Nonparametric Identification of Dynamic Games with Discrete and Continuous Choices*

Jason R. Blevins

Duke University

November 28, 2009†

Abstract. This paper considers nonparametric identification of dynamic games of incomplete information in which players make both discrete and continuous choices. Such models are commonly used in applied work in industrial organization where, for example, firms make discrete entry and exit decisions followed by continuous investment decisions. We first review existing identification results for single agent dynamic discrete choice models before turning to single-agent models with an additional continuous choice variable and finally to multi-agent models with both discrete and continuous choices. We provide new conditions for nonparametric identification of the utility function in both cases.

Keywords: dynamic games, dynamic discrete choice, nonparametric identification.

JEL Classification: C5, C14, C73.

1. Introduction

In this paper we present new nonparametric identification results for both single-agent models and games in which players make both discrete and continuous choices. Such models are routinely used, for example, in industrial organization where firms in dynamic oligopoly models typically make discrete entry and exit decisions and continuous investment, pricing, or quantity choices. For example, in the theoretical framework of Ericson and Pakes (1995), each period incumbent firms first decide whether to continue in the industry or exit, and conditional upon continuing, they make a continuous investment decision.

Previous work regarding identification of dynamic structural models has focused primarily on discrete choice models. Identification of single-agent dynamic discrete choice

*I am especially grateful to Han Hong and the Department of Economics at Stanford University for hosting me in July 2008, during which the majority of this paper was written. I also thank Shakeeb Khan and Andrew Sweeting for many useful discussions.

†First draft: July 30, 2008. This revision: 98a63c0.

models has been studied by Rust (1994) and Magnac and Thesmar (2002). Several authors have also considered nonparametric identification of dynamic discrete games. Pesendorfer and Schmidt-Dengler (2007) provide a rank condition for identification in models with discrete state spaces and Bajari, Chernozhukov, Hong, and Nekipelov (2007) show that models with continuous state spaces are identified under an exclusion restriction. Other identification results have been established in related models. Jofre-Bonet and Pesendorfer (2003) consider nonparametric identification of the cost distribution in a dynamic auction game with continuous choices. Heckman and Navarro (2005) consider semiparametric identification of dynamic discrete choice models and dynamic treatment effect models.

Our contribution relative to the existing literature is to establish conditions for the nonparametric identification of both single-agent and multi-agent models in which agents also make continuous choices in addition to the usual discrete choices. Given that identification of discrete choice games has been established, it may not seem surprising at first that models with an additional continuous choice are also identified, since observing a continuous choice should provide more information than a observing a discrete one. However, in the continuous choice framework, for each state, the unknown primitives are infinite-dimensional functionals rather than finite-dimensional vectors. Thus, although more information is available, the objects of interest are much more complex.

We build on the insights of Hotz and Miller (1993), who develop a method for estimating single agent dynamic discrete choice models which is based on a mapping from (observable) conditional choice probabilities to differences (with respect to a normalizing choice) in the choice-specific value function. Bajari et al. (2007) use this idea in a preliminary step in establishing nonparametric identification of dynamic discrete games of incomplete information. They then show that the choice specific value function can be recovered in levels by establishing that the functional operator in the recursive definition of the value function for the normalized choice is a contraction, and therefore has a unique fixed point. The utility function is then identified trivially by definition of the value function.

We follow a similar approach but we must account for an additional layer of complication due to the introduction of a second, continuous choice. We define a *discrete* choice specific value function and show that it can be recovered similarly, in differences, from the conditional choice probabilities. We show that the relevant functional operators are also contractions in the models we consider, and can thus be used to identify the discrete choice specific value function in levels. The utility function can then be identified up to a normalization, using the first order condition implied by agents' optimal choice of the continuous variable. Although we will explicitly only consider models in which all components of the state vector are either serially independent or fully observable to the researcher, in light of recent work by Hu and Shum (2008a,b), our results can also be applied to models with serially correlated unobserved state variables.

This paper proceeds as follows. Section 2 introduces a general modeling framework and some fundamental assumptions. We then turn to nonparametric identification of the structural primitives in specific models. We approach the main result for dynamic games in three steps, each of which adds one level of complexity. Section 3 begins with a discussion of identification of single agent discrete choice models in order to build intuition. Then, we consider single agent models with the addition of a continuous choice in Section 4. Finally, we extend these results to multi-agent dynamic games in Section 5. Section 6 concludes.

2. Framework and Basic Assumptions

We consider a general class of discrete-time dynamic models with N players, indexed by $i = 1, \ldots, N$, over an infinite time horizon $t = 1, 2, \ldots, \infty$. The state of the market at time t can be summarized by a state vector $s_t \in \mathcal{S}$ which is common knowledge to all players and evolves according to a first order Markov process. At the beginning of the period, players observe vectors of private choice-specific shocks $\varepsilon_{it} \in \mathcal{E}_i \subseteq \mathbb{R}^{K+1}$ and simultaneously make discrete choices $d_{it} \in \mathcal{D}_i = \{0, 1, \ldots, K\}$. Next, players observe private shocks $\eta_{it} \in \mathcal{H}_i \subseteq \mathbb{R}$ and simultaneously make continuous choices $c_{it} \in \mathcal{C}_i \subseteq \mathbb{R}$. Let $a_{it} = (d_{it}, c_{it})$ and $\nu_{it} = (\varepsilon_{it}, \eta_{it})$ denote, respectively, the vectors of choices and private shocks, and let a_t denote the vector consisting of the actions of all players at time t. For simplicity, we assume that all players have the same choice sets and that the support of each of the player-specific shocks is identical across players. We occasionally omit the time subscript on variables when the context is clear.

