

Semiparametric Estimation of a Dynamic Game of Incomplete Information

Patrick Bajari and Victor Chernozhukov and Han Hong*
University of Michigan and NBER
MIT
Duke University

September 5, 2005

Abstract

Recently, empirical industrial organization economists have proposed estimators for dynamic games of incomplete information. In these models, agents choose from a finite number actions and maximize expected discounted utility in a Markov perfect equilibrium. Previous econometric methods estimate the probability distribution of agents' actions in a first stage. In a second step, a finite vector of parameters of the period return function are estimated. In this paper, we develop semiparametric estimators for dynamic games allowing for continuous state variables and a nonparametric first stage. The estimates of the structural parameters are $T^{1/2}$ consistent (where T is the sample size) and asymptotically normal even though the first stage is estimated nonparametrically. We also propose sufficient conditions for identification of the model.

*We would like to thank the National Science Foundation and the Sloan Foundation for generous research support. We would also like to thank John Krainer for detailed comments on an earlier draft.

1 Introduction.

In economic theory it is common place, if not standard, to model market equilibrium as a game. Game theory has profoundly influenced how economic theorists conceptualize markets and regularly influences policy debates. By comparison, the impact of game theory on most applied fields in economics has been much less significant. While there is a large empirical literature that tests the predictions of certain games, there is by comparison much less work that formally models the relationship between endogenous and exogenous variables in light of game theory. Rigorously understanding the econometric implications of game theory is clearly a necessary condition for coherence between theoretical and empirical work in economics.

Over the past decade, structural estimation of game theoretic models has been a topic of active research within the subfield of empirical industrial organization. See Akerberg, Benkard, Berry, and Pakes (2005) for an excellent survey. In particular, researchers have recently proposed two-step estimation methods for dynamic games. Like the models surveyed in Rust (1994) or studied in Keane and Wolpin (1997), agents choose from a finite set of actions and utility at a particular point in time is a function of covariates and stochastic preference shocks. Agents are forward looking and maximize expected discounted utility. However, unlike Rust, agents interact strategically and play a Markov perfect equilibrium to a dynamic game. See Pakes, Ostrovsky and Berry (2003), Aguirregabiria and Mira (2002), Pesendorfer and Schmidt-Dengler (2003) and Bajari, Benkard and Levin (2003). Substantive applications of dynamic games estimators include Jenkins, Liu, McFadden, and Matzkin (2004) to the browser war and Ryan (2005) to regulation in a concentrated industry. The problem of equilibrium existence is also considered in Jenkins et. al. (2004) as well.

Like Hotz and Miller (1993) or Guerre, Perrigne and Vuong (2000), estimation in these models essentially takes place in two steps. In the first step, the economist flexibly estimates agents' reduced form policy functions. In many cases, this involves estimating the probability that each of the finite set of actions is played conditional on a finite vector of state variables. In the second step, the structural parameters of the model are estimated. Typically, these are a finite vector that parameterize the period utility as a function of actions and states. These estimators are attractive for applied work since they impose a relatively low computational burden and are straightforward to program.

The formal econometric arguments used in these papers are based standard methods for the analysis of parametric models. Such methods are appropriate for problems where (i) there are a finite number of discrete state variables or (ii) there are continuous state variables but the parametric first stage is correctly specified. However, the econometric theory has not been worked out for the case of a nonparametric first stage and continuous state variables. This case is important since many problems have state variables that are naturally modeled as continuous. For instance, the state variable in models of strategic firm entry/exit usually includes productivity which the literature usually models as continuous.

One could of course ignore the fact that the state variables are continuous and discretize the state space. However, increasing the number of grids in estimating the first stage choice probabilities has two offsetting effects. It reduces the bias in the first stage estimation but on the other hand increases the variance. When the dimension of the continuous state variables is larger than four, it is impossible to obtain \sqrt{T} consistent (where T is the sample size) and asymptotically normal estimators for the second stage parameters through discretization. It can be shown that the variance of the discretization procedure is of the magnitude of $1/\sqrt{Th^d}$ where d is the dimension of the continuous state variables and h is the window size used in the discretization. The bias, on the hand, is of the magnitude of $\sqrt{Th^2}$. \sqrt{T} consistency of the parameter estimator requires that both the variance and the bias decrease to zero as the sample size increases to ∞ . It can easily be shown that this is impossible, however, when d is larger than 4.

In this paper, we consider the problem of semiparametric estimation of a dynamic game of incomplete information similar to models discussed by Pakes, Ostrovsky and Berry (2003), Aguirregabiria and Mira (2002), Pesendorfer and Schmidt-Dengler (2003) and Bajari, Benkard and Levin (2003). As in these earlier papers, the goal of estimation is the recovery of a finite number of parameters in the players' period utility functions. Estimation takes place in four steps and, like the earlier literature imposes a fairly low computational burden. In the first step, the econometrician uses sieve methods to nonparametrically estimate choice probabilities as a function of the state variables. (See Ai and Chen (2003) and Chen (2005)). In the second step, using standard formulas that relate choice probabilities to choice-specific value functions, the econometrician recovers an estimate of the choice-specific value function. (See Hotz and Miller (1993) and Aguirregabiria and Mira (2002)). In the third step, the econometrician generates choice-specific value functions consistent with a guess about the parameters of the period utility function. In the final step, the econometrician minimizes the distance between the choice-specific value functions derived in the second and third steps.

While the construction of our estimator is in many ways inspired by insights from the earlier literature, our approach to establishing the asymptotic properties of the estimator is quite distinct. Building on the analysis of Newey (1994), we prove that our estimators are \sqrt{T} consistent and asymptotically normal even if the first stage is estimated nonparametrically. For expositional simplicity, we consider the case of period utility functions that are linear in the parameters as in Bajari, Benkard and Levin (2003). However, we demonstrate in the last section that our methods can be generalized to models that do not impose this assumption.

Finally, we consider the problem of identification. Nonparametric identification results for dynamic discrete games are developed by Aguirregabiria and Mira (2002) and Pesendorfer and Schmidt-Dengler (2003) in the context of discrete state space models. Recent works by Heckman and Navarro (2005) and Aguirregabiria (2005) present identification results for dynamic discrete choice models allowing for continuous state variables. In general, the class of models

we study is not identified for general period return functions. We follow Bajari, Hong and Ryan (2004) and Bajari and Krainer (2004), we demonstrate that if appropriate exclusion restrictions are made on payoffs identification is possible under fairly weak assumptions about the nature of payoffs.

2 The Model.

The model is an infinitely repeated game of incomplete information. We restrict attention to stationary environments and Markov perfect equilibrium.¹ In the model, there are a finite number of players $i = 1, \dots, n$ and an infinite number of discrete time periods $T = 1, \dots, \infty$. During each time period, players simultaneously choose an action $a_i \in \{0, 1, \dots, K\}$ out of a finite set. We restrict players to have the same set of actions for notational simplicity. However, all of our results will generalize to the case where all players have different finite sets of actions. Let $A = \{0, 1, \dots, K\}^n$ denote the vector of possible actions for all players and let $a = (a_1, \dots, a_n)$ denote a generic element of A . As is common in the literature, we shall let $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ denote a vector of strategies for all players excluding player i . There are no mixed strategies since with probability one players will have a unique best response.

Let $s_i \in S_i$ denote the state variable for player i which is common knowledge to all players in the game. The state variable is assumed to be a real valued vector, but unlike most of the previous literature, S_i is not required to be a finite set. Let $S = \prod_i S_i$ and let $s = (s_1, \dots, s_n) \in S$ denote a vector of state variables for all n players. For each agent, there are also $K + 1$ state variables which we label as $\epsilon_i(a_i)$ which are private information. These state variables are distributed i.i.d. across time periods, agents and actions. Let ϵ_i denote the $1 \times (K + 1)$ vector of the individual $\epsilon_i(a_i)$. The density of $\epsilon_i(a_i)$ will be denoted as $f(\epsilon_i(a_i))$, however, we shall sometimes abuse notation and denote the density for $\epsilon_i = (\epsilon_i(0), \dots, \epsilon_i(K))$ as $f(\epsilon_i)$. Let $g(s'|s, a)$ denote the density for the realization of next periods state, s' , conditional on the current state, s , and the vector of actions, a .

The period utility function for player i is:

$$u_i(a, s, \epsilon_i; \theta) = \Pi_i(a_i, a_{-i}, s; \theta) + \epsilon_i(a_i). \quad (1)$$

The utility (1) is identical to commonly used discrete choice models such as the multinomial logit. Player i 's utility is the sum of two terms. The first term, $\Pi_i(a_i, a_{-i}, s; \theta)$ is a deterministic function of the players' actions $a = (a_i, a_{-i})$, the state, $s = (s_1, \dots, s_n)$ which depends on the parameters θ . In the previous literature, $\Pi_i(a_i, a_{-i}, s; \theta)$ has been parameterized as a linear combination of the actions and states. The second term, $\epsilon_i(a_i)$, is i 's private information which is commonly interpreted as an unobserved state variable (see Rust (1994)). In many applications, this will be drawn from an extreme value distribution as in the logit model.

¹Certain aspects of the notation will follow Rust (1994) and Pesendorfer and Schmidt-Dengler (2003).

In what follows, we shall assume that $\Pi_i(a, s_t; \theta)$ is a linear function of θ . Suppose that the “deterministic” part of utility takes the form:

$$\Pi_i(a, s_t; \theta) = \Phi_i(a_i, a_{-i}, s)' \theta \quad (2)$$

where $\Phi_i(a_i, a_{-i}, s) = (\mu_1(a_i, a_{-i}, s), \dots, \mu_l(a_i, a_{-i}, s))$ is a collection of l basis functions and $\Pi_i(a, s_t; \theta)$ is formed as the linear combination of this basis. This assumption may initially seem quite restrictive. We shall impose this restriction for three reasons. The first is that this will generate a considerable savings in terms of both our notation. The second is that in almost all related applications in the literature payoffs are assumed to be linear. Indeed, if the set of basis functions is sufficiently rich, (2) can approximate a continuous utility function arbitrarily well. We invoke the linearity assumption primarily to simplify the exposition of the estimator. In Section 5, we formally discuss extensions of the estimator to the case where utility is nonlinear in θ .

