^{-1} q^{k(m)}(s)q^{k(m)}(s)'\right\} \\
&\qquad = \text{trace}\left\{m^{2q}\left(I_k - \beta\Gamma^{-1/2}G\Gamma^{-1/2}\right)^{-1}\tilde{\Omega}_\lambda\left(I_{k(m)} - \beta\Gamma^{-1/2}G\Gamma^{-1/2}\right)^{-1}\tilde{q}^{k(m)}(s)\tilde{q}^{k(m)}(s)'\right\}.
\end{aligned}$$

This means that

$$\omega_1^2 = \lim_{m\to\infty} \text{trace}\left\{m^{2q}\left(\Gamma - \beta G_i\right)^{-1}\Gamma\Omega_\lambda\Gamma\left(\Gamma - \beta G_i\right)^{-1} q^{k(m)}(s)q^{k(m)}(s)'\right\} < \frac{\sigma_\psi^2}{(1-\beta)^2},$$

and it does not vanish. This proves that the rate of convergence of the non-parametric estimate for $V_i(\cdot, 0)$ is the same as the rate for $h_i(\cdot)$.

In the analysis so far we assume that the convergence rate of the estimator for the transitional density $g_i\left(\cdot \mid s, 0\right)$ is fast enough so that we can ignore the error associated with its estimation. We assume that we are using the same spectral representation for this conditional density. The relevant characteristic for our analysis is matrix G_i showing how transitional density changes the basis vectors. Then the following theorem establishes the asymptotic normality of the error in the value function associated with the error in the estimation of matrix G_i: By our assumption, the estimator for $\varphi(\cdot)$ is poitwise asymptotically normal. Consider local representation of the error in the value function:

$$\begin{aligned}
&\widehat{V}_i^{k(m)}(s,0) - V_i^{k(m)}(s,0) = q^{k(m)}(s)'\Gamma^{-1/2}\left(I - \beta\Gamma^{-1/2}\hat{G}_i\Gamma^{-1/2}\right)^{-1}\Gamma^{1/2}\lambda_i^{k(m)} - V_i(s) \\
&= \beta q^p(s)'\Gamma^{-1/2}\left(I_{k(m)} - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2}\right)^{-1}\Gamma^{-1/2} \\
&\qquad\qquad \times\left(\hat{G}_i - G_i\right)\Gamma^{-1/2}\left(I_{k(m)} - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2}\right)^{-1}\Gamma^{1/2}\lambda^p + o_p\left(\left\|\hat{G}_i - G_i\right\|\right),
\end{aligned}$$

where $\left\|\hat{G}_i - G_i\right\|$ is a standard matrix norm. Consider next the rotation of the basis and the coefficients by the matrix $\Gamma^{1/2}$ and denote the rotated variables by tildes. By the assumption, we know that

$$\sigma_\varphi(s) = \lim_{m\to\infty}\left\{m^{2r}\text{Var}\,\text{trace}\left(\left[\widehat{\tilde{G}}_i - \tilde{G}_i\right]\right)\right\} < \infty.$$

Then from the positive semi-definiteness of $\left(I_{k(m)} - \beta\Gamma^{-1/2}G\Gamma^{-1/2}\right)^{-1}$ we conclude that

$$\omega_2^2 \leq (1-\beta)^{-4}\lim_{m\to\infty}\left\{m^{2r}\text{Var}\,\text{trace}\left(\left[\widehat{\tilde{G}}_i - \tilde{G}_i\right]\Gamma^{1/2}\lambda_i^{k(m)}\lambda_i^{k(m)\prime}\Gamma^{1/2}\right)\right\} < \infty.$$

In this expression $\Gamma^{1/2}\lambda^{k(m)}\lambda^{k(m)\prime}\Gamma^{1/2}$ is *a priori* finite due to the Bessel's inequality. The normality of the error follows from the assumption about the distribution of $g_i(\cdot \mid s_0, 0)$ and the fact that matrix multiplication is a linear operation over the elements of the matrix.