Upon making the choices a_t, each player i receives a payoff $U_i(a_t, s_t, \nu_{it})$ associated with making choice a_{it} in state s_t given that player i's rivals make choices a_{-it}, where in a slight abuse of notation we define $a_{-it} \equiv (a_{1t}, \ldots, a_{i-1,t}, a_{i+1,t}, \ldots, a_{Nt})$. Players are forward looking and discount future payoffs. We assume that players share a common discount factor $\beta \in [0, 1)$. Players choose actions a_{it} in order to maximize their expected discounted future utility which. When the market is state s_t this can be written as

$$\mathrm{E}\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} U_i(a_\tau, s_\tau, \nu_{i\tau}) \,\middle|\, s_t\right],$$

where the expectation is taken over the infinite sequence of actions, states, and private shocks. Note that we have implicitly assumed that the utility functions and state transitions probabilities are time invariant. Combined with the Markov assumption, this implies that agents' optimal decision rules are stationary.

Before proceeding we make several standard assumptions to make the model more tractable (cf. Rust, 1994; Aguirregabiria and Mira, 2009). Rust (1994) imposes a conditional independence assumption which requires ν_t to be conditionally independent of ν_{t-1} given the state s_t. However, we make a stronger assumption that the private shocks are iid, an assumption which is commonly used in practice. For example, many applications assume the discrete choice specific shocks are iid and have a type I extreme value distribution.

Assumption 1 (Private Information)**.** The private shocks ν_{it} are independent across i and t with known distribution $G_i(\cdot \mid s_t)$. Furthermore, each component of ν_{it} is independent, has full support on $\mathbb{R}$, and has a finite first moment.

Additive separability assumptions are standard in discrete choice analysis, both in static and dynamic models. Given the potential addition of a continuous choice, we make the following modified additive separability condition.

Assumption 2 (Additive Separability)**.** The utility function is additively separable in ε_i:

$$U_i(a, s, \nu_i) = u_i(a, s, \eta_i) + \varepsilon_{id_i}.$$

Additionally, if a continuous choice is made, then u_i is further separable as

$$u_i(a, s, \eta_i) = u_{1i}(d, c, s) + u_{2i}(d, c, s)\eta_i$$

where $u_{2i}(d, c, s)$ is a known function. Otherwise, $u_i(a, s, \eta_i)$ is independent of η_i.

For the discrete choice, this assumption is equivalent to the usual additive separability assumption on the discrete choice specific shocks. In models where a continuous choice is made, we make the following additional assumption on the utility function.

Assumption 3 (Monotone Choice)**.** For all i, u_i is twice differentiable and $\partial^2 u_i / \partial c_i \partial \eta_i > 0$.

The monotone choice assumption is standard and will be used for identification in models with continuous choices. This assumption implies that agents' continuous choice policy rules are monotonic in η_i (Bajari, Benkard, and Levin, 2007). Note that we can always use $\tilde{\eta}_i \equiv -\eta_i$ in cases where the policy function is decreasing in η_i.

In the following sections, we analyze several common special cases of the general model described above. In each case, we first describe the model and then provide conditions for nonparametric identification accompanied by constructive proofs which can be used to motivate estimators. Our approach extends previous results for discrete-choice models and so in Section 3 we briefly review existing results in the context of a single-agent discrete choice model. Next, we consider identification in a similar single agent model with the addition of a continuous choice variable in Section 4. Finally, we establish the main identification result for dynamic games with both discrete and continuous choices in Section 5.

3. Single Agent Dynamic Discrete Choice

Single agent dynamic discrete choice models are an important special case of the more general multi-agent model discussed above. These models have a long history in applied microeconomics, beginning with the pioneering work of Miller (1984), Wolpin (1984), Pakes (1986), and Rust (1987). See Eckstein and Wolpin (1989) for a survey of the early literature. Rust (1994) also provides a survey, discusses identification, and develops a general framework for estimating such models. Hotz and Miller (1993) develop two-step methods for estimating these models which are based on first-step estimates of the conditional choice probabilities. Aguirregabiria and Mira (2002) extend this approach to develop a class of nested pseudo-likelihood estimators. Again, our focus is on nonparametric identification and in this section we review certain results of Hotz and Miller (1993) and Bajari et al. (2007) which we build upon in later sections.

Since there is only a single player ($N = 1$) we omit the i subscript from states and payoffs in this section. Furthermore, since there is only a discrete choice we have $\nu = \varepsilon$. Assumption 2 simplifies to the usual additive separability condition. Note that Assumption 3 has no meaning here since there is no continuous choice.

The value function for this model can be expressed recursively as follows:

$$V(s,\varepsilon) = \max_{d\in\mathcal{D}} \left[u(d,s) + \varepsilon_d + \beta \iint V(s',\varepsilon')\, G(d\varepsilon' \mid s')\, P(ds' \mid s,d) \right].$$

Under Assumption 2, following Rust (1994), we also define the choice-specific value function

$$v(d,s) \equiv u(d,s) + \beta \iint V(s',\varepsilon')\, G(d\varepsilon' \mid s')\, P(ds' \mid d,s) \tag{1}$$

which gives the expected discounted utility in the current period and all future periods resulting from choosing d when the current state is s, excluding the iid shock ε_d. Assumption 2 allows us to express this problem in a more compact form resembling a static discrete choice problem with the choice-specific value function playing the role of the period utility function. Let $\sigma(s,\varepsilon)$ denote the agent's optimal choice of d in state (s,ε). Then,

$$\sigma(s,\varepsilon) = \arg\max_{d\in\mathcal{D}} \left[v(d,s) + \varepsilon_d \right].$$

In this setting the ex-ante value function becomes

$$\bar{V}(s) \equiv \mathrm{E}\left[V(s,\varepsilon) \mid s \right] = \mathrm{E}\left[\max_{d\in\mathcal{D}} \left(v(d,s) + \varepsilon_d \right) \middle| \, s \right]. \tag{2}$$

This is social surplus function of McFadden (1981). Using this notation we can rewrite the choice specific value function from (1) as

$$v(d,s) = u(d,s) + \beta \int \bar{V}(s')\, P(ds' \mid s,d). \tag{3}$$

Assuming β is known, the structural primitive of interest here is the utility function $u(d,s)$.

Example 1. A canonical dynamic discrete choice model is the bus engine replacement model of Rust (1987). The state variable s_t is the accumulated mileage of the bus. Each period, a manager must decide whether or not to replace the engine of a bus. The manager pays a cost $u(1,s_t)$ upon replacement, and pays a cost $u(0,s_t)$ when keeping a bus with mileage s_t active. Associated with each choice d is a random component ε_{dt} which represents unobserved costs or benefits incurred in period t. The manager chooses d in order to minimize his expected discounted costs.