Player i 's decision rule is a function $a_i = \delta_i(s, \epsilon_i)$. Note that the decision rule is not indexed by time because of the Markovian assumption. Also, i 's decision does not depend on ϵ_{-i} , since these shocks are private information to the other players in the game. Define $\sigma_i(a_i|s)$ as:

$$\sigma_i(a_i = k|s) = \int 1 \{ \delta_i(s, \epsilon_i) = k \} f(\epsilon_i) d\epsilon_i. \quad (3)$$

In the above expression, $1 \{ \delta_i(s, \epsilon_i) = k \}$ is the indicator function that player i 's action is k given the vector of state variable (s, ϵ_i) . Therefore, $\sigma_i(a_i = k|s)$ is the distribution of i 's actions conditional on the state variables that are public information. We will define the distribution of a given s as $\sigma(a|s) = \prod_{i=1}^n \sigma_i(a_i|s)$. Next define $\pi_i(a_i, s, \epsilon_i; \theta)$ as:

$$\pi_i(a_i, s, \epsilon_i; \theta) = \sum_{a_{-i}} \Pi_i(a_i, a_{-i}, s, \theta) \sigma_{-i}(a_{-i}|s) + \epsilon_i(a_i) \quad (4)$$

$$\text{where } \sigma_{-i}(a_{-i}|s) = \prod_{j \neq i} \sigma_j(a_j|s). \quad (5)$$

In (4), $\pi_i(a_i, s, \epsilon_i; \theta)$ is player i 's expect utility from choosing a_i when the vector of parameters is θ . Since i does not know the private information shocks, ϵ_j for the other players, his beliefs about their actions are given by $\sigma_{-i}(a_{-i}|s)$.

Players maximize expected discounted utility using the discount factor β . Given a state s and private information ϵ_i , player i 's value function is:

$$W_i(s, \epsilon_i; \sigma_i, \sigma_{-i}) = \max_{a_i \in A_i} \left\{ \pi_i(a_i, s, \epsilon_i; \theta) + \beta \int \sum_{a_{-i}} W_i(s', \epsilon'_i; \sigma_i, \sigma_{-i}) g(s'|s, a_i, a_{-i}) \sigma_{-i}(a_{-i}|s) f(\epsilon'_i) d\epsilon'_i \right\} \quad (6)$$

The optimal choice of a_i depends on the expected period utility, $\pi_i(a_i, s, \epsilon_i; \theta)$ plus the discounted continuation value. Note that the term $\sum_{a_{-i}} g(s'|s, a_i, a_{-i}) \sigma_{-i}(a_{-i}|s)$ is the density for the state variable in the next period given that player i chooses the action a_i today.

Definition: A Markov Perfect Equilibrium is a collection of decision rules $\delta_i(s, \epsilon_i)$, $i = 1, \dots, n$ such that for all i , all s and all ϵ_i , $\delta_i(s, \epsilon_i)$ maximizes $W_i(s, \epsilon_i; \sigma_i, \sigma_{-i})$ where $\sigma_{-i}(a_{-i}|s)$ is given by (3) and (5).

2.1 Expected and Choice-Specific Value Functions.

Following Rust (1994), we will next define the expected and choice-specific value functions. The expected value function, $V_i(s)$ is defined as the expected value of $W_i(s, \epsilon_i)$ marginalizing out ϵ_i . It follows from (6) that:

$$V_i(s) = \sum_{a \in A} \sigma(a|s) \Pi_i(a, s; \theta) + \sum_{k=0}^K E(\epsilon_i(a_i) | a_i = k \text{ is chosen}, s) \sigma_i(a_i = k | s) \\ + \beta \int \sum_{a \in A} V_i(s') g(s'|s, a) \sigma(a|s) ds'.$$

The expected value function is defined recursively as a sum of three terms. The first term, $\sum_{a \in A} \sigma(a|s) \Pi_i(a, s; \theta)$, is the expected value of the deterministic part of the period return function. The second term, $\sum_{k=0}^K E(\epsilon_i(a_i) | a_i = k \text{ is chosen}, s) \sigma_i(a_i = k | s)$, is the expected value of the error term conditional on the chosen action $a_i = k$ and the state s . The third term, $\beta \sum_{a \in A} V_i(s') g(s'|s, a) \sigma(a|s)$, is the expected discounted continuation value.

Let the choice-specific value function, $V_i(a_i, s)$ be defined as:

$$V_i(a_i, s) = \sum_{a_{-i}} \sigma(a_i, a_{-i}) \Pi_i(a_i, a_{-i}, s, \theta) + \beta \int \sum_{a_{-i}} V_i(s') g(s'|s, a) \sigma(a_{-i}|s) ds'$$

The choice-specific value function is the expected utility that the agent received from taking the action a_i in the current period and reverting to the optimal action in the future periods. Note that the choice-specific value function, however, does not include the error term in the period return function $\sum_{k=0}^K E(\epsilon_i(a_i) | a_i = k \text{ is chosen}, s) \sigma_i(a_i = k | s)$.

2.2 Choice Probabilities and Choice-Specific Value Functions.

Arguing as in Rust (1994), it is straightforward to show that the equilibrium in our model must satisfy:

$$\delta_i(s, \epsilon_i) = k \text{ if and only if } V_i(k, s) + \epsilon_i(k) > V_i(k', s) + \epsilon_i(k') \text{ for all } k' \neq k. \quad (8)$$

The choice-specific value function is useful since it allows us to characterize the optimal decision rule in our fully dynamic model in an analogous way as in a static discrete choice model. That is, action k is chosen if and only if the choice-specific value function and error term associated with k is greater than the analogous values for $k' \neq k$.

An implication of (8) is that the equilibrium choice probabilities $\sigma_i(a|s)$ must satisfy:

$$\sigma_i(a_i|s) = \Pr \{ \epsilon_i(a_i) + V_i(a_i, s) - V_i(0, s) > \epsilon_i(k) + V_i(k, s) - V_i(0, s), \forall k = 0, \dots, K, k \neq a_i \} \quad (9)$$

Equation (9) is a simple consequence of (8). The equilibrium probability that action a_i is chosen is precisely the probability that inequality (8) holds. Since the inequalities (8) depend only on the differences between the choice-specific value functions, we can subtract $V_i(0, s)$ from both sides.

For example, suppose that we generate $\epsilon_i(a_i)$ from an extreme value distribution as in the logit model. Then:

$$\sigma_i(a_i|s) = \frac{\exp(V_i(a_i, s) - V_i(0, s))}{\sum_{k=0}^K \exp(V_i(k, s) - V_i(0, s))} \quad (10)$$

A key insight emphasized by Hotz and Miller (1993) is that equation (10) implies that the equilibrium choice probabilities, $\sigma_i(a_i|s)$, have a one-to-one relationship to the choice-specific value functions, $V_i(a_i, s) - V_i(0, s)$. Simple algebra implies that for any k, k' :

$$\log(\sigma_i(k|s)) - \log(\sigma_i(k'|s)) = V_i(k, s) - V_i(k', s)$$

This equation is central to the estimation algorithms proposed by Aguirregabiria and Mira (2002) and Pesendorfer and Schmidt-Dengler (2003). It will also play a crucial role in the algorithm that we discuss below.

The one-to-one mapping between choice probabilities and choice-specific value functions holds more generally than in just the simple case of the logit model. Evidently, the mapping holds whenever the distribution of ϵ_i has full support. We let $\Gamma : \{0, \dots, K\} \times S \rightarrow [0, 1]$ denote the map in general from choice-specific value functions to choice probabilities, i.e.

$$(\sigma_i(0|s), \dots, \sigma_i(K|s)) = \Gamma(V_i(1, s) - V_i(0, s), \dots, V_i(K, s) - V_i(0, s))$$

We will denote the inverse mapping by Ω :

$$(V_i(1, s) - V_i(0, s), \dots, V_i(K, s) - V_i(0, s)) = \Omega_i(\sigma_i(0|s), \dots, \sigma_i(K|s)). \quad (11)$$

2.3 Outline of Estimation Strategy.

Estimation proceeds in four steps. In the first step, the economist estimates $\sigma_i(k|s)$ flexibly using a sieve estimation strategy. In the second step, the economist evaluates equation (11) using the estimated choice probabilities. This generates an estimate of the choice-specific value functions that is consistent with the observed choices in the data. In the third step, given a guess θ of the true value of the utility parameters, θ_0 , the economist evaluates equations similar to (7). This generates choice-specific value functions that are consistent with θ . In the fourth step, the economist then minimizes the distance between the choice-specific value functions found in the second and third steps. We will describe the estimator in detail in what follows below.

2.3.1 First Step: Estimation of Choice Probabilities.

Suppose that the economist has access to time series data on $t = 1, \dots, T$ repetitions of the dynamic game. During each time period, the economist observes the actions and state variables for each agent $(a_{i,t}, s_{i,t})$. In the first step we form an estimate $\hat{\sigma}_i(k|s)$ of $\sigma_i(k|s)$ using sieve series expansions (Ai and Chen (2003)). Let $s = (s^d, s^c)$ denote the discrete and continuous components of s . Also let $\#d$ and $\#c$ denote the dimension of s^d and s^c .

Let $\{q_l(s^c), l = 1, 2, \dots\}$ denote a sequence of known basis functions that can approximate a real valued measurable function of s^c arbitrarily well for a sufficiently large value of l . The sieve could be formed using splines, Fourier Series or orthogonal polynomials (see Chen (2005) for a survey of sieve estimation). We let the basis become increasingly flexible as T becomes large. Let $\kappa(T)$ denote the number of basis functions to be used when the sample size is T . We shall assume that $\kappa(T) \rightarrow \infty$, $\kappa(T)/T \rightarrow 0$ at an appropriate rate to be specified below. Denote the $1 \times \kappa(T)$ vector of basis functions as

$$q^{\kappa(T)}(s^c) = (q_1(s^c), \dots, q_{\kappa(T)}(s^c)), \quad (12)$$

Let s^d be a particular value of the discrete state variables. Define the vector $Q_T(s^d)$ as:

$$Q_T(s^d) = (q^{\kappa(T)}(s_1^c)1(s_1^d = s^d), \dots, q^{\kappa(T)}(s_T^c)1(s_T^d = s^d)).$$

The number of elements in this vector is $\kappa(T)$ times the number of time periods T times the number of discrete values that s^d can assume times $\kappa(T)$. The indicator function $1(s_t^d = s^d)$ is equal to one if the value of the discrete state in the data is equal to s^d and zero otherwise. Thus, the vector $Q_T(s^d)$ is comprised of the individual $q^{\kappa(T)}(s_t^c)$ for $t = 1, \dots, T$ in those periods when the state is s^d and a zero vector of length $\kappa(T)$ in periods t when $s_t^d \neq s^d$.

