

Game Theory and Econometrics: A Survey of Some Recent Research

Patrick Bajari

University of Minnesota and NBER

Han Hong

Stanford University

Denis Nekipelov

University of California at Berkeley *

December 3, 2010

Abstract

We survey an emerging literature on the econometric analysis of static and dynamic models of strategic interactions. Econometric methods of identification and estimation allow researcher to make use of observed data on individual choice behavior and on the conditional transition distribution of state variables to recover the underlying structural parameters of payoff functions and discount rates nonparametrically without imposing strong functional form assumptions. We also discuss the progress that the literature has made on understanding the role of unobserved heterogeneity in the estimation analysis of these models, and other related issues.

1 Introduction

In this paper, we shall survey an emerging literature at the intersection of industrial organization, game theory, and econometrics. In theoretical models in industrial organization, game theory is by far the most common tool used to model industries. In such models, the researcher specifies a set of players, their strategies, information, and payoffs. Based on these choices, the researcher can use equilibrium concepts to derive positive and normative economic predictions. The application of game theory to industrial organization has spawned a large and influential theoretical literature (see Tirole (1988) for a survey). Game theory can be used to model a very broad set of economic problems. However, this flexibility has sometimes proved problematic for researchers. The predictions of game theoretic models

*This survey article draws heavily from a related paper that is also coauthored with Victor Chernozhukov and Denis Nekipelov. We acknowledge insightful discussion and comments by Martin Pesendorfer. The manuscript has been substantially revised to account for these comments.

often depend delicately on the specification of the game. Researchers may not be able to agree, a priori, on which specification is most reasonable and theory often provides little guidance on how to choose between multiple equilibria generated by a particular game.

The literature that we shall survey attempts to address these problems by letting the data tell us the payoffs which best explain observed behavior. In the literature we survey, the econometrician is assumed to observe data from plays of a game and exogenous covariates which could influence payoffs or constraints faced by the agent. The payoffs are specified as a parametric or nonparametric function of the actions of other agents and exogenous covariates. The estimators that we discuss "reverse engineer" payoffs to explain the observed behavior. Over the past decade, researchers have proposed methods to estimate games for a diverse set of problems. These include static games where agents choose among a finite number of alternatives (see Seim (2006), Sweeting (2009), Bajari, Hong, Krainer, and Nekipelov (2006)), dynamic games where the choice set of agents is discrete (see Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2003), Pesendorfer, Schmidt-Dengler, and Street (2008), Bajari, Chernozhukov, Hong, and Nekipelov (2009)) and dynamic games with possibly continuous strategies (see Bajari, Benkard, and Levin (2007)).

While these methods have studied a diverse set of problems, most have relied on a common insight. Games can be quite complicated objects and it may take weeks of computational time to compute even a single equilibrium to dynamic games in particular (see Benkard (2004), Doraszelski and Judd (2007), and Ericson and Pakes (1995)). A brute force approach which repeatedly computes the equilibrium for alternative parameter values will not be computationally feasible for many models of interest. Researchers realized that the numerical burden of estimation would be lessened if estimation was broken into two steps. In a first step, the reduced form to the game is estimated using a flexible method. The reduced form is an econometric model of how the choice of an agent depends on exogenous or predetermined variables. In many models of interest, this boils down to estimating canonical models from applied econometrics. The second step attempts to recover the structural parameters of the model, that is, how payoffs depend on actions and control variables. As we shall discuss in the paper, if we condition on the reduced forms, estimation of the structural parameters can be viewed as a single agent problem. In the simplest static case, the second stage will be no more complicated than estimating McFadden's conditional logit. In the dynamic case, estimation can be performed using very well understood methods from single agent dynamic discrete choice (see Rust (1994)).

An advantage of this approach is that it allows the economist to ground the specification of the game in the data, rather than on the prior beliefs of the economist of what is reasonable. To be clear, we are not claiming that these methods are a substitute for traditional approaches in IO theory. Rather, we view these approaches as a complement which can be highly valuable in applied research. For instance, in a merger, regulation, or litigation application, the economist is often interested in using game theory to analyze a very specific market. The estimates that we shall describe can allow the economist to build

”crash test dummies” of what might happen under alternative policies or allow the economist assist firms within those industries in decision making. In many applied policy settings, there is very little or no evidence from quasi experimental sources to inform decision making. Also, the policy changes may be too expensive or complex to engage in small scale experiments to gauge their effects.