The identification argument proceeds as follows. Hotz and Miller (1993) show that there is a one-to-one mapping between the conditional choice probabilities $\Pr(d \mid s)$, which are observable, and differences in the choice-specific value function, $\Delta(d,s) \equiv v(d,s) - v(0,s)$. Without loss of generality, we work in differences with respect to the choice $d = 0$. Bajari et al. (2007) show that this mapping can be used to recover $v(0,s)$ itself, through the use of a contraction mapping in the choice-specific value function. Then, the entire choice-specific value function can be recovered in levels which in turn allows one to recover the utility function, the primary structural primitive of interest, up to a standard normalization. We state the result before proceeding with the argument.

Theorem 1. *Suppose Assumptions 1 and 2 are satisfied. Then the payoff function $u(d,s)$ is nonparametrically identified up to the normalization $u(0,s) = 0$ for all s.*

Proof of Theorem 1. Hotz and Miller (1993) show that there is a one-to-one mapping Ψ from the conditional choice probabilities to differences in the choice-specific value function (1):

$$(v(1,s) - v(0,s), \dots, v(K,s) - v(0,s)) = \Psi\left(\Pr(d = 1 \mid s), \dots, \Pr(d = K \mid s)\right)$$

(see also Rust, 1994, Lemma 3.1). This mapping depends on the distribution G and, given the choice probabilities, it is sufficient to identify the differences $\Delta(k,s) \equiv v(k,s) - v(0,s)$ for each $k = 1, \dots, K$ and $s \in \mathcal{S}$.

For some function $f : \mathcal{D} \times \mathcal{S} \to \mathbb{R}$, define

$$\tilde{H}\left(f(0,s), f(1,s), \dots, f(K,s)\right) \equiv \int_{\mathcal{E}} \max_{d \in \mathcal{D}} [f(d,s) + \varepsilon_d]\, G(d\varepsilon).$$

When f is the utility function in a static discrete choice model, $\tilde{H}$ is McFadden's social surplus function. $\tilde{H}$ has the following additivity property (Rust, 1994, Theorem 3.1):

$$\tilde{H}\left(f(0,s) + \alpha, f(1,s) + \alpha, \dots, f(K,s) + \alpha\right) = \tilde{H}\left(f(0,s), f(1,s), \dots, f(K,s)\right) + \alpha.$$

In particular, when we take $f = v$, we define

$$H(\Delta(1,s), \dots, \Delta(K,s)) \equiv \tilde{H}\left(0, \Delta(1,s), \dots, \Delta(K,s)\right) \tag{4}$$

and note that

$$\begin{aligned} H(\Delta(1,s),\dots,\Delta(K,s)) + v(0,s) &= \tilde{H}\left(0,\Delta(1,s),\dots,\Delta(K,s)\right) + v(0,s) \\ &= \tilde{H}\left(0, v(1,s) - v(0,s),\dots, v(K,s) - v(0,s)\right) + v(0,s) \\ &= \tilde{H}\left(v(0,s), v(1,s),\dots, v(K,s)\right). \end{aligned}$$

We can now write the functional mapping for the choice-specific value function in (3) as a function of $\Delta(d,s)$, $u(d,s)$, and $v(0,s)$:

$$v(d,s) = u(d,s) + \beta \int \left[H\left(\Delta(1,s'),\dots,\Delta(K,s')\right) + v(0,s')\right] \, p(s' \,|\, s,d)\, ds$$

where H is defined above. Again, this function is specific to the distribution G. Note that $p(s' \mid s,d)$ is identified since it is observable. If we normalize $u(0,s) = 0$, then for $d = 0$, the only remaining unknown is the functional $v(0,s)$. Thus, $v(0,s)$ is identified as the unique[1] fixed point to this functional equation. Given $v(0,s)$, we can identify the remainder of the choice-specific value functions in levels since $v(d,s) = \Delta(d,s) + v(0,s)$.

It remains to identify the utility function $u(d,s)$ for $d = 1,\dots,K$. We can express the utility function in terms of the choice-specific value function and other identified quantities as

$$u(d,s) = v(d,s) - \beta \int \left[H\left(\Delta(1,s'),\dots,\Delta(K,s')\right) + v(0,s')\right] \, P(ds' \,|\, s,d)$$

for any d and s. Thus, $u(d,s)$ is also identified up to the normalization $u(0,s) = 0$. ■

Note that this identification result is not in conflict with the non-identification result of Rust (1994) since the utility normalization rules out alternative specifications of the form $\tilde{u}(d,s) = u(d,s) + f(s) - \beta\,\mathrm{E}[f(s') \mid d,s]$.

4. Single Agent Dynamic Discrete-Continuous Choice

In this section we consider the single agent dynamic discrete choice model from the previous section with the addition of a continuous choice. The choice sets are $\mathcal{D}$ for the discrete choice and $\mathcal{C}$ for the continuous choice. The associated random shocks in each period are ε_t and η_t. We assume that the discrete choice is made at the beginning of the period, prior to making the continuous choice and prior to learning the value of η_t.

The value function from the perspective of the beginning of the period is thus

$$V(s,\varepsilon) = \max_d \mathrm{E}_\eta \left\{ \sup_c \left[u(d,c,s,\eta) + \varepsilon_d + \beta\,\mathrm{E}\left[V(s',\varepsilon') \,|\, d,c,s\right]\right] \middle|\, d,s,\varepsilon \right\},$$

[1] See Rust (1994) for a proof that this functional operator is a contraction.

and the corresponding *discrete* choice specific value function can be written

$$v(d,s) = \mathrm{E}_{\eta}\left\{ \sup_{c} \left[u(d,c,s,\eta) + \beta\,\mathrm{E}\left[\bar{V}(s') \mid d,c,s \right] \right] \,\middle|\, d,s \right\}, \tag{5}$$

where we have made use of the ex-ante value function as defined in (2) and the fact that ε_d is independent of η.