One potential sieve estimator for $\hat{\sigma}_i(k|s)$, $k = 1, \dots, K$ is a linear probability model, i.e.:

$$\hat{\sigma}_i(k|s) = \sum_{t=1}^T 1(a_{it} = k, s_t^d = s^d) q^{\kappa(T)}(s_t^c) (Q_T(s^d)' Q_T(s^d))^{-1} q^{\kappa(T)}(s^c). \quad (13)$$

Equation (13) is the standard formula for a linear probability model where the regressors are the sieve functions $\kappa(T)$ in equation (12). The sieve estimator $\hat{\sigma}_i(k|s)$ will converge to the true $\sigma_i(k|s)$ at a nonparametric rate which is slower than \sqrt{T} . This results in no loss of generality as long as we define the sieve basis functions properly to include dummy variables that indicate the discrete state variables. In what follows, we shall assume for the ease of exposition that $s = s^c$. Since there are a finite number of discrete states, our rate of convergence and asymptotic theory will be unaffected in the more general case where $s = (s^c, s^d)$. Other link functions could also be used. For example, we could estimate a sieve logit model in the first step. The asymptotic theory is notationally simpler for the case of a sieve linear probability model. However, a similar strategy could

be used to establish our asymptotic results for alternative estimators in the first step.

The first step estimator implicitly abstracts away from problems that might be caused by the multiplicity of equilibrium. If our data consists of a time series for a fixed set of agents, multiplicity is not an issue. However, in many standard applications in empirical industrial organization the data are pooled across markets (see, for example, the entry models of Bresnahan and Reiss (1991,1992), Berry (1992) and Ciliberto and Tamer (2005)). Since it is likely that different markets may be in different equilibria, lack of uniqueness could be an issue. In this case, it is not possible even in principal to recover a single rule $\sigma_i(k|s)$ that describes behavior in all markets.

In practice, there are two potential solutions to the multiplicity problem. The first is to use estimators that explicitly accommodate multiplicity as in Bjorn and Vuong (1984), Bresnahan and Reiss (1990,1991), Berry (1992), Moro (2003), Akerberg and Gowrisankaran (2002), Tamer (2002), Sweeting (2004), Ciliberto and Tamer (2005) or Bajari, Hong and Ryan (2005). Unfortunately, many of these estimators either require the econometrician to be able to analytically characterize certain properties of the equilibrium or to compute the entire equilibrium set. These strategies do not generalize to the class of dynamic models that we are considering.

The second approach to dealing with multiplicity is to assume uniqueness. This is a strong assumption, as even single agent problems can generate multiple solutions. Typically, very strong strict convexity and concavity assumptions are required for uniqueness in many single agent models. The results of Komunjer and Echenique (2005) suggest that many standard regressions, ranging from estimates of production functions to the returns from education are highly problematic if multiplicity is present. The basic idea is that the observed endogenous variables are determined by how the equilibrium is selected given a fixed set of primitives. This source of uncertainty is rarely accounted for in either econometric theory or applied econometrics. While the uniqueness assumption is potentially strong, we note that it is made implicitly in many applications. We shall make this assumption in what follows, but not that accounting for multiplicity important topic for future research.

2.3.2 Second Step: Inversion.

In our second step, we take as given our estimates $\hat{\sigma}_i(k|s)$ of the equilibrium choice probabilities. We then form an estimate of the choice-specific value functions, $\hat{V}_i(k, s_t) - \hat{V}_i(0, s_t)$ for $k = 1, \dots, K$ and $t = 1, \dots, T$. This can be done by evaluating (11) using $\hat{\sigma}_i(k|s)$ in place of $\sigma_i(k|s)$. That is:

$$\left(\hat{V}_i(1, s_t) - \hat{V}_i(0, s_t), \dots, \hat{V}_i(K, s_t) - \hat{V}_i(0, s_t) \right) = \Phi_i(\hat{\sigma}_i(0|s_t), \dots, \hat{\sigma}_i(K|s_t))$$

In the specific case of the logit model, this inversion would simply be:

$$\hat{V}_i(k, s_t) - \hat{V}_i(0, s_t) = \log(\hat{\sigma}_i(k|s_t)) - \log(\hat{\sigma}_i(0|s_t)) \quad (14)$$

In an alternative model, such as one with normal shocks, we would need to solve a nonlinear system.

In our second step, we will also want to compute $\widehat{E}(\epsilon_i(a_i)|a_i \text{ is chosen}, s_t)$ for all i and all t . This is an estimate of the expected value of the error term conditional on the action and the observed value of the state variable. Fix i and s_t and draw $r = 1, \dots, R$ simulated values of $\epsilon_i^{(r)}$ of the stochastic preference shocks. Since the distribution of the preference shocks is known, this can easily be done. Equation (9) implies that given choice-specific value functions $\widehat{V}_i(k, s_t) - \widehat{V}_i(0, s_t)$:

$$a_i^{(r)} = \delta(s_t, \epsilon_i^{(r)}) \text{ if and only if } \epsilon_i^{(r)}(a_i) + \widehat{V}_i(a_i, s) - \widehat{V}_i(0, s) > \epsilon_i^{(r)}(k) + \widehat{V}_i(k, s) - \widehat{V}_i(0, s), \forall k = 0, \dots, K, k \neq a_i.$$

By taking the average value of $\epsilon_i^{(r)}$ when $a_i^{(r)} = \delta(s_t, \epsilon_i^{(r)})$, we can estimate $\widehat{E}(\epsilon_i(a_i)|a_i \text{ is chosen}, s_t)$. In principal, such a simulations could introduce error into our estimator. However, we will assume that the investigator has access to a sufficiently powerful computer to simulate the objects so that this source of error can be effectively ignored.

2.3.3 Third Step: Computation of Choice-Specific Value Function.

In the third step, we find choice-specific value functions that are consistent with a particular value of θ . In what follows, we make the assumption (2) that utility is a linear function of the underlying parameters θ . Define

$$\Lambda_i(s) = \sum_k \widehat{E}(\epsilon_i(a_i)|a_i = k \text{ is chosen}, s) \widehat{\sigma}_i(a_i = k|s), \quad (15)$$

$$\Phi_i(s) = \int \sum_a \Phi_i(a, s) \widehat{\sigma}(a|s) ds, \quad (16)$$

$$\overline{\Lambda}_i(s) = \Lambda_i(s) + \beta E \overline{\Lambda}_i(s'|s), \quad (17)$$

$$\overline{\Phi}_i(s) = \Phi_i(s) + \beta E \overline{\Phi}_i(s'|s). \quad (18)$$

In equation (15), $\Lambda_i(s)$ is the expected value of the error term after we have marginalized out player i 's action a_i . The function $\Phi_i(s)$ is the expected value of the basis functions after we have marginalized out the actions of all players. Equations (17) and (18) has the same recursive structure as a Bellman's equation. The term $E \overline{\Lambda}_i(s'|s)$ is defined as the expected value of $\Lambda_i(s')$ given that the current state is s :

$$E \overline{\Lambda}_i(s'|s) = \int \sum_a \Lambda_i(s') g(s'|s, a) \sigma(a) ds. \quad (19)$$

The term $E \overline{\Phi}_i(s'|s)$ is defined similarly.

Let $V_i^*(a_i, s; \theta)$ denote the choice-specific value function. Using standard arguments in dynamic programming (see Rust (1994)), it can easily be shown that:

$$V_i^*(a_i, s; \theta) = \Phi_i(a_i, s)' \theta + \beta \tilde{\Lambda}_i(a_i, s) + \beta \tilde{\Phi}_i(a_i, s)' \theta, \quad (20)$$

$$\text{where } \Phi_i(a_i, s) = \hat{E} [\Phi_i(a_i, a_{-i}, s) | a_i, s] \quad (21)$$

$$\tilde{\Phi}_i(a_i, s) = \hat{E} [\bar{\Phi}_i(s') | a_i, s] \quad (22)$$

$$\tilde{\Lambda}_i(a_i, s) = \hat{E} [\bar{\Lambda}_i(s') | a_i, s]. \quad (23)$$

In equations (21)-(23) \hat{E} will denote a consistent “estimate” of the objects in parentheses. For instance, $\tilde{\Lambda}_i(a_i, s)$ will denote an estimate of $\bar{\Lambda}_i(s')$ given that the current action is a_i and the state is equal to s . Equation (20) says that the choice-specific value function is equal to the sum of three terms. The first term, $\Phi_i(a_i, s)' \theta$ is the expected value of $\Phi_i(a_i, a_{-i}, s)$ given that player i 's action is a_i and the state is s . The second term, $\beta \tilde{\Lambda}_i(a_i, s)$ is the expected continuation value of future values of the error term in equation (17) given that the current choice is a_i . The final term, $\beta \tilde{\Phi}_i(a_i, s)' \theta$, is the expected discounted value of the deterministic part of utility, given that the current choice is a_i .

We wish to avoid the computational burden of simulating objects similar to (19) or using a contraction mapping to directly solve equations (17)-(18). Therefore, we will use a sieve estimator to estimate the right hand sides of (21)-(23). For example, to estimate the right hand side of (23) we regress $\bar{\Lambda}_i(s_{t+1})$ on $\bar{\Lambda}_i(s_t)$ conditional on the fact that the action at time t is a_i . The fitted value of the regression is our consistent estimate of the right hand sides of (23).

Let $Q_T(a_i)$ be formed in a similar way to Q_T in step one, except that only observations in which the choice of agent i is equal to a_i are used. Instead of having T rows, $Q_T(a_i)$ only has number of rows equal to the total number of observations with realized choice a_i . We then form our estimates by:

$$\hat{E} [\bar{\Lambda}_i(s') | a_i, s] = q^{\kappa(T)}(s)' (Q_T(a_i)' Q_T(a_i))^{-1} \sum_{\tau=1, a_{i\tau}=a_i}^T q^{\kappa(T)}(s_\tau) \bar{\Lambda}_i(s_{\tau+1}). \quad (24)$$

$$\hat{E} [\Phi_i(a_i, a_{-i}, s) | a_i, s] = q^{\kappa(T)}(s)' (Q_T(a_i)' Q_T(a_i))^{-1} \sum_{\tau=1, a_{i\tau}=a_i}^T q^{\kappa(T)}(s_\tau) \Phi_i(a_{i\tau}, a_{-i\tau}, s_\tau). \quad (25)$$

$$\hat{E} [\bar{\Phi}_i(s') | a_i, s] = q^{\kappa(T)}(s)' (Q_T(a_i)' Q_T(a_i))^{-1} \sum_{\tau=1, a_{i\tau}=a_i}^T q^{\kappa(T)}(s_\tau) \bar{\Phi}_i(s_{\tau+1}). \quad (26)$$

Given the above equations, value function iteration is not required to solve for $\bar{\Lambda}_i(s_t)$ or $\bar{\Phi}_i(s_t)$ for $t = 1, \dots, T$. To see why, substituting (24) into equation (17) yields to following system:

$$\bar{\Lambda}_i(s_t) - \beta q^{\kappa(T)}(s_t)' (Q_T(a_i)' Q_T(a_i))^{-1} \sum_{\tau=1, a_{i\tau}=a_i}^T q^{\kappa(T)}(s_\tau) \bar{\Lambda}_i(s_{\tau+1}) = \Lambda_i(s_t).$$

Suppose that we estimate $\Lambda_i(s_t)$ by simulating (15). Substitute this simulation estimate $\widehat{\Lambda}_i(s_t)$ of $\Lambda_i(s_t)$ into the above to yield:

$$\bar{\Lambda}_i(s_t) - \beta q^{\kappa(T)}(s_t)' (Q_T' Q_T)^{-1} \sum_{\tau=1}^{T-1} q^{\kappa(T)}(s_\tau) \bar{\Lambda}_i(s_{\tau+1}) = \widehat{\Lambda}_i(s_t).$$

$$A_T \begin{pmatrix} \bar{\Lambda}_i(s_1) \\ \vdots \\ \bar{\Lambda}_i(s_T) \end{pmatrix} = \begin{pmatrix} \widehat{\Lambda}_i(s_1) \\ \vdots \\ \widehat{\Lambda}_i(s_T) \end{pmatrix}$$

It is straightforward to demonstrate that the T by T matrix A_T is invertible and $\bar{\Lambda}_i(s_t)$ can be solved for uniquely.² By analogous arguments, $\bar{\Phi}_i(s_t)$, $t = 1, \dots, T$ can be computed as the solution to a linear system.