In what follows, our goal is to introduce the key insights in the literature to the broadest audience possible by restricting attention to models and methods that are particularly easy to understand. We do not strive for the most elaborate or econometrically sophisticated estimation strategy. Instead, we hope to introduce strategies that are simple enough so that after reading this article, the nonspecialist with a working knowledge of econometrics can program the estimators using canned statistical packages without too much difficulty. Where appropriate, we shall direct the reader to more advanced papers which discuss more refined estimation procedures. We note that these more advanced papers typically rely on the insights which make our simple estimators possible. Therefore, we hope that our survey will be useful to the more sophisticated reader by focusing attention on what we believe are the key principals required for estimation.

Many of the papers that we will discuss are less than a decade old and this literature is clearly in its infancy. However, we hope to demonstrate both the generality and computational simplicity of many of these methods. The domain of applicability is not limited to industrial organization. These methods could be applied to estimating seemingly complicated structural models in other fields such as macro, labor, public finance, or marketing. The conceptual framework discussed in this survey offers a viable alternative to inversion nested swapping fixed point algorithms. By treating the ex ante value as a parameter which can be nonparametrically identified, rather than a nuisance function that is difficult to compute, the efficient and flexible estimators can be developed in the context of conditional moment models.

2 Motivating Example

In this section, we shall describe a simple econometric model of a game and discuss the heuristics of estimating this model. Our example is not intended to be particularly general or realistic. Our intention here is to exposit the key principles in formulating and estimating these models in the simplest possible manner. In the next sections, we shall demonstrate that these principles extend to much more general settings.

We first consider the static decision by a firm to enter a market, similar to the static entry models of Bresnahan and Reiss (1991), Berry (1992), Seim (2006), Jia (2008a), Ishii (2005), among other. For concreteness, suppose that we observe the decisions of two big box retailers, such as Walmart and Target. For each retailer $i = 1, 2$ we observe whether they enter some geographically separated markets $m = 1, \dots, M$. This set of potential markets is frequently defined as spatially separated population centers

from Census or related data. Let $a_{i,m} = 1$ denote a decision by retailer i to enter and $a_{i,m} = 0$ to not enter.

In the model, firms simultaneously choose their entry decisions. The model is static and firms receive a single period flow of profits. Economic theory suggests that a firm will enter a market m if the expected profit from entry is greater than the profit from not entering. Oligopoly models suggest that a firm's profits should depend on three factors. The first is consumer demand, which is commonly measured by market size and represented by POP_m , the population of market m . Other variables that would proxy for demand include demographic features such as income per capita. However, in order to keep our notation simple, we use only a single control for demand and ignore possibly richer specifications.

A second factor which enters profit is cost. Holmes (2008) argues that distribution is a key factor in the success of big box retailers. Walmart was founded in Bentonville, Arkansas and its subsequent entry decisions display a high degree of spatial autocorrelation. Specifically, Walmart followed a pattern of opening up new stores in close proximity to exiting stores and gradually fanned out from central United States. Holmes argues that proximity to existing distribution centers explains this entry pattern. Let $DIST_{im}$ denote the closest distribution center of firm i to market m . We shall use this as a measure of costs. A fully developed empirical model would use a richer specification.

A third factor which enters profits is competitive interactions. Suppose that Target is considering entering a market with 5,000 people. If Target believes that Walmart will also enter this market, there will be 2,500 customers per store. This is unlikely to be an adequate number of customers for both stores to be profitable. More generally, most oligopoly models predict that entry by a competitor will depress profits through increased competition. Therefore, a_{-i} , the entry decision by i 's competitor, should be an argument in the profit function. Oligopoly models suggest that other actions of competitors, such as pricing decisions and product choice, may also matter. However, we keep our specification parsimonious for illustration.

Following the above discussion, we specify the profits of firm i as:

$$\begin{cases} u_{im} = \alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot a_{-i,m} + \varepsilon_{im} & \text{if } a_{im} = 1 \\ u_{im} = 0 & \text{if } a_{im} = 0 \end{cases} \quad (1)$$

In the above, we normalize the profits of not entering to zero. The profit from entering depends on the exogenous covariates we discussed above and parameters α , β , and δ . These parameters index the contribution to profits of demand, cost, and competitive factors respectively. We let ε_{im} denote an iid shock to the profitability of firm i 's entry decision in market m . Inclusion of such shocks is standard in econometric models where agents make discrete choices. Failure to include such shocks lead to degenerate models that make deterministic predictions and hence will be trivially rejected by the data.