Example 2. Timmins (2002) considers the problem of a municipal water utility administrator who chooses the price of water each period. The price may either be zero, or it may be some positive value. Consider a simplified version of the model in which $d_t \in D = \{0,1\}$ represents the decision of whether or not to set the price at zero, and, conditional on not choosing zero, $c_t \in C = \mathbb{R}^+$ represents the choice of the (positive) price. Associated with each discrete choice j is a random shock ε_{ij}. The continuous choice specific shock η_{it} represents unobservables affecting the cost of extracting water in period t.

As with the discrete choice model, we show that the utility primitives of this model are nonparametrically identified up to an obvious normalization. The steps in the proof correspond conceptually to those of the previous proof in that we first use a conditional choice probability inversion to identify differences in the discrete choice specific value function. This is a different function from the previous model in that it represents the value before the continuous choice shock is known. Then, we show that the discrete choice specific value function is identified in levels by showing that it is the unique fixed point to a similar functional equation, appropriately modified to account for the continuous choice. Finally, we can identify the utility function up to a normalization using the first order condition and the fact that the continuous choice is monotonic in the corresponding shock. This result is stated formally in the following theorem.

Theorem 2. *Suppose Assumptions 1–3 are satisfied. Then the payoff function u_1 is nonparametrically identified up to the normalization $u(0,c,s) = 0$ for all c and s.*

Proof of Theorem 2. After defining the discrete choice specific quantities above, the first few steps of the proof are similar to those of the proof of Theorem 1. Note that as before, we can now write the discrete choice probabilities in terms of the discrete choice specific value function as

$$\Pr(d = k \mid s) = \int 1\{k = \arg\max_{d} \left(v(d,s) + \varepsilon_d \right)\} G(d\varepsilon).$$

We can thus use the mapping Ψ from choice probabilities to $\Delta(d,s)$ to identify differences in the discrete choice specific value function for all d and s.

We show that $v(0,s)$ can be identified as before, through the use of a similar contraction mapping. First, note that by definition of the ex-ante value function $\bar{V}(s)$ and the function H, defined as in (4), we have

$$\begin{aligned}
\mathrm{E}_{s'|d,c,s}\,\bar{V}(s') &= \mathrm{E}_{s',\varepsilon'|d,c,s}\left[\max_{d'\in\mathcal{D}}\left(v(d',s')+\varepsilon'_{d'}\right)\right] \\
&= \mathrm{E}_{s'|d,c,s}\left[H\left(\Delta(1,s),\ldots,\Delta(K,s)\right)+v(0,s)\right].
\end{aligned}$$

Using this identity in the expression above for $v(d,s)$ evaluated at $d=0$ yields a functional equation for $v(0,s)$:

$$\begin{aligned}
v(0,s) &= \mathrm{E}_\eta\left\{\sup_c\left[u(0,c,s,\eta)+\beta\,\mathrm{E}_{s'|d=0,c,s}\,\bar{V}(s')\right]\right\}. \\
&= \mathrm{E}_\eta\left\{\sup_c\left[u(0,c,s,\eta)+\beta\,\mathrm{E}_{s'|d=0,c,s}\left[H\left(\Delta(1,s),\ldots,\Delta(K,s)\right)+v(0,s)\right]\right]\right\}.
\end{aligned}$$

Under the normalizations imposed in the hypothesis, everything in this expression is known except for $v(0,s)$. Lemma 1 (see Appendix A) establishes that this mapping is a contraction and thus it identifies $v(0,s)$ as the unique fixed point. We can then use the identity $v(d,s)=\Delta(d,s)+v(0,s)$ to identify $v(d,s)$ for all $d=1,\ldots,K$.

It remains to identify $u(d,c,s,\eta)$, the only remaining unknown. In the case of single agent models with only discrete choices, we were able to identify u directly from the choice-specific value function. With the addition of the continuous choice, however, we cannot simply identify u from $v(d,s)$ as before, due to the presence of the additional private shock η and the $\sup_c$ operator. The monotone choice assumption allows us to overcome the first problem. It guarantees a one-to-one relationship between c and η so that given values of d, s, and c we can infer the value of η. We address the second problem by working with the first-order condition.

First, by the monotone choice assumption (Assumption 3) the policy function $c=\sigma_c(d,s,\eta)$ provides a one-to-one relationship between the private shock η and the continuous choice c conditional on d and s. Let $\eta(d,s,c)\equiv\sigma_c^{-1}(d,s,c)$ denote the inverse mapping. The distribution $F_{c|d,s}$ is identified, since it is observable, and the distribution G of η is known, so we have

$$\begin{aligned}
F_{c|ds}(c\mid d,s) &= \Pr\left(\sigma_c(d,s,\eta)\le c\mid d,s\right) \\
&= \Pr\left(\eta\le\sigma_c^{-1}(d,s,c)\,\middle|\,d,s\right) \\
&= G\left(\sigma_c^{-1}(d,s,c)\right)
\end{aligned}$$

Thus, if we observe the continuous choice c made when the discrete choice is d and the state is s, the value of η must have been

$$\eta=\sigma_c^{-1}(d,s,c)=G^{-1}\circ F_{c|ds}(c\mid d,s).$$

From now on we focus on identifying $u_1(d, c, s)$. The optimal choice of c, given by the policy rule $\sigma_c(d, s, \eta)$, satisfies

$$\sigma_c(d, s, \eta) = \arg\sup_{c} \left[u(d, c, s, \eta) + \phi(d, c, s) \right],$$

where

$$\phi(d, c, s) \equiv \beta \, \mathrm{E}\left[H(\Delta(1, s'), \ldots, \Delta(K, s')) + v(0, s') \mid d, c, s \right]$$

is an identified function. Therefore, c satisfies the corresponding first-order condition

$$\frac{\partial}{\partial c} u(d, c, s, \eta) + \frac{\partial}{\partial c} \phi(d, c, s) = 0.$$