The approximation for the value function can be expressed in terms of subsequent projections. From the Bellman's equation it follows that

$$\widehat{V}_i^{k(m)}(s,0) - V_i(s,0) - \beta E\left[\hat{V}_i(s',0) - V_i(s',0)\big|s\right] = \beta\left(\hat{E}\left[V_i(s',0)\big|s\right] - E\left[V_i(s',0)\big|s\right]\right) + \Delta, \quad (29)$$

with the residual Δ. Using the spectral representation for the expectation in the basis $q^{k(m)}(\cdot)$ (where the coefficients of $V_i(\cdot,0)$ in this basis are denoted $\theta^{k(m)}$) we obtain that up to the error of order smaller than Δ:

$$\hat{E}\left[V(s',0) \mid s\right] - E\left[V(s',0) \mid s\right] = q^{k(m)}(s)'\Gamma^{-1}\left(\hat{G}_i - G_i\right)\theta^{k(m)},$$

$$E\left[\hat{V}(s') - V(s')\big|s\right] = q^{k(m)}(s)'\Gamma^{-1}G_i\left(\hat{\theta}^{k(m)} - \theta^{k(m)}\right).$$

From spectral representation of the Bellman's equation it follows that (up to the series approximation error):

$$\theta^{k(m)} = \left(\Gamma - \beta G_i\right)^{-1}\Gamma\lambda^{k(m)}.$$

Substitution of these expressions into (29) gives:

$$\widehat{V}_i^{k(m)}(s,0) - V_i^{k(m)}(s,0) = \beta q^{k(m)}(s)'\Gamma^{-1/2}\left(I - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2}\right)^{-1}\Gamma^{-1/2}$$
$$\times\left(\hat{G}_i - G_i\right)\Gamma^{-1/2}\left(I - \beta\Gamma^{-1/2}G_i\Gamma^{-1/2}\right)^{-1}\Gamma^{1/2}\lambda_i^{k(m)}.$$

This suggests that the method of approximating value function by consecutive conditional expectations (29) is equivalent to the spectral approach up to approximation error.

Now we will discuss the case where we substitute the stationary measure $G_i(\cdot)$ by its empirical analog. In this case for the sample $\{s_l\}_{l=1}^m$ the inner product for $f, g \in C^p(S)$ can be defined as:

$$\langle f,\, g\rangle^m = \sum_{l=1}^m f(s_l)\, g(s_l).$$

We can describe the quality of approximation only outside the kernel of the seminorm in $C^{pm}(S)$. In that part of the subspace the norm of the elements of the basis is well-defined. For this reason, we can write the same expressions for the coefficients for expansion of the value function in the basis $q^{k(m)}(\cdot)$ but in terms of matrices Γ^m and G^m defined by the inner product in $C^{pm}(\mathcal{S})$. In this case, the problem of evaluation of the difference between the estimate of the value function obtained from Γ^m and G^m and the true value reduces to two separate problems. The first one is evaluation of the error due to series approximation, which was considered above. The second one is evaluation of the quality of approximation when using empirical measure instead of the true stationary measure. The general results regarding these properties are given, for instance, in (Billingsley, 1968). Here we will consider a special case when the stationary and empirical measures have densities. We can evaluate the quality of approximation of the value function as:

$$V_i^{m,k(m)}(s) - V_i^{k(m)}(s) = \beta q^p(s)\left(\Gamma - \beta G_i\right)^{-1}\left(G_i^m - G_i\right)\left(\Gamma - \beta G_i\right)^{-1}\Gamma\lambda^p$$
$$+\beta q^{k(m)}(s)\left(\Gamma - \beta G_i\right)^{-1}\left(\Gamma^m - \Gamma\right)\left[I - \left(\Gamma - \beta G_i\right)^{-1}\Gamma\right]\lambda_i^{k(m)} + o\left(\|\Gamma^m - \Gamma\|, \|G_i^m - G_i\|\right),$$

where the norm in the residual term is a standard matrix norm. This expression has similar structure as the expression for the errors due to estimation of $h_i(\cdot)$. From Assumption 5 it follows that traces of matrices $\Gamma^m - \Gamma$ and $G_i^m - G_i$ approach to zero faster than $m^{\max\{q,r\}}$. This means that in the asymptotic expansion the corresponding term vanishes as well.

This result proves that we can, in general, substitute the matrices G_i and Γ by their sample versions without affecting the asymptotic variance. The estimate of the value function will take the form:

$$\hat{V}_i^{k(m)}(s,0) = q^{k(m)}(s)'\left(\hat{\Gamma} - \beta\hat{G}_i\right)^{-1}\hat{\Gamma}'\hat{\lambda}_i^{k(m)},$$

where $\hat{\Gamma}$ and $\hat{G}_i$ are sample averages for estimating Γ and G. For example:

$$\hat{G}_i = \frac{1}{m}\sum_{j=1}^{m}\frac{1}{T}\sum_{t=1}^{T-1} q^{k(m)}\left(s_{j,t+1}\right) q^{k(m)}\left(s_{j,t}\right)'.$$