Since both (17) and (18) can be solved as a linear system, it follows that $V_i^*(s_t; \theta) = \bar{\Lambda}_i(s_t) + \bar{\Phi}_i(s_t)' \theta$ for $t = 1, \dots, T$ can be solved for as a linear function of θ . We note that neither $\bar{\Lambda}_i(s_t)$ nor $\bar{\Phi}_i(s_t)$ depend on θ . Thus, while it may initially take some effort to compute these objects, once these $2T$ scalars are computed, we never need to compute them again in order to solve for $V_i^*(s_t; \theta)$!

We note that other strategies are possible for finding the expected choice-specific value functions. For instance, one could simulate continuation values as in Bajari, Benkard and Levin (2003) in order to avoid the need to invert the $T \times T$ matrix. In samples with many observations, this may be a preferable procedure since less memory is required.

2.3.4 Final Step: Minimization.

The final step that we perform is to find a value of θ which minimizes the distance between the choice-specific value functions found in the second and third steps. As before we will use $\widehat{V}_T(s)$ to denote the vector of functions recovered from inverting the vector of conditional choice probabilities, for each s :

$$\widehat{V}_T(s) = \left[\widehat{V}_i(k, s) - \widehat{V}_i(0, s), \quad k = 1, \dots, K, i = 1, \dots, n \right], \quad (27)$$

and let $\widetilde{\Lambda}_T^*(s)$ denote the column vector of functions computed in our third step which we define as:

$$\widetilde{\Lambda}_T^*(s) = \beta \left[\widetilde{\Lambda}_i(k, s) - \widetilde{\Lambda}_i(0, s), k = 1, \dots, K, i = 1, \dots, n \right].$$

Also, define $\widehat{\Phi}_T^*(s)$ by

$$\widehat{\Phi}_T^*(s) = \left[\Phi_i(k, s) - \Phi_i(0, s) + \beta \left(\widetilde{\Phi}_i(k, s) - \widetilde{\Phi}_i(0, s) \right), k = 1, \dots, K, i = 1, \dots, n \right]. \quad (28)$$

²If we examine the coefficient matrix on $\bar{\Lambda}_i(s_t)$, $t = 1, \dots, T$, we will easily see that it is a nonsingular and invertible matrix. Alternatively, we can use an iterative procedure to solve for $\bar{\Lambda}_i(s_t)$ for all t . It is not difficult to see that for each T , this defines a contraction mapping in the integrated square norm $1/T \sum_{t=1}^T \bar{\Lambda}_i(s_t)$, so that the value function iteration procedure will converge to a unique fixed point.

Let $A(s)$ be a $\dim(\theta) \times (n \times K)$ dimension matrix of instruments which is sufficiently rich to identify θ . We can then define our parameter estimate as:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{T} \sum_{t=1}^T A(s_t) \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta \right] = 0. \quad (29)$$

The estimator $\hat{\theta}$ then can be written as

$$\hat{\theta} = \left(\frac{1}{T} \sum_{t=1}^T A(s_t) \hat{\Phi}_T^*(s_t)' \right)^{-1} \frac{1}{T} \sum_{t=1}^T A(s_t) \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) \right].$$

3 Practical Inference

The theory of sieve approximation provides rigorous conditions for controlling the bias term from the fact that the first stage is approximated using a finite number of basis functions. It turns out, as we shall rigorously justify in the next sections, that the statistical properties of our estimator can be performed as if the first stage was estimated parametrically. In this section we describe in details how one can perform practical inference to obtain a consistent estimate of the confidence intervals and standard errors. In the next section, we will discuss the formal theory of semiparametric variances and the required regularity conditions.

In equation (29), a simple approach would be to define our instruments as $A(s_t) = \hat{\Phi}_T^*(s_t)$. This would be equivalent to running a least squares regressions to recover $\hat{\theta}$ by minimizing

$$\frac{1}{T} \sum_{t=1}^T \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta \right]' \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta \right].$$

This is equivalent to choosing the instrument matrix $A(s_t) = \hat{\Phi}_T^*(s_t)$, and one can compute $\hat{\theta}$ analytically by

$$\hat{\theta} = \left(\frac{1}{T} \sum_{t=1}^T \hat{\Phi}_T^*(s_t) \hat{\Phi}_T^*(s_t)' \right)^{-1} \frac{1}{T} \sum_{t=1}^T \hat{\Phi}_T^*(s_t) \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) \right]. \quad (30)$$

In practice, it is possible to use bootstrap or other resampling schemes to obtain consistent standard errors. When T is small, the value function iteration step can be solved by just inverting a $T \times T$ matrix, the computation cost of repeatedly computing $\hat{\theta}$ should not be overly demanding.

In what follows, we describe a method for computing standard errors that does not require resampling. This can potentially save computational time if T is large and difficult to store in the memory of the researcher's computer. By equation (30) it follows that:

$$\sqrt{T} (\hat{\theta} - \theta_0) = \left(\frac{1}{T} \sum_{t=1}^T A(s_t) \hat{\Phi}_T^*(s_t)' \right)^{-1} \frac{1}{\sqrt{T}} \sum_{t=1}^T A(s_t) \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta_0 \right]. \quad (31)$$

By standard arguments, the variance of $\sqrt{T}(\hat{\theta} - \theta_0)$ can be consistently estimated by $\hat{G}_T^{-1}\hat{\Omega}_T\hat{G}_T^{-1'}$, where

$$\hat{G}_T = \left(\frac{1}{T} \sum_{t=1}^T A(s_t) \hat{\Phi}_T^*(s_t)' \right),$$

and $\hat{\Omega}_T$ is a consistent estimate of the variance of

$$\ell(\theta_0) = \frac{1}{\sqrt{T}} \sum_{t=1}^T A(s_t) \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta_0 \right].$$

To obtain $\hat{\Omega}_T$ we examine the structure of $\ell(\theta_0)$. Let $V_T(s_t)$, $\Lambda^*(s_t)$ and $\Phi_T^*(s_t)$ denote the population analog of $\hat{V}_T(s_t)$, $\hat{\Lambda}^*(s_t)$ and $\hat{\Phi}_T^*(s_t)$. First we note that

$$V_T(s_t) - \Lambda^*(s_t) - \Phi_T^*(s_t)' \theta_0 \equiv 0$$

as an identity relation. Then we can rewrite $\ell(\theta_0)$ as

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T A(s_t) \left[\hat{V}_T(s_t) - V_T(s_t) - (\hat{\Lambda}^*(s_t) - \Lambda^*(s_t)) - (\hat{\Phi}_T^*(s_t) - \Phi_T^*(s_t))' \theta_0 \right].$$

The estimation uncertainty in $\ell(\theta_0)$ all comes from the first stage estimation errors in estimating $\hat{V}_T(s_t)$, $\hat{\Lambda}_T^*(s_t)$ and $\hat{\Phi}_T^*(s_t)$. To describe this, we will use $\hat{\Psi}_T(s_t) - \Psi(s_t)$ to denote any one of the elements of

$$\left[\hat{V}_T(s_t) - V_T(s_t), \hat{\Lambda}^*(s_t) - \Lambda^*(s_t), \hat{\Phi}_T^*(s_t) - \Phi_T^*(s_t) \right].$$

In a parametric setup for the first stage where the parametric functional form is assumed to be correctly specified, $\hat{\Psi}_T(s_t) - \Psi(s_t)$ depends on a set of least squares regression coefficients α , such that $\Psi_T^*(s_t) = \Psi_T^*(s_t; \alpha_0)$, and $\hat{\Psi}_T^*(s_t) = \hat{\Psi}_T^*(s_t; \hat{\alpha})$. Each component of $\ell(\theta_0)$ is then approximately distributed as

$$\hat{W}_T \sqrt{T} (\hat{\alpha} - \alpha_0) + o_p(1).$$

where the first component

$$\hat{W}_T = \frac{1}{T} \sum_{t=1}^T A(s_t) \frac{\partial \hat{\Psi}_T^*(s_t; \hat{\alpha})}{\partial \alpha}. \quad (32)$$

can be evaluated by numerical derivatives (The appendix gives more details about computing \hat{W}_T). The variance of $\sqrt{T}(\hat{\theta} - \theta_0)$ can then be estimated by $\hat{G}_T^{-1} \hat{W}_T \hat{\Sigma} \hat{W}_T' \hat{G}_T^{-1'}$, where $\hat{\Sigma}$ is an estimate of the variance of $\sqrt{T}(\hat{\alpha} - \alpha_0)$. Since $\hat{\alpha}$ is just a set of least square regression coefficients, $\hat{\Sigma}$ can easily be computed by either bootstrapping, or Huber-White robust standard error type calculation.

In the later case, let y_t denote the vector of the collection of dependent variables used in these linear regression, such as $\Lambda_i(s_{t+1})$ for all i , and use e_t to denote the vector of residuals in these linear regressions. Then we can estimate

$$\hat{\Sigma} = I \otimes (Q_T' Q_T)^{-1} \frac{1}{T} \sum_{t=1}^T e_t e_t' \otimes q^{\kappa(T)}(s_t) q^{\kappa(T)}(s_t)' I \otimes (Q_T' Q_T)^{-1}.$$

Naturally, one can also estimate Σ by the empirical variance covariance matrix of bootstrapped $\hat{\alpha}$. The next question will be related to the efficient choice of $A_T(s_t)$, the instrument matrix. We can choose $A_T(s_t)$ to minimize the asymptotic variance of $\sqrt{T}(\hat{\theta} - \theta_0)$, which is

$$(E [A_T(s) \Phi_T^*(s)'])^{-1} E \left[A_T(s) \frac{\partial \Psi_T^*(s; \alpha)}{\partial \alpha} \right] \Sigma E \left[\frac{\partial \Psi_T^*(s; \alpha)'}{\partial \alpha} A_T(s)' \right] (E [\Phi_T^*(s) A_T(s)'])^{-1}.$$

With a correctly specified parametric first step α , it is not immediate to simplify the solution for the optimal instruments $A_T(s)$, except in the case where s are all discrete, in which case $A_T(s)$ can be solved by a system of linear equations. As we will see in the parametric case it is easier to obtain the optimal $A_T(s)$ in the general case including continuous state variables s .