Applying Assumption 2, we have the equivalent condition

$$\frac{\partial}{\partial c} u_1(d, c, s) + \frac{\partial}{\partial c} u_2(d, c, s) \eta + \frac{\partial}{\partial c} \phi(d, c, s) = 0.$$

Rearranging and using the fact that $\eta = \sigma^{-1}(d, s, c)$, we have

$$\frac{\partial}{\partial c} u_1(d, c, s) = -\frac{\partial}{\partial c} u_2(d, c, s) \, \sigma_c^{-1}(d, s, c) - \frac{\partial}{\partial c} \phi(d, c, s).$$

$u_1(d, c, s)$ is now identified (up to a normalizing constant) for each d and s since

$$u_1(d, c, s) = -\int_{\underline{c}}^{c} \left[\frac{\partial}{\partial c} u_2(d, \tilde{c}, s) \, \sigma_c^{-1}(d, s, \tilde{c}) + \frac{\partial}{\partial c} \phi(d, \tilde{c}, s) \right] d\tilde{c} + u_1(d, \underline{c}, s)$$

where $\underline{c} = \inf \mathcal{C}$. ■

5. Dynamic Games with Discrete and Continuous Choices

In this section we consider dynamic games with discrete and continuous choices involving $N > 1$ players. These models are becoming increasingly important in empirical work in applied microeconomics, especially in industrial organization, and many methods have been developed to estimate them. Aguirregabiria and Mira (2007) propose pseudo maximum likelihood estimators for dynamic discrete games and Pakes, Ostrovsky, and Berry (2007) consider two-step method of moments based estimators. Pesendorfer and Schmidt-Dengler (2007) discuss a general class of asymptotic least squares estimators for such games. Bajari et al. (2007) develop simulation-based methods for estimating dynamic games with both discrete and continuous choices based on revealed preference conditions. See Aguirregabiria and Mira (2009) for a survey of this literature. We focus on nonparametric identification of these models.

In models with multiple players, each player's optimal decision depends on the expectations that player holds about the actions of the other players and so we require some sort of equilibrium concept. We assume that players use strategies that are consistent with a Markov perfect equilibrium (MPE). A Markov strategy for player i in this model is a vector-valued mapping $\sigma_i = (\sigma_{di}, \sigma_{ci})$ where $\sigma_{di} : \mathcal{S} \times \mathcal{E}_i \to \mathcal{D}_i$ denotes the discrete choice of player i in each state and $\sigma_{ci} : \mathcal{S} \times \mathcal{D} \times \mathcal{H}_i \to \mathcal{C}_i$ denotes the continuous choice of player i in each state, conditional on the discrete choices d. Player i's beliefs about the strategies of rival players can be represented by a strategy profile $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_N)$. Given these beliefs, when behaving optimally, the present discounted value of player i's payoffs in state (s, ε_i) at the beginning of the period is

$$(6) \qquad V_i(s, \varepsilon_i \mid \sigma_{-i}) = \max_{d_i \in \mathcal{D}_i} \mathrm{E}\left[W_i(d_i, \sigma_{d,-i}, s, \eta_i \mid \sigma_{c,-i}) + \varepsilon_{id_i} \mid d_i, s, \varepsilon_i\right],$$

where

$$(7) \qquad W_i(d_i, d_{-i}, s, \eta_i \mid \sigma_{c,-i}) = \sup_{c_i \in \mathcal{C}_i} \mathrm{E}\left\{u_i(d, c_i, \sigma_{c,-i}, s, \eta_i) + \beta \, \mathrm{E}\left[V_i(s', \varepsilon_i' \mid \sigma_{-i}) \mid d, c, s\right] \,\middle|\, d, c_i, s, \eta_i\right\}.$$

The value function V_i gives the present discounted value of player i's payoffs when facing the discrete decision in state s after the discrete choice specific shocks ε_i are known. The expectation is with respect to ε_{-i}, the discrete choice specific shocks of player i's rivals, and η_i, player i's own continuous choice specific shock. W_i is the value function when facing the continuous choice in state s after η_i is revealed, given that discrete choices d_i and d_{-i} were made. These functions provide a succinct recursive representation of the dynamic decision of players in this model. The expectation is with respect to η_{-i}, the continuous choice specific shocks of player i's rivals.

A *Markov perfect equilibrium* in this model is a strategy profile $(\sigma_1, \ldots, \sigma_N)$ such that for all $i = 1, \ldots, N$ and $s \in \mathcal{S}$, σ_{di} assigns the optimal discrete choice in (6) given beliefs σ_{-i} and σ_{ci} assigns the optimal continuous choice in (7) given beliefs σ_{-i}.

The primitives of the model are the discount factor β, the distribution of private shocks G_i for each i, the utility functions $U_1, \ldots, U_N$, and the state transition kernel $P(ds' \mid s, a)$. The state transition kernel is observable and we assume β and the distributions G_i are known. Thus, we focus on identifying the utility functions.

The following example outlines a model of the type we consider.

Example 3. Consider a dynamic game in which each period firms first choose whether or not to remain in the market (d_{it}) and then choose quantities for competing in a product market. Suppose that there is learning by doing in that firms' marginal costs are decreasing in their past cumulative production. The continuous choice c_{it} is thus firm i's quantity in period t and the choice of quantity now has dynamic implications since more production today results in lower marginal costs in the future. We thus have $\mathcal{D}_i = \{0, 1\}$ and $\mathcal{C}_i = \mathbb{R}^+$

for all i. The state vector is $s_{it} = (x_t, C_{it})$ where x_t is a market-wide state variable and C_{it} is the past cumulative production of firm i which evolves as $C_{i,t+1} = C_{it} + c_{it}$. Suppose for simplicity that C_{i1}, firm i's past cumulative production in the initial period, is known. Let $p(d_t, c_t, s_t)$ denote the inverse demand function. Marginal costs are a function of cumulative production so the utility (profit) function (conditional on continuing) is

$$u_i(d_{it} = 1, d_{-i}, c_t, s_t, \eta_i) = c_{it}\left[p(d_t, c_t, s_t) - (\mu - \theta C_{it} + \eta_i)\right] + \varepsilon_{it1}$$

where μ is the baseline marginal cost, θ is the amount by which an additional unit of past cumulative production decreases the marginal cost, and η_i is a private marginal cost shock. ε_{it1} is the private choice-specific shock for continuing. Upon exit ($d_{it} = 0$), firms receive a scrap payment ϕ and a random shock ε_{it0}.