In the previous step we have estimated $V_i(s,l) - V_i(s,0)$ non-parametrically as $q^{k(m)\prime}\gamma_{i,l}^{k(m)}$. This means that the non-parametric estimate for the choice-specific value function is a combination of the obtained estimate for $V_i(s,0)$ and this difference and:

$$\hat{V}_i^{k(m)}(s,l) = q^{k(m)\prime}\left(\hat{\theta}_i^{k(m)} + \hat{\gamma}_{i,l}^{k(m)}\right).$$

This variable will be normal as it is non-degenerate and computed as a sum of two asymptotically normal estimates. This fact becomes straightforward if we explicitly express coefficients $\theta_i^{k(m)}$ in terms of $\gamma_{i,l}^{k(m)}$. Let $\gamma_i^{k(m)} = \left(0, \gamma_{i,1}^{k(m)}, \ldots, \gamma_{i,K}^{k(m)}\right)$ be the stacked matrix of coefficients in the expansions for $V_i(s,l) - V_i(s,0)$. We introduce the following vector of logit probabilities:

$$\Lambda = \left(\frac{\exp(V_i(s,l) - V_i(s,0))}{\sum\limits_{j=0}^{K} \exp(V_i(s,j) - V_i(s,0))} \right)_{l=1,\ldots,K}$$

Then we can express $\lambda_i^{k(m)}$ (up to the error of approximation) as:

$$\lambda_i^{k(m)} = \Gamma^{-1} G_i \gamma_i^{k(m)} \Lambda.$$

Therefore, the corresponding coefficients for the value function can be expressed as:

$$\theta_i^{k(m)} = (\Gamma - \beta G_i)^{-1} G_i \gamma_i^{k(m)} \Lambda.$$

Value function can be explicitly estimated from coefficients $\hat{\gamma}_i^{k(m)}$ and matrices G_i and Γ as:

$$\hat{V}_i^{k(m)}(s,l) = q^{k(m)\prime} \left(\hat{\gamma}_{i,l}^{k(m)} + \left(\hat{\Gamma} - \beta \hat{G}_i\right)^{-1} \hat{G}_i \hat{\gamma}_i^{k(m)} \hat{\Lambda} \right).$$

From this estimate one can see that the estimate for the value function is obtained from the estimates for the choice-specific probabilities by permuting them by bounded linear transformations (as $\sum_t \Lambda_t = 1$ and $\Lambda_t > 0$, while the operator represented by the matrix $I - \beta \Gamma^{-1/2} G_i \Gamma^{-1/2}$ is bounded as shown above). This motivates asymptotic normality with non-degenerate distribution for their estimates. Estimated profit will be, again, a non-degenerate linear combination of the estimates for the choice-specific probabilities, and pointwise normality of the estimate with the rate of convergence, corresponding to the minimum of the convergence rate for the choice specific probability or transition density.

To formalize this recall that we can compute the profit function from the value function by the formula:

$$\Pi_i(s,l) = V_i(s,l) - \beta E\left[V_i(s') \,\middle|\, s, a_i = l\right].$$

Let $G_i^{(l)}$ be the matrix corresponding to the state transition density $g_i(s' \mid s,l)$ such that $G_{i,tr}^{(l)} = \int\int g_i(s' \mid s,l)\, q_t^{k(m)}(s') q_r^{k(m)}(s) \pi(ds)\, ds'$. We can then express the spectral representation for the profit as:

$$\Pi_i^{k(m)}(s,l) = q^{k(m)\prime}(s) \left(\gamma_{i,l}^{k(m)} + \left\{ \left[I_{k(m)} - \Gamma^{-1} G_i^{(l)} \right] \left[(\Gamma - \beta G_i)^{-1} G_i + I_{k(m)} \right] - I_{k(m)} \right\} \gamma_i^{k(m)} \Lambda \right).$$

Then we can transform the expression for the profit function as:

$$\Pi_i^{k(m)}(s,l) = q^{k(m)\prime}(s) \gamma_{i,l}^{k(m)} + \tilde{q}^{k(m)\prime}(s) \Big\{ \left[I_{k(m)} - \Gamma^{-1/2} G_i^{(l)} \Gamma^{-1/2} \right]$$
$$\times \left[\left(I_{k(m)} - \beta \Gamma^{-1/2} G_i \Gamma^{-1/2} \right)^{-1} + \Gamma^{1/2} G_i^{-1} \Gamma^{1/2} \right] \Gamma^{-1/2} G_i \Gamma^{-1/2} - \Gamma^{-1/2} \Big\} \tilde{\gamma}_i^{k(m)} \Lambda.$$

In this expression tildes denote the rotation of the basis considered before. The matrix in the second expression represents a bounded linear transformation due to assumption 3. Therefore the estimate for the profit function is a bounded transformation of the estimate of the choice probabilities.