4 Semiparametric Variance

The estimator that we consider falls within the class of semiparametric estimators considered by Newey (1994). He demonstrates that, under appropriate regularity conditions, the second stage asymptotic variance will be independent of the particular choice of nonparametric method used to estimate the first stage (e.g. sieve or kernel). This suggests that we can derive the form of the semiparametric asymptotic variance of our estimator that is independent of the nonparametric methods that are used to estimate the choice probabilities. Deriving this semiparametric variance is important because it validates the practical inference methods we described in the previous section. The validity of these parametric inference methods depends on both the knowledge of the limit semiparametric variances, and a set of regularity conditions that require the choice probabilities and transition probabilities to be sufficiently smooth functions of the state variables.

We are concerned with the asymptotic variance of the estimator defined in (31), where the first step parametric estimation is considered to be a sieve approximation that expands as a function of the sample size. Since $\hat{\Phi}_T^*(s_t)$ consistently estimates $\Phi_T^*(s_t)$, it should be intuitively clear that the first component of the Jacobian term converges in probability: $\hat{G}_T \xrightarrow{p} G$, where $G = E [A(s_t) \Phi_T^*(s_t)]$. If we can derive the asymptotic variance Ω of

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T m(s_t, \hat{h}) = \frac{1}{\sqrt{T}} \sum_{t=1}^T A(s_t) \left[\hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta \right]$$

then it follows that $\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, G^{-1}\Omega G^{-1'})$. In the above \hat{h} is used to denote the set of sieve least square projections that are used to estimate h , the set of conditional choice probabilities and conditional transition processes given the state variables, that are used in forming

$$\hat{\Psi}_T^*(s_t) = \hat{V}_T(s_t) - \hat{\Lambda}^*(s_t) - \hat{\Phi}_T^*(s_t)' \theta_0.$$

For this purpose, the following proposition verifies the linearization requirement in Newey (1994) for the dynamic discrete model that we considered:

Proposition 1 *There exists a set of linear functionals $D(s_t, h)$ such that for any parametric sub-path $h(\theta)$ through the space of nonparametric functions h :*

$$\frac{\partial}{\partial \theta} Em(s_t, h(\theta)) = \frac{\partial}{\partial \theta} ED(s_t, h(\theta)),$$

and we can find a set of functions $\delta(s_t)$ such that for all h :

$$ED(s_t, h) = E\delta(s_t)h(s_t).$$

The complete set of functions h , the linear functions $D(z, h)$, and the set of influence functions $\delta(z)$, are given in the next section.

Based on the calculations that lead to this proposition, it then follows immediately from proposition 4 in Newey (1994) that $\Omega = Var(\alpha(s_t))$, where $\alpha(s_t) = \delta(s_t) \otimes (y_t - h(s_t))$, and y_t is the vector of dependent variables used in the first stage linear regressions, and

$$h(s_t) = E(y_t | s_t).$$

In the previous section about practical inference, given that $\hat{\alpha} - \alpha$ is essentially

$$(Q_T' Q_T)^{-1} \sum_{t=1}^T (y_t - h_\alpha(s_t)) \otimes q^{\kappa(T)}(s_t),$$

where $h_\alpha(s_t)$ is the parametric approximation of $h(s_t)$, in (32) we are effectively estimating $\alpha(s_t)$ by the linear square projection of

$$\hat{\alpha}(s_t) = \hat{W}_T (Q_T' Q_T)^{-1} q^{\kappa(T)}(s_t) \otimes (y_t - h_\alpha(s_t)).$$

This is because the \hat{W}_T in (32) is approximately

$$\frac{1}{T} \sum_{t=1}^T D(s_t, q^{\kappa(T)}(\cdot)) = \frac{1}{T} \sum_{t=1}^T \delta(s_t) q^{\kappa(T)}(s_t),$$

therefore the first part $\hat{W}_T (Q_T' Q_T)^{-1} q^{\kappa(T)}(s_t)$ in $\hat{\alpha}(s_t)$, being the fitted value of a least square projection of $\delta(s_t)$ on $q^{\kappa(T)}(s_t)$, should be close to $\delta(s_t)$.

The verification of the semiparametric asymptotic variance indicates that other approximations can be used to estimate the asymptotic variance of $\hat{\theta}$. For example, a kernel based nonparametric regression or other sieve basis functions can be used to consistently estimate $h(s_t)$ and $\delta(s_t)$, which are feasibly given the analytical forms of $\delta(s_t)$ derived in the appendix. They can be chosen based on computational tractability and the plausibility of the results.

4.1 Semiparametric Influence Functions

This subsection derives the linear asymptotic influence functions $D(s_t, h)$ and $\delta(s_t)$ that are used in proposition 1. This subsection presumes familiarity with the arguments in Newey (1994) and can be skipped by those readers interested in implementation issues and not the proof of proposition 1. We will denote the collection of the influence functions as

$$h_J = (h^j, j = 1, \dots, J),$$

where each set of $h^j(s_t)$ corresponds to a step in the nonparametric estimations before the last stage of fitting the parametric utility functions.

The influence function $h^1(s_t)$ The function $h^1(s_t)$ corresponds to the estimation of $\hat{\sigma}_i(k|s)$ in (14). Since the functions $\hat{\sigma}_i(k|s)$ enter the calculations of $\hat{V}_T(s_t)$ through (14) and $\hat{\Lambda}_T^*(s_t)$ through (15), (17), (23) and (27). Hence in correspondence to $\hat{V}_T(s_t)$ and $\hat{\Lambda}_T^*(s_t)$ we can separate

$$D(s_t, h^1) = D_1(s_t, h^1) + D_2(s_t, h^1).$$

Define $\frac{\partial \Phi(\sigma)}{\partial \sigma}(s_t)$ to be the $(K \times n)$ by $(K) \times n$ block diagonal derivative matrix of $\hat{V}_T(s_t)$ with respect to all $\sigma_i(k|s_t)$ of equation (14), and define $\hat{\sigma}(s_t)$ to be the $K \times n$ vector of the collections of all $\hat{\sigma}_i(k|s_t)$ for all i and all k , then

$$D_1(s_t, h^1) = A(s_t) \frac{\partial \Phi(\sigma)}{\partial \sigma}(s_t) \hat{\sigma}(s_t).$$

The corresponding linear influence function is then

$$A(s_t) \frac{\partial \Phi(\sigma)}{\partial \sigma}(s_t) (a_t - \hat{\sigma}(s_t))$$

with $\delta(s_t) = A(s_t) \frac{\partial \Phi(\sigma)}{\partial \sigma}(s_t)$ and a_t is the stacked vector of a_{it} for all $i = 1, \dots, n$.

Next define $\frac{\partial \Lambda(\sigma)}{\partial \sigma}(s_t)$ to be the n by $K \times n$ derivative of (15). Note that $\frac{\partial \Lambda(\sigma)}{\partial \sigma}(s_t)$, (17), (23) and (27) define a (system of) linear functional from $\sigma(s_t)$ to $\hat{\Lambda}_T^*(s_t)$. Denote this linear mapping by $\hat{\Lambda}_T^*(\hat{\sigma}(\cdot))(s_t)$:

$$D_2(s_t, h^1) = A(s_t) \hat{\Lambda}_T^*(\hat{\sigma}(\cdot))(s_t).$$

The linear influence function for $D_2(s_t, h^1)$ will be derived below.

The influence function $h^2(s_t)$ The influence function $h^2(s_t)$ corresponds to the conditional expectation operator in (17) and (18) that are used to compute the fixed point functions $\bar{\Lambda}_i(s)$ and $\bar{\Phi}_i(s)$. The transformations in (17) and (18) are nonlinear but can be linearized as follows. Let Ψ to denote generically either Λ or Φ . We can replace (17) and (18) by

$$\bar{\Psi}_i(s) - \Psi_i^0(s) = \beta \left[\hat{E}(\Psi_i^0(s')|s) - E^0(\Psi_i^0(s')|s) \right] + \beta E \left[{}^0\Psi_i(s') - \Psi_i^0(s')|s \right]. \quad (33)$$

This defines a linear mapping from $\hat{E}(\Psi_i^0(s')|s)$ to $\bar{\Psi}_i(s)$. Together (33) and (23) and (22) define a linear mapping from $\hat{E}(\Psi_i^0(s')|s)$ to $\hat{\Psi}_T^*(s)$, which was stacked in a vector $\hat{\Psi}_i(a_i, s)$. Denote this linear transformation by $\hat{\Psi}_T^*(\hat{E}[\Psi^0(s')|\cdot])(s_t)$:

$$D(s_t, h^2) = A(s_t) \hat{\Psi}_T^* \left(\hat{E}[\Psi^0(s')|\cdot] \right) (s_t).$$

Its corresponding linear influence function is also discussed below.

The influence function $h^3(s_t)$ The influence function $h^3(s_t)$ corresponds to the conditional expectation operator in (22) and (23) that are used to compute the choice-specific expectations. Since these are already linear transformations.

$$D(s_t, h^3) = A(s_t) \hat{E}[\Psi^0(s')|s_t, a_i],$$

and the corresponding linear influence function is therefore

$$A(s_t) \left(\Psi_0(s_{t+1}) - \hat{E}[\Psi^0(s')|s_t, a_i] \right) \quad \text{where} \quad \delta(s_t) = A(s_t).$$

The influence function $h^4(s_t)$ The influence function $h^4(s_t)$ corresponds to the estimation of

$$\hat{E}_{a_{-i}} [\Phi_i(a_i, a_{-i}, s)|a_i, s]$$

in (21). This enters the calculation of $\hat{\Phi}_T^*(s_t)$ directly through its first terms in (28). by (21), and also through the second term in (28) through both (21) the value function iterations (22). Therefore we can correspondingly partition

$$D(s_t, h^4) = D_1(s_t, h^4) + D_2(s_t, h^4),$$

Obviously, the first component can be written as

$$D_1(s_t, h^4) = A(s_t) E_{a_{-t}} [\Phi(a_t, a_{-t}, s_t)|a_t, s_t],$$

where $E_{a_{-t}} [\Phi(a_t, a_{-t}, s_t)|a_t, s_t]$ is the stacked vector of $\hat{E}_{a_{-i}} [\Phi_i(a_i, a_{-i}, s)|a_{it}, s_t]$.