With multiple players, we require an exclusion restriction to identify the utility function. Namely, we assume that the state vector can be written as $s = (s_1, \ldots, s_N)$ where s_i are the relevant state variables for player i.

Assumption 4 (Exclusion Restriction)**.** The utility function for player i satisfies

$$u_i(a, s, \eta_i) = u_i(a, s_i, \eta_i)$$

where s_i has at least one continuous component for each i.

Thus, some components of s_i may be common to all players but there must be at least one continuous component that is specific to player i. A similar exclusion restriction was used for identification purposes in dynamic discrete choice games by Bajari et al. (2007). Our condition is slightly different in that we require that there be at least one continuous component of s_i for each player. Note that the utility function in Example 3 satisfies this assumption.

Theorem 3. *Suppose Assumptions 1–4 are satisfied. Furthermore, suppose that both u_i and $\partial u_i / \partial c_i$ are continuous for all i and that the conditional expectation operator $\mathrm{E}_{c_{-i}|d,s}$ is one-to-one. Then the payoff function u_i is nonparametrically identified up to the normalization $u_i(0, d_{-i}, c_i, c_{-i}, s_i) = 0$.*

Proof of Theorem 3. For simplicity, we drop the strategy profile notation and simply treat d_{-i} and c_{-i} as random variables distributed according to the appropriate equilibrium beliefs. We can define the discrete choice specific value function as

$$v_i(d_i, s) \equiv \mathrm{E}_{\eta_i|d_i,s}\, \mathrm{E}_{d_{-i}|d_i,s} \sup_{c_i} \left\{ \mathrm{E}_{c_{-i}|d,c_i,s}\left[u_i(d_i, d_{-i}, c_i, c_{-i}, s, \eta_i) + \beta\, \mathrm{E}_{s'|d,c,s}\, \bar{V}_i(s')\right]\right\}$$

where $\bar{V}_i(s) = \mathrm{E}\left[V_i(s, \varepsilon_i) \mid s\right]$ is the ex-ante value function for player i. If we define H as in (4), then

$$\bar{V}_i(s) = \mathrm{E}\left[\max_{d_i} \left\{v_i(d_i, s) + \varepsilon_{id_i}\right\} \middle|\, s\right] = H(\Delta_i(1, s), \ldots, \Delta_i(K, s)) + v_i(0, s).$$

Then, for $d_i = 0$, after normalizing $u_i(0, d_{-i}, c_i, c_{-i}, s, \eta_i) = 0$, we have

$$
\begin{aligned}
v_i(0,s) &= \mathrm{E}_{\eta_i|d_i,s}\,\mathrm{E}_{d_{-i}|d_i,s} \sup_{c_i} \left\{ \beta\, \mathrm{E}_{c_{-i}|d,c_i,s} \left[\mathrm{E}_{s'|d,c,s}\, \bar{V}_i(s') \right] \right\} \\
&= \mathrm{E}_{\eta_i|d_i,s}\,\mathrm{E}_{d_{-i}|d_i,s} \sup_{c_i} \left\{ \beta\, \mathrm{E}_{c_{-i}|d,c_i,s}\, \mathrm{E}\left[H(\Delta_i(1,s'),\dots,\Delta_i(K,s')) + v_i(0,s') \mid d,c,s \right] \right\}.
\end{aligned}
$$

Note that all quantities except the function $v_i(0,s)$ are known, including β, $\Delta_i(d,s)$, the distributions of η_i, $d_{-i} \mid d_i, s$, and $c_{-i} \mid d, c_i, s$, and the transition density of $s' \mid d, c, s$. As established by Lemma 2 (see Appendix A), this defines a contraction mapping which identifies $v_i(0,s)$ and in turn, knowing $v_i(0,s)$ allows us to identify $v_i(d_i,s) = \Delta_i(d_i,s) + v_i(0,s)$ for all $d_i > 0$.

We now turn to identifying the utility function u. For any d, s, and η_i, the optimal choice of c_i maximizes

$$
\mathrm{E}_{c_{-i}|d,c_i,s}\left[u_{i1}(d_i, d_{-i}, c_i, c_{-i}, s) + u_{i2}(d_i, d_{-i}, c_i, c_{-i}, s)\eta_i \right] + \phi(d, c_i, s)
$$

where $\phi(d, c_i, s)$ is the known function

$$
\phi(d, c_i, s) \equiv \beta\, \mathrm{E}_{c_{-i}|d,c_i,s}\, \mathrm{E}_{s'|d,c,s}\left[H(\Delta_i(1,s'),\dots,\Delta_i(K,s')) + v_i(0,s') \right].
$$

Therefore, c_i must satisfy the first order condition

$$
\frac{\partial}{\partial c_i} \int \left[u_{i1}(d_i, d_{-i}, c_i, c_{-i}, s) + u_{i2}(d_i, d_{-i}, c_i, c_{-i}, s)\eta_i \right] p(c_{-i} \,|\, d, s)\, dc_{-i} + \frac{\partial}{\partial c_i}\phi(d, c_i, s) = 0
$$

Here, we have used the fact that $p(c_{-i} \mid d, c_i, s)$ does not depend on c_i conditional on s and d since the relevant shocks are iid. Under the maintained assumptions, we can interchange the order of integration and differentiation:

$$
\int \left[\frac{\partial}{\partial c_i} u_{i1}(d_i, d_{-i}, c_i, c_{-i}, s) + \frac{\partial}{\partial c_i} u_{i2}(d_i, d_{-i}, c_i, c_{-i}, s)\eta_i \right] p(c_{-i} \,|\, d, s)\, dc_{-i} + \frac{\partial}{\partial c_i}\phi(d, c_i, s) = 0
$$

The partial derivatives of ϕ and u_{i2} are identified since both are known functions. Furthermore, due to the monotone choice assumption we can replace η_i with $\sigma_{ci}^{-1}(d, s, c_i)$, also a known quantity.

$\mathrm{E}_{c_{-i}|d,s}$ is one-to-one by assumption so we can apply the inverse operator to the right-hand side above to recover u_i, however we need to use the exclusion restriction $u_i(d_i, d_{-i}, c_i, c_{-i}, s) = u_i(d_i, d_{-i}, c_i, c_{-i}, s_i)$ if we hope to recover a *function* of d_{-i} and c_{-i}. The conditional expectation operator maps functions of $(d_i, d_{-i}, c_i, c_{-i}, s_i)$ (such as the utility function) to functions of (d_i, d_{-i}, c_i, s). The inverse operator reverses this mapping so that for fixed values of (d, c_i, s_i) it maps functions of s_{-i} to functions of c_{-i}.

Now for each (d, c_i, s_i) we can apply the inverse operator $\mathrm{E}^{-1}_{c_{-i}|d,s}$ to $-\frac{\partial}{\partial c_i}\phi(d, c_i, s)$ to recover

$$\frac{\partial}{\partial c_i} u_{i1}(d_i, d_{-i}, c_i, c_{-i}, s_i) + \frac{\partial}{\partial c_i} u_{i2}(d_i, d_{-i}, c_i, c_{-i}, s_i)\, \sigma^{-1}_{ci}(d, s_i, c_i).$$

Integrating this with respect to c_i and using the fact that we know u_{i2} and $\sigma^{-1}_{ci}(d, s_i, c_i)$ identifies u_{i1} up to a constant normalization for each c_i. ■

Intuitively, the assumption that the conditional expectation operator $E_{c_{-i}|d,s}$ is one-to-one requires there to be sufficient variation in the conditional distribution of c_{-i} for different values of d and s. Assumptions of this type have been used by Hu and Shum (2008a,b) in identifying the Markov kernel in dynamic discrete choice models and dynamic games with unobserved state variables. They are also key conditions for identification in many nonparametric econometric models such as instrumental regression models (Newey and Powell, 2003; Darolles, Florens, and Renault, 2007; Blundell, Chen, and Kristensen, 2007) and classical measurement error models (Chen and Hu, 2006; Hu and Schennach, 2008).

6. Conclusion

We have established conditions for nonparametric identification of both single agent models and dynamic games of incomplete information in which agents make both discrete and continuous choices. Our nonparametric identification results can serve as a point of reference for practitioners estimating parametric models. In the absence of conditions for parametric identification, our conditions, though likely stronger than necessary for highly parametrized models, can serve as a benchmark. Furthermore, our proofs are constructive and suggest the possibility of nonparametric or semiparametric estimators for such models.

A. Auxiliary Results

Lemma 1. *Under the assumptions of Theorem 2, the functional mapping*

$$v(0, s) = \mathrm{E}_\eta \left\{ \sup_c \left[u(0, c, s, \eta) + \beta\, \mathrm{E}_{s'|d=0,c,s} \left[H\left(\Delta(1, s), \ldots, \Delta(K, s)\right) + v(0, s) \right] \right] \right\}$$

is a contraction with modulus β.

Proof of Lemma 1. First we simplify the notation, defining $w(s) \equiv v(0, s)$ and

$$\Gamma(w)(s) \equiv \sup_c \left[\psi(0, c, s) + \beta\, \mathrm{E}_{s'|d=0,c,s}\, w(s') \right],$$

where

$$\psi(d,c,s) \equiv \beta\, \mathrm{E}_{s'|d,c,s}\left[H\left(\Delta(1,s),\ldots,\Delta(K,s)\right)\right]$$

Note that in light of the utility normalization, this expression no longer depends on η and so we can drop the outer expectation. Furthermore, under the assumptions of Theorem 2, the function $\psi(d,c,s)$ is identified.

We must show that for any two functions w and $\tilde{w}$, $\|\Gamma w - \Gamma\tilde{w}\| \leq k\,\|w-\tilde{w}\|$ for some $0<k<1$ where $\|\cdot\|$ is the sup norm, $\|f\| \equiv \sup_{s\in\mathcal{S}} |f(s)|$. We have

$$\begin{aligned}
\|\Gamma w - \Gamma\tilde{w}\| &= \sup_s |\Gamma w(s) - \Gamma\tilde{w}(s)| \\
&= \sup_s \left|\sup_c \left[\psi(0,c,s) + \beta\, \mathrm{E}_{s'|d=0,c,s}\, w(s')\right] - \sup_c \left[\psi(0,c,s) + \beta\, \mathrm{E}_{s'|d=0,c,s}\, \tilde{w}(s')\right]\right| \\
&\leq \sup_s \sup_c \left|\left[\psi(0,c,s) + \beta\, \mathrm{E}_{s'|d=0,c,s}\, w(s')\right] - \left[\psi(0,c,s) + \beta\, \mathrm{E}_{s'|d=0,c,s}\, \tilde{w}(s')\right]\right| \\
&= \beta \sup_s \sup_c \left|\mathrm{E}_{s'|d=0,c,s}\left[w(s') - \tilde{w}(s')\right]\right| \\
&\leq \beta \sup_s \sup_c \mathrm{E}_{s'|d=0,c,s}\left|w(s') - \tilde{w}(s')\right| \\
&\leq \beta \sup_{s'} \left|w(s') - \tilde{w}(s')\right| \\
&= \beta\,\|w - \tilde{w}\|\,.
\end{aligned}$$

The first two equalities follow by definition while the third line follows from the properties of the supremum: for real valued functions f and g $|\sup f(x) - \sup g(x)| \leq \sup|f(x)-g(x)|$. The next two lines follow from properties of the integral: we know that $\left|\int f\right| \leq \int |f|$ and that for any measure μ, $\int_E f\,d\mu \leq \mu(E)\sup_{x\in E} f(x)$. The last equality holds by definition of the norm. ■

Lemma 2. *Under the assumptions of Theorem 3, the functional mapping*

$$v_i(0,s) = \mathrm{E}_{\eta_i|d_i=0,s}\, \mathrm{E}_{d_{-i}|d_i=0,s} \sup_{c_i} \left\{\beta\, \mathrm{E}_{c_{-i}|d,c_i,s}\, \mathrm{E}\left[H(\Delta_i(1,s'),\ldots,\Delta_i(K,s')) + v_i(0,s') \mid d,c,s\right]\right\}.$$

is a contraction with modulus β.