For the second part $D_2(s_t, h^4)$, both (21) and (22) are linear transformations. They define a linear functional from $\hat{E}_{a_{-i}} [\Phi_i(a_i, a_{-i}, s)|a_i, s]$ to $\hat{\Phi}_T^*(s_t)$, which we denote by $\hat{\Phi}_T^* \left(\hat{E}_{a_{-i}} [\Phi_i(a_i, a_{-i}, \cdot)|a_i, \cdot] \right) (s_t)$. Hence,

$$D_2(s_t, h^4) = A(s_t) \hat{\Phi}_T^* \left(\hat{E}_{a_{-i}} [\Phi_i(a_i, a_{-i}, \cdot)|a_i, \cdot] \right) (s_t).$$

The linear influence functions for $D_2(s_t, h^1)$, $D(s_t, h^2)$, and $D_2(s_t, h^4)$. In the above we have explicitly given the asymptotic linear influence function representations for $D_1(s_t, h^1)$, $D(s_t, h^3)$ and $D_1(s_t, h^4)$. We are now left to specify the linear influence functions that correspond to $D_2(s_t, h^1)$, $D(s_t, h^2)$, and $D_2(s_t, h^4)$. These influence functions have a common structure which we exploit now. This common structure begins with a value function contraction mapping:

$$\hat{g}(s_t) - \beta E[\hat{g}(s_{t+1})|s_t] = \hat{h}(s_t).$$

In the second step the forward choice-specific conditional expectation of $g(s_t)$ is computed and instrumented, so that we are concerned with the asymptotic representation of the moment conditions:

$$E[A(s_t)E[\hat{g}(s_{t+1})|a_t = k, s_t]].$$

For ease of exposition, we analyze an equivalent set of moment conditions:

$$\begin{aligned} & E[A(s_t)p(a_{it} = k|s_t)E[\hat{g}(s_{t+1})|a_{it} = k, s_t]] \\ = & E[A(s_t)E[1(a_{it} = k|s_t)\hat{g}(s_{t+1})|s_t]] = E[A(s_t)1(a_{it} = k|s_t)\hat{g}(s_{t+1})]. \end{aligned}$$

In the spirit of Newey (2004), we are looking for a set of functions $\delta(s_t)$ such that we can write

$$E[A(s_t)1(a_{it} = k)\hat{g}(s_{t+1})] = E[\delta(s_t)\hat{h}(s_t)]. \quad (34)$$

To describe $\delta(s_t)$, first we define the function

$$\tilde{A}(s_t) = E[A(s_{t-1})1(a_{it-1} = k)|s_t].$$

Then we can define $\delta(s_t)$ as the unique function solution to the following contraction mapping:

$$\delta(s_t) - \beta E[\delta(s_{t-1})|s_t] = \tilde{A}(s_t). \quad (35)$$

To see why this satisfies (34), note that we can write its left hand side as

$$E[E[A(s_t)1(a_{it} = k)|s_{t+1}]\hat{g}(s_{t+1})] = E[\tilde{A}(s_{t+1})\hat{g}(s_{t+1})],$$

which can also be written $E[\tilde{A}(s_t)\hat{g}(s_t)]$ because of stationarity. Now using the definition of $\delta(s_t)$, this can be written as

$$\begin{aligned} & E[(\delta(s_t) - \beta E[\delta(s_{t-1})|s_t])\hat{g}(s_t)] = E\delta(s_t)\hat{g}(s_t) - \beta E\delta(s_{t-1})\hat{g}(s_t) \\ = & E\delta(s_t)\hat{g}(s_t) - \beta E\delta(s_t)\hat{g}(s_{t+1}) = E\delta(s_t)[\hat{g}(s_t) - \beta E(\hat{g}(s_{t+1})|s_t)]. \end{aligned}$$

The definition of $g(s_t)$ in (35) then verifies that this is equal to $E\delta(s_t)\hat{h}(s_t)$. With this definition of $\delta(s_t)$, we can then write the asymptotic linear representation as

$$\delta(s_t)(y_t - h(s_t))$$

where y_t are the dependent variables used in the construction of $h(s_t) = E(y_t|s_t)$, including $\frac{\partial \Phi(\sigma)}{\partial \sigma}(s_t)a_t$ in $D_2(s_t, h^1)$, $\beta \Psi_i^0(s_{t+1})$ for $D(s_t, h^2)$, and $\Phi_i(a_{it}, a_{-it}, s_t)$ in $D(s_t, h^4)$.

The linearity representation functions also allow us to address the issue of efficient choice of instruments. In general as we show above $\delta(s_t)$ is a linear functional of the instrument functions $A(s_t)$, which we denote $\delta[A(\cdot)](s_t)$. Also let $\Sigma(s_t) = \text{Var}(y_t - h(s_t)|s_t)$ denote the conditional variance matrix of the dependent variables in the conditional expectation calculations. Then we can in general write the asymptotic variance in the form of

$$(EA(s_t)\Phi_T^*(s_t))^{-1} E(\delta[A(\cdot)](s_t)\Sigma(s_t)\delta[A(\cdot)](s_t)') (E\Phi_T^*(s_t)'A(s_t)')^{-1}.$$

The efficient choice of $A(s_t)$ minimizes this asymptotic variance, which will equate

$$(EA(s_t)\Phi_T^*(s_t)) = E(\delta[A(\cdot)](s_t)\Sigma(s_t)\delta[A(\cdot)](s_t)').$$

Alternatively, we can also choose $\delta(s_t)$ efficiently, noting in turn that $A(s_t)$ can be written as a linear functional of $\delta(s_t)$ by inverting the mapping from $A(s_t)$ to $\delta(s_t)$. It is straightforward but tedious to describe the explicit functional form of $A(s_t)$, because $\delta[A(\cdot)](s_t)$ involves both a simple pointwise derivative and a value function iteration. We omit such details.

4.2 Regularity conditions

Newey (1994) provided a set of sufficient conditions that rigorously justify the validity of semiparametric variances and the validity of the use of sieve parametric models to approximate the limiting semiparametric variance. We verify these regularity conditions for the dynamic model that we study.

The required regularity conditions can be broadly classified into two categories. The first category contains conditions on the sieve functions used and the degree of smoothness of the underlying function that is estimated nonparametrically, so that the conditional expectations can be estimated at a sufficiently fast rate of convergence. The second category of regularity conditions requires that the second stage semiparametric parameter is a sufficiently smooth functional of the conditional expectation functions that are being estimated nonparametrically in the first stage.

Assumption 1. For each of the $h(s) = h^j(s)$ conditional expectation function that is being estimated nonparametrically in the first stage,

$$\sup_{s \in \mathcal{S}} |h(s) - h^\kappa(s)| \leq C\kappa^{-\alpha}$$

where $h^\kappa(s) = q^\kappa(s)' (Eq^\kappa(s)q^\kappa(s)')^{-1} Eq^\kappa(s)h(s)$. The class of sieve approximating functions satisfies the following conditions:

$$\sqrt{T}\zeta_0(\kappa)^2 \left[\left(\frac{\kappa}{T} \right) + \kappa^{-2\alpha} \right] \rightarrow 0$$

where $\zeta_0(\kappa) = \sup_{z \in \mathcal{Z}} \|q^\kappa(s)\|$.

This assumption implies that the first stage nonparametric estimation converges to the truth at a rate faster than $T^{1/4}$ because Newey (1994b) showed that under this assumption:

$$\sup_{z \in \mathcal{Z}} |\hat{h}(z) - h(z)| = O_p \left(\zeta_0(\kappa) \left[\sqrt{\frac{\kappa}{T}} + \kappa^{-\alpha} \right] \right).$$

Assumption 2. For each of the $\delta(s) = \delta^j(s)$, define $\delta_\kappa(s) = q^\kappa(s) (E q^\kappa(s) q^\kappa(s)')^{-1} E q^\kappa(s) \delta(s)$. Then

$$TE [|\delta(s) - \delta_\kappa(s)|^2] E [|h(s) - h_\kappa(s)|^2] \rightarrow 0.$$

In addition, $\zeta_0^4 \frac{\kappa}{T} \rightarrow 0$, $E (\delta(s) - \delta_\kappa(s))^2 \rightarrow 0$, and

$$\zeta_0(\kappa)^2 E (h(s) - h_\kappa(s))^2 \rightarrow 0.$$

Newey (1994) showed that this assumption implies that

$$\sqrt{T} ED \left(s_t, \hat{h} - h \right) = \frac{1}{\sqrt{T}} \sum_{t=1}^T \delta(s_t) (y_t - h(z_t)) + o_p(1).$$

Note that both of these two assumptions apply generically to similar two step semiparametric models. Once we made these two assumptions, there is no need to verify them for our model. In contrast, the next two assumptions stated below will need to be verified for the model we consider. The contraction mapping properties turn out to be useful for verifying these assumptions.

Assumption 3. $(\sum_{k=1}^{\kappa} |q_k(s)|_0^2)^{1/2} \left[\left(\frac{\kappa}{T} \right)^{1/2} + \kappa^{-\alpha} \right] \rightarrow 0$, and for each $h = h^j$, $D(s_t, h)$ satisfies

$$|D(s_t, h) - D(s_t, h_0)| \leq b(s_t) \sup_s |h(s) - h_0(s)|. \quad \text{where } Eb(s_t)^2 < \infty.$$

The first part of this assumption is a condition on the sieve functions and is not model specific. In our model, the second part of the above assumption is satisfied as long as the instrument matrix $A(s_t)$ has finite variance, because of the contraction mapping property. Note that in our model $D(s_t, h - h_0)$ typically takes the form of

$$A(s_t) (g(s_t) - g_0(s_t)) \leq A(s_t) \sup_s |g(s) - g_0(s)|,$$

where $(g(s) - g_0(s))$ is defined as the unique fixed point solution to the functional iteration:

$$g(s) - g_0(s) = h(s) - h(s_0) + \beta E [g(s') - g_0(s')] \cdot |s|$$

As long as $\beta < 1$, it follows immediately that

$$\sup_s |g(s) - g_0(s)| \leq \frac{1}{1 - \beta} |h(s) - h_0(s)|.$$

Hence this assumption is satisfied.

The last assumption requires that the sample moment condition used in our estimation procedure is approximated well enough by the $D(s, h)$ that are linear functionals of the first stage conditional expectations $h = h^j$ for all j 's.

Assumption 4. For each $h = h^j$, $\sqrt{T} \sup_s |m(s, h) - m(s, h_0) - D(s, h - h_0)| = o_p(1)$.