Proof of Lemma 1. For simplicity, define $w(s) \equiv v(0,s)$ and

$$\Gamma(w)(s) \equiv \mathrm{E}_{\eta_i|d_i=0,s}\, \mathrm{E}_{d_{-i}|d_i=0,s} \sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \left\{\psi(0,c,s) + \beta\, \mathrm{E}_{s'|d,c,s}\, w(s')\right\},$$

where

$$\psi(d,c,s) \equiv \beta\, \mathrm{E}_{c_{-i}|d,c_i,s}\, \mathrm{E}_{s'|d,c,s}\left[H\left(\Delta(1,s),\ldots,\Delta(K,s)\right)\right]$$

Note that in light of the utility normalization, this expression no longer depends on η_i and so we can drop the outermost expectation. Furthermore, under the assumptions of Theorem 3, the function $\psi(d, c, s)$ is identified.

We must show that for any two functions w and $\tilde{w}$, $\|\Gamma w - \Gamma \tilde{w}\| \leq k \|w - \tilde{w}\|$ for some $0 < k < 1$ where $\|\cdot\|$ is the sup norm, $\|f\| \equiv \sup_{s \in \mathcal{S}} |f(s)|$. We have

$$\begin{aligned}
\|\Gamma w - \Gamma \tilde{w}\| &= \sup_s |\Gamma w(s) - \Gamma \tilde{w}(s)| \\
&= \beta \sup_s \left| \mathrm{E}_{d_{-i}|d_i=0,s} \left[\sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \mathrm{E}_{s'|d,c,s} w(s') - \sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \mathrm{E}_{s'|d,c,s} \tilde{w}(s') \right] \right| \\
&\leq \beta \sup_s \mathrm{E}_{d_{-i}|d_i=0,s} \left| \sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \mathrm{E}_{s'|d,c,s} w(s') - \sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \mathrm{E}_{s'|d,c,s} \tilde{w}(s') \right| \\
&\leq \beta \sup_s \mathrm{E}_{d_{-i}|d_i=0,s} \sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \mathrm{E}_{s'|d,c,s} \left| w(s') - \tilde{w}(s') \right| \\
&\leq \beta \sup_s \mathrm{E}_{d_{-i}|d_i=0,s} \sup_{c_i} \mathrm{E}_{c_{-i}|d,c_i,s} \sup_{s'} \left| w(s') - \tilde{w}(s') \right| \\
&\leq \beta \sup_{s'} \left| w(s') - \tilde{w}(s') \right| \\
&= \beta \|w - \tilde{w}\|.
\end{aligned}$$

The first equality follows by definition of the norm. The second follows since ψ is known and by the linearity of $\mathrm{E}_{d_{-i}|d_i=0,s}$. The remaining inequalities from properties of the integral, the uniform continuity of sup, and the definition of the norm. ■

References

Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica 70*, 1519–1543. [5]

Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica 75*, 1–53. [10]

Aguirregabiria, V. and P. Mira (2009). Dynamic discrete choice structural models: a survey. *Journal of Econometrics*. [4, 10]

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica 75*, 1331–1370. [4, 5, 6, 10, 12]

Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2007). Semiparametric estimation of dynamic games of incomplete information. Unpublished manuscript, University of Minnesota. [2]

Blundell, R., X. Chen, and D. Kristensen (2007). Semi-nonparametric iv estimation of shape-invariant engel curves. *Econometrica 75*, 1613–1669. [14]

Chen, X. and Y. Hu (2006). Identification and inference of nonlinear models using two samples with aribrary measurement errors. Discussion Paper 1590, Cowles Foundation. [14]

Darolles, S., J.-P. Florens, and E. Renault (2007). Nonparametric instrumental regression. Unpublished manuscript, University of North Carolina, Chapel Hill. [14]

Eckstein, Z. and K. I. Wolpin (1989). The specification and estimation of dynamic stochastic discrete choice models: A survey. *The Journal of Human Resources 24*, 562–598. [5]

Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economics and Statistics 62*, 53–82. [1]

Heckman, J. J. and S. Navarro (2005). Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics 136*(2), 341–396. [2]

Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies 60*, 497–529. [2, 5, 6]

Hu, Y. and S. M. Schennach (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica 76*, 195–216. [14]

Hu, Y. and M. Shum (2008a). Identifying dynamic games with serially-correlated unobservables. Unpublished manuscript, Johns Hopkins University. [3, 14]

Hu, Y. and M. Shum (2008b). Nonparametric identification of dynamic models with unobserved state variables. Unpublished manuscript, Johns Hopkins University. [3, 14]

Jofre-Bonet, M. and M. Pesendorfer (2003). Estimation of a dynamic auction game. *Econometrica 71*, 1443–1489. [2]

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica 70*, 801–816. [2]

McFadden, D. (1981). Econometric models of probabilistic choice. In C. F. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data*. Cambridge, MA: MIT Press. [5, 6]

Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political Economy 92*(6), 1086–1120. [5]

Newey, W. K. and J. L. Powell (2003). Instrumental variable estimation of nonparametric models. *Econometrica 71*, 1565–1578. [14]

Pakes, A. (1986). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica 54*, 755–784. [5]

Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics 38*, 373–399. [10]

Pesendorfer, M. and P. Schmidt-Dengler (2007). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies 75*, 901–928. [2, 10]

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*, 999–1013. [5, 6]

Rust, J. (1994). Structural estimation of Markov decision processes. In R. F. Engle and D. L. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, Amsterdam. North Holland. [2, 4, 5, 6, 7]

Timmins, C. (2002). Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid west. *Econometrica 70*, 603–629. [8]

Wolpin, K. I. (1984). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy 92*(5), 852–874. [5]