To see why this assumption is satisfied in our model, note that $m(s_t, h)$ typically takes the form of $A(s_t)\hat{g}(s_t)$, where $\hat{g}(s_t)$ is the solution to the sample value function iteration:

$$\hat{g}(s) = \hat{h} + \beta \hat{E}(\hat{g}(s')|s).$$

The true $g_0(s)$ solves the population analog

$$g_0(s) = h_0(s) + \beta E(g_0(s')|s).$$

On the other hand, $D(s, h)$ takes the form of $A(s_t)\bar{g}(s_t)$, where $\bar{g}(s)$ solves the following iteration:

$$\bar{g}(s) - g_0(s) = \hat{h}(s) - h_0(s) + \beta \left(\hat{E}(g_0(s')|s) - E(g_0(s')|s) \right) + \beta E(\bar{g}(s') - g_0(s')|s).$$

The difference between $m(s, h)$ and $D(s, h)$ hence is driven by the difference $\hat{g}(s) - \bar{g}(s)$, which can be written as

$$\hat{g}(s) - \bar{g}(s) = \beta \left(\hat{E}(\hat{g}(s') - \bar{g}(s')|s) \right) + \beta \left(\hat{E}(\bar{g}(s') - g_0(s')|s) - E(\bar{g}(s') - g_0(s')|s) \right).$$

Under the smoothness conditions stated in assumption 1, it is not difficult to show that

$$\sup_s \left| \hat{E}(\bar{g}(s') - g_0(s')|s) - E(\bar{g}(s') - g_0(s')|s) \right| = \left(\zeta_0(\kappa) \left[\sqrt{\frac{\kappa}{T}} + \kappa^{-\alpha} \right] \right)^2.$$

and that

$$\beta \hat{E}(\hat{g}(s') - \bar{g}(s')|s) \leq (\beta + o_p(1)) E(\hat{g}(s') - \bar{g}(s')|s).$$

By combining these, we have therefore shown that

$$(1 - \beta + o_p(1)) \sup_s \sqrt{T} |\hat{g}(s) - \bar{g}(s)| \leq \left(\zeta_0(\kappa) \left[\sqrt{\frac{\kappa}{T}} + \kappa^{-\alpha} \right] \right)^2.$$

Hence the last required assumption is verified because of assumption 1 and the above properties of our model.

5 Nonparametric Identification and Alternative Estimators

In the previous sections, we have assumed that the period utility function has a parametric representation $\Pi_i(a, s; \theta)$. Identification in the estimator Section 2 formally required a sufficiently rich set of instruments. In this section, we discuss the problem of identifying the period return function if parametric assumptions are not imposed. Our identification strategy will suggest an alternative set of estimators which we will briefly discuss at the end of this section. Identification of models with discrete state spaces has been discussed by Hotz and Miller (1993), Berry, Pakes, and Ostrovsky (2003), Aguirregabiria and Mira (2002) and Pesendorfer and Schmidt-Dengler (2003). Recent work by Heckman and Navarro (2005) and Aguirregabiria (2005) discuss identification in models with continuous state variables. The arguments we present here are closely related to the ideas of Bajari, Hong and Ryan (2004) and Bajari and Krainer (2004) which propose exclusion restrictions to identify static games.

Formally, we consider the problem of recovering the function $\Pi_i(a_i, a_{-i}, s)$ without specifying it parametrically as $\Pi_i(a_i, a_{-i}, s; \theta)$. We begin by assuming that the econometrician has knowledge of the distribution of the ϵ_i . As we discussed in Bajari, Hong and Ryan (2004), this assumption is required for identification even in a static model. The intuition is straightforward. For instance, a textbook binary probit is a special case of our dynamic model where $\beta = 0$, there are two choices and the error term is normally distributed. We can think of s as the set of covariates. Let $\sigma_1(a_1 = 0|s)$ be the probability that the choice is equal to one in the probit. Then we can rationalize this probability by setting $\Pi_1(a_1, s) = F^{-1}(\sigma_1(a_1 = 0|s))$ where F is the normal c.d.f. Obviously, if knowledge of the distribution of the error term is required for a static, single agent problem, it must be required for our model that nests this as an extremely special case. Furthermore, we assume that the economist has knowledge of the discount factor, β . Rust (1994) discusses why this is not identified even in single agent discrete choice problems.

We break the problem of identification into two steps. In the first step, we seek to identify for all $a_i = 0, \dots, K$:

$$\Pi_i(a_i, s) \equiv \sum_{a_{-i}} \Pi_i(a_i, a_{-i}, s) \sigma_{-i}(a_{-i}|s) \quad (36)$$

In the second step, we ask what restrictions, such as exclusion restrictions, can be used to identify $\Pi_i(a_i, a_{-i}, s)$ from knowledge of $\Pi_i(a_i, s)$.

5.1 Identification of $\Pi_i(a_i, s)$

Identification of the first step of $\Pi_i(a_i, s)$ follows from arguments along the lines of Aguirregabiria and Mira (2002) and Magnac and Thesmar (2002). The basic idea is quite simple. The definition of equilibrium implies that an agent makes a best response to his equilibrium expectations about the actions of the

other agents. By focusing on (36), we are identifying an agent's period utility in equilibrium from choosing an action a_i . This is similar to identifying the structural parameters in a single agent problem:

Recall (7) implies that for each $a_i = 0, \dots, K$:

$$V_i(a_i, s) = \Pi_i(a_i, s) + \beta E [V_i(s') | s, a_i].$$

Furthermore, from equations (7) and (20), we can write

$$V_i(s) = \Lambda_i(s) + \sum_{k=0}^K \sigma_i(k|s) V_i(k, s).$$

The combination of these two relations shows that

$$V_i(a_i, s) = \Pi_i(a_i, s) + \beta E \left[\Lambda_i(s') + \sum_{k=0}^K \sigma_i(k|s') V_i(k, s') | s, a_i \right]. \quad (37)$$

By writing $V_i(k, s') = V_i(k, s') - V_i(0, s') + V_i(0, s')$, we can rearrange this relation so that

$$\begin{aligned} & V_i(a_i, s) - \beta E \left[\Lambda_i(s') + \sum_{k=1}^K (V_i(k, s') - V_i(0, s')) \sigma_i(k|s') | s, a_i \right] \\ &= \Pi_i(a_i, s) + \beta E [V_i(0, s') | s, a_i]. \end{aligned}$$

Next, we impose the normalization that $\Pi_i(0, a_{-i}, s) = 0$ for all a_{-i} . This is similar to the assumption that there is an outside good in a single agent discrete choice model. In an entry model, if 0 corresponded to the decision not to enter a market, then this assumption could be interpreted as the profit from not entering a market being zero. Since this assumption is required for identification in much simpler static models (see Bajari, Hong and Ryan (2004) for a formal argument), it is not surprising that it is also required for identification in dynamic models. We can then rewrite the above equation as:

$$V_i(0, s) - \beta E [V_i(0, s') | s, 0] = \beta E \left[\Lambda_i(s') + \sum_{k=1}^K (V_i(k, s') - V_i(0, s')) \sigma_i(k|s') | s, 0 \right]. \quad (38)$$

Based on arguments similar to Section 2, it is clear that the right hand side can be recovered from the population. The left hand side obviously satisfies Blackwell's sufficient conditions and is a contraction that can be used to recover $V_i(0, s)$ uniquely. Therefore $V_i(0, s)$ is identified. In section 2, we established that there is an inversion between choice probabilities and $V_i(k, s) - V_i(0, s)$ for $k = 1, \dots, K$. Given knowledge of $V_i(0, s)$, we can clearly recover $V_i(k, s)$. Thus, our choice-specific value functions are identified.

Once these quantities are known, $\Pi_i(a_i, s)$ can then be identified since all of the terms on the right hand side of the equation below can be recovered from the population:

$$\Pi_i(a_i, s) = V_i(a_i, s) - \beta E \left[\Lambda_i(s') + \sum_{k=0}^K V_i(k, s') \sigma_i(k|s') | s, a_i \right] \quad (39)$$

5.2 Identification of $\Pi_i(a_i, a_{-i}, s)$

Next we turn to the problem of identifying $\Pi_i(a_i, a_{-i}, s)$ from knowledge of $\Pi_i(a_i, s)$ and $\sigma_{-i}(a_{-i}|s)$:

$$\Pi_i(a_i, s) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s) \Pi_i(a_i, a_{-i}, s), \forall i = 1, \dots, n, a_i = 1, \dots, K. \quad (40)$$

Even with the normalization that $\Pi_i(0, a_{-i}, s) \equiv 0$, it is clear that $\Pi_i(a_i, a_{-i}, s)$ is not identified. Holding the state vector s fixed, determining the utilities of all agents involves solving for $n \times K \times (K+1)^{n-1}$ unknowns. That is, there are n agents, for each action $k = 1, \dots, K$, utility depends on the $(K+1)^{n-1}$ possible actions of the other agents. However, the left hand side of (40) only contains information about $n \times (K+1)$ scalars holding s fixed. It is clearly not possible to invert this system in order to identify $\Pi_i(a_i, a_{-i}, s)$ for all i , all $k = 1, \dots, K$ and all $a_{-i} \in A_{-i}$. Related nonidentification results have been found by Bresnahan and Reiss (1991,1992) and Pesendorfer and Schmidt-Dengler (2003).

Obviously, there must be cross equation restrictions across either i or k in order to identify the system. An obvious way to identify the system is to impose exclusion restrictions. Partition $s = (s_i, s_{-i})$, and suppose $\Pi_i(a_i, a_{-i}, s) = \Pi_i(a_i, a_{-i}, s_i)$ depends only on the subvector s_i . An example of this might be in an entry model. In this type of model the state is usually a vector of productivity shocks. While we might expect the profit of firm i to depend on the entry decisions of other agents, it should not depend on the productivity shocks of other agents. See Bajari and Krainer (2003) and Bajari, Hong and Ryan (2004) for other examples of possible exclusion restrictions that can be used in applications. If such an exclusion restriction is possible, we can then write

$$\hat{\Pi}_i(a_i, s_{-i}, s_i) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s_{-i}, s_i) \hat{\Pi}_i(a_i, a_{-i}, s_i).$$

Clearly, a sufficient identification condition is that for each s_i , there exists $(K+1)^{n-1}$ points in the support of the conditional distribution of s_{-i} given s_i , such that this system of equations form by these $(K+1)^{n-1}$ points given s_i is invertible. In other words, Let $s_{-i}^1, \dots, s_{-i}^{(K+1)^{n-1}}$ denote these points, then identification requires that the matrix

$$\left[\sigma(a_{-i}|s_{-i}^j, s_i), a_{-i} = 1, \dots, (K+1)^{n-1}, j = 1, \dots, (K+1)^{n-1} \right]$$

be nonsingular and invertible. Note that this assumption will be satisfied as long as s_{-i} contains a continuously distributed variable with sufficient variability.

Obviously, there must be cross equation restrictions across either i or k in order to identify the system. An obvious way to identify the system is to impose exclusion restrictions. Partition $s = (s_i, s_{-i})$, and suppose $\Pi_i(a_i, a_{-i}, s) = \Pi_i(a_i, a_{-i}, s_i)$ depends only on the subvector s_i . An example of this might be in an entry model. In empirical studies of entry, the profit of firm i is usually

modeled as a function of i 's entry decision and the entry decision of other firms. The productivity shocks of other firms $-i$ are not typically included in i 's profits. However, the productivity shocks of other firms may influence their own entry decisions. Thus, these generate a set of variables that might be plausibly excluded from s_i but which might enter s . Such exclusion restrictions are often difficult to find in practice and may be controversial. However, they are required for identification when simultaneity is present in many other settings. We next characterize which exclusion restrictions are sufficient for identification in games as well.

The two step identification structure that we discuss above also suggests simple identification conditions for parametric utility models. Suppose we parameterize utilities as $\Pi_i(a_i, a_{-i}, s; \theta)$ such that the normalization constraint $\Pi(0, a_{-i}, s; \theta) \equiv 0$ at all values of θ . Since as in (40), the left hand side $\Pi_i(a_i, s)$ and the choice probabilities $\sigma_{-i}(a_{-i}|s)$ are identified, parametric identification can be stated as requiring that for any $\theta \neq \theta_0$,

$$\sum_{a_{-i}} \sigma_{-i}(a_{-i}|s) \Pi_i(a_i, a_{-i}, s; \theta) \neq \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s) \Pi_i(a_i, a_{-i}, s; \theta_0)$$

for some $a_i = 1, \dots, K$ and for a set of s with positive probabilities.

5.3 Alternative Semiparametric Estimators

The identification procedure discussed above suggests an alternative approach to estimating the model that does not rely on the linearity assumptions that we imposes earlier in the paper. Using the notation developed in the previous sections, we briefly describe the steps involved in constructing this alternative estimator.

Step 1: Estimate choice probabilities as in step 1 in section 2.3.1.

Step 2: Estimate $\hat{V}_i(k, s_t) - \hat{V}_i(0, s_t)$ for $k = 1, \dots, K$, $i = 1, \dots, n$ and $t = 1, \dots, T$ through the inversion step described in section 2.3.2. As an immediate consequence, we can also estimate $\hat{\Lambda}_i(s)$, the expected unobserved utilities conditional on optimal choice of the agents.

Step 3: Use the sample analog of (38) to obtain an estimate of $\hat{V}_i(0, s)$ for all i :

$$\hat{V}_i(0, s) - \beta \hat{E}[\hat{V}_i(0, s')|s] = \beta \hat{E} \left[\hat{\Lambda}_i(s') + \hat{E} \left[(\hat{V}_i(k, s') - \hat{V}_i(0, s')) | s' \right] | s, 0 \right].$$

Consistent estimation of the expectation objects \hat{E} in the above can be obtained by sieve based least square regressions described in section 2.3.3. The value function iteration for $\hat{V}_i(0, s)$ on the left hand side can be obtained by either recursive least square projection that iterates to convergence, or by inverting a $T \times T$ matrix as described in section 2.3.3.

Step 4: Use equation (39) or (37) to recover a nonparametric estimate of the conditional expected per period utility $\hat{\Pi}_i(a_i, s)$.

Step 5: For each θ , use a least square projection to estimate $\hat{\Pi}_i(a_i, s; \theta)$:

$$\hat{\Pi}_i(a_i, s; \theta) = q^{\kappa(T)}(s)' (Q_T(a_i)' Q_T(a_i))^{-1} \sum_{\tau=1, a_i, \tau=a_i}^T q^{\kappa(T)}(s_\tau) \Pi_i(a_{i\tau}, a_{-i\tau}, s_\tau; \theta). \quad (41)$$

Note that when $\Pi_i(a_i, a_{-i}; \theta)$ is specified as a linear function of θ as in (2), we can estimate $\hat{\Pi}_i(a_i, s; \theta)$ by $\hat{\Phi}_i(a_i; s)' \theta$, where the least square projection is applied to the individual index functions:

$$\hat{\Phi}_i(a_i; s) = q^{\kappa(T)}(s)' (Q_T(a_i)' Q_T(a_i))^{-1} \sum_{\tau=1, a_i, \tau=a_i}^T q^{\kappa(T)}(s_\tau) \Phi_i(a_{i\tau}, a_{-i\tau}, s_\tau).$$

Step 6: Minimize a proper norm of the distance between the nonparametric estimate of $\hat{\Pi}_i(a_i, s)$ and the semiparametric estimate of $\hat{\Pi}_i(a_i, s; \theta)$ in the above step 4:

$$\hat{\theta} = \min_{\theta} \|\hat{\Pi}(a, s) - \hat{\Pi}(a, s; \theta), \forall i = 1, \dots, n, \forall a_i = 0, \dots, K\|.$$

where $\hat{\Pi}_T(s)$ and $\hat{\Pi}_T(s, \theta)$ are the vectors of collections $\hat{\Pi}_i(a_i, s)$ and $\hat{\Pi}_i(a_i, s; \theta)$ for all $i = 1, \dots, n$ and $a_i = 0, \dots, K$. As before, we can use the sample weights to specify the norms and use smooth norms so that $\hat{\theta}$ is asymptotically defined by the solution to a set of moment conditions with a $\dim(\theta) \times (n \times (K + 1))$ dimension instrument matrix $A(s)$:

$$\frac{1}{T} \sum_{t=1}^T A(s_t) \left[\hat{\Pi}_T(s_t) - \hat{\Pi}_T(s_t; \theta) \right].$$

For example, a semiparametric nonlinear least square estimator can be defined as

$$\hat{\theta} = \arg \min \frac{1}{T} \sum_{t=1}^T \left[\hat{\Pi}_T(s_t) - \hat{\Pi}_T(s_t; \theta) \right]^2.$$

This will be asymptotically equivalent to the ‘‘IV’’ estimator using the instrument matrix:

$$A(s_t) = \frac{\partial \Pi_T(s_t; \theta_0)}{\partial \theta}.$$

In particular, with the linear in parameter specification of the per period utility function, the estimator for $\hat{\theta}$ can be computed analytically:

$$\hat{\theta} = \left(\frac{1}{T} \sum_{t=1}^T A(s_t) \hat{\Phi}_T(s_t) \right)^{-1} \frac{1}{T} \sum_{t=1}^T A(s_t) \hat{\Pi}_T(s_t).$$

where $\hat{\Phi}_T(s_t)$ is the vector of $\hat{\Phi}_i(a_i, s_t)$ constructed in (41).

Practical inference methods and the asymptotic distribution theory can be written similar to sections 3 and 4. Since they are completely analogous, we do not reproduce the results here. By choosing the instrument matrix $A(s_t)$ efficiently, we can also show that this estimator is as efficient as the previous one we studied in sections 2, 3 and 4.

6 Conclusion

In this paper, we have proposed a semiparametric estimator for dynamic games of incomplete information. The estimator is influenced by earlier work by Pakes, Ostrovsky and Berry (2003), Aguirregabiria and Mira (2002), Pesendorfer and Schmidt-Dengler (2003) and Bajari, Benkard and Levin (2003). However, unlike the earlier literature, the econometric approach that we take exploits the modern theory of semiparametric estimation, particularly sieve estimation (see Ai and Chen (2003) and Chen (2005)) and the theoretical results of Newey (1994). Therefore, we are able to estimate models that allow for a nonparametric first step and continuous state variables. Also, our results clarify the identification of these models for the case of continuous state variables. In general these models are underidentified, however, with appropriate restrictions on payoffs, identification is possible.

7 References.

References

- [1] Akerberg, D., and Gowrisankaran, G. (2003): “Quantifying Equilibrium Network Externalities in the ACH Banking Industry,” University of Arizona and Ohlin School of Business, Washington University at St. Louis Working Paper.
- [2] Aguirregabiria, V. (2005): “Another Look at the Identification of Dynamic Discrete Decision Processes,” Boston University, Department of Economics.
- [3] Aguirregabiria, V., and Mira, P. (2003): “Sequential Simulated-Based Estimation of Dynamic Discrete Games,” Boston University Working Paper.
- [4] Ai, C. and Chen, X. (2003): “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, Econometric Society, vol. 71(6), 1795-1843.
- [5] Bajari, P., Benkard, C.L., and Levin, J. (2003): “Estimating Dynamic Models of Imperfect Competition,” Stanford and Duke University Working Paper.
- [6] Bajari, P., and Krainer, J. (2004): “An Empirical Model of Analysts’ Stock Recommendations: Market Fundamentals, Conflicts of Interest and Peer Effects,” mimeo, Duke University.
- [7] Berry, S. (1992): “Estimation of a Model of Entry in the Airline Industry,” *Econometrica* , 60, 889–917.
- [8] Berry, S., Ovstrovsky, M. and Pakes, A. (2003): “Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry / Exit Examples),” Harvard University Working Paper.
- [9] Bresnahan, T., and P. Reiss (1990): “Entry in Monopoly Markets,” *Review of Economic Studies*, 57, 531–553.
- [10] Bresnahan, T., and P. Reiss (1991): “Empirical Models of Discrete Games,” *Journal of Econometrics*, 48, 57-81.
- [11] Chen, X. “Large Sample Sieve Estimation of Semi-Nonparametric Models forthcoming in Handbook of Econometrics,” Vol. 6, eds J. Heckman and E. Leamer .
- [12] Ciliberto, F., and Tamer, E. (2003): “Market Structure and Multiple Equilibria in Airline Markets,” Princeton and North Carolina State University Working Paper.
- [13] Davis, P. (2005) “Estimation of Quantity Games in the Presence of Indivisibilities and Heterogeneous Firms,” forthcoming in *Journal of Econometrics*.

- [14] Heckman, J., and S. Navarro (2005): “Dynamic Discrete Choice and Dynamic Treatment Effects,” University.
- [15] Hotz, V., and Miller, R. (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *The Review of Economic Studies* 60, 497-531.
- [16] 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.
- [17] Magnac, T., and D. Thesmar (2002): “Identifying dynamic discrete decision processes,” *Econometrica*, 70(2), 801–816.
- [18] Newey, W. (1990): “Efficient Instrumental Variables Estimation of Nonlinear Models,” *Econometrica*, 58(4), 809–837.
- [19] Newey, W. (1994): “The Asymptotic Variance of Semiparametric Estimators,” *Econometrica*, 62, 1349–82.
- [20] Newey, W., and D. McFadden (1994): “Large Sample Estimation and Hypothesis Testing,” in *Handbook of Econometrics*, Vol. 4, ed. by R. Engle, and D. McFadden, pp. 2113–2241. North Holland.
- [21] Pesendorfer, M., and Schmidt-Dengler, P. (2003): “Identification and Estimation of Dynamic Games,” London School of Economics Working Paper.