In our model, we shall assume that ε_{im} is private information to firm i . In practice, this means that firm $-i$ is unable to perfectly forecast firm i 's profits. In practice, this is a realistic assumption

in many markets. Some researchers have considered the case where ε_{im} is common knowledge, see for example Tamer (2003) Ciliberto and Tamer (2007) and Bajari, Hong, and Ryan (2004). The choice to model preference shocks as private information has two practical advantages. First, estimation is much easier. Many properties of the model can be studied in closed form, which then leads to clean identification. Second, very flexible versions of the model can be estimated almost trivially in standard software packages, such as STATA or Matlab.¹

Suppose that the econometrician has access to data on entry decisions and exogenous covariates from a cross section of markets $(a_{1,m}, a_{2,m}, POP_m, DIST_{im})$ for $m = 1, \dots, M$. The goal of estimation will be to learn the parameters α, β , and δ of the game. That is, we shall attempt to recover the game being played from the observed behavior of firms in the marketplace. Economic theory generally starts by specifying payoffs and then solving for equilibrium behavior. However, in econometrics, we study the inverse problem of recovering the game from observed actions, rather than deriving the actions from the specification of the game.

Let $\sigma(a_{i,m} = 1)$ denote the probability that a firm i enters market m . Firm i will make a best response to its equilibrium beliefs about $-i$'s equilibrium entry decision. Therefore, i 's decision rule is:

$$a_{i,m} = 1 \iff \alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1) + \varepsilon_{im} > 0 \quad (2)$$

That is, firm i will enter if its expected profits from doing so is greater than zero. Note that firm i does not know firm $-i$'s profits exactly because it does not observe ε_{-im} . Therefore, we will be studying the Bayes-Nash equilibrium to the game.

In practice, it is common for researchers to assume that ε_{im} has an extreme value distribution as in the conditional logit model. If we make this distributional assumption, then the standard result from discrete choice equation 2 simplifies that:

$$\sigma(a_{im} = 1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1))} \quad (3)$$

$$\sigma(a_{im} = 0) = \frac{1}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \sigma(a_{-im} = 1))} \quad (4)$$

Note that this closely resembles the standard binary logit model, where choice probabilities can be expressed using the exponential function in closed form. The formula depends on exogenous covariates and parameters in a closed form manner. The main difference is that the logit probabilities for i depend

¹However, the assumption that shocks are private information has substantive implications. Bajari, Hong, Krainer, and Nekipelov (2006) show that the number of equilibrium to the model will typically be much smaller. Indeed, in the examples they study, the average number of equilibria appears to decrease with the number of players in the game. A game with complete information would have an increasing average number of equilibrium in the number of players or other measures of the complexity of the game (see McLennan (2005)). Therefore, the predictions of the model are changed in ways that are not completely understood from a theoretical viewpoint. Obviously, this is an important arena for future research. Navarro and Takahashi (2010a) proposes a specification test for the private information model.

on the decisions of $-i$ through $\sigma_{-i}(a_{-i,m} = 1)$. Therefore, instead of being a single agent problem, the decisions of the agents are determined simultaneously.

Note that the equilibrium probabilities add up to one. Therefore, one of the equations in (3) and (4) is co-linear. As a result, we can express the Bayes-Nash equilibrium to this model as a system of two equations in two unknowns.

$$\sigma_1(a_{1m} = 1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{1m} + \delta \cdot \sigma_2(a_{2m} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{1m} + \delta \cdot \sigma_2(a_{2m} = 1))} \quad (5)$$

$$\sigma_2(a_{2m} = 1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{2m} + \delta \cdot \sigma_1(a_{1m} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{2m} + \delta \cdot \sigma_1(a_{1m} = 1))} \quad (6)$$

Note that in general, this defines a system of equations for each different market m each of which admits a distinct choice probability model. Pesendorfer and Schmidt-Dengler (2003) showed that the assumption of uniqueness of equilibrium in the data is more convincing when the markets are defined across different time period for a given geographical location than when markets are defined across different geographical location. Typically the unique equilibrium assumption is unlikely to hold in spatially heterogeneous markets, except in very special cases such as in the herding model of (Bajari, Hong, Krainer, and Nekipelov, 2006) where an equilibrium shifter is clearly identifiable.

Under the assumption that only one equilibrium is played out in the data, this system is extremely convenient to work with econometrically. First, note that the equilibrium can be written down in a closed form. This is much more convenient than complete information games where the equilibrium set cannot be characterized and is often quite complex when the number of players is large. A second advantage is that the equilibrium will be locally smooth at all but a measure zero set of covariates and parameters. This is because the equilibrium will inherit the smoothness of the logit model. This facilitates the econometric analysis of the model. Finally, our model is a generalization of the standard binary logit model, which is one of the best studied models in econometrics. As a result, many of the well worn tools from this literature will be applicable to our model. This is compelling for an applied researcher since the econometric analysis of discrete games to a large extent turns out to be a reasonably straightforward extension of discrete choice.

2.1 Two-Step Estimators

Much of the literature on the empirical analysis of games relies on multistep estimators. In what follows, we shall describe an approach that has the greatest computational simplicity, rather than focusing on estimators that are more efficient or have other desirable econometric properties at the cost of being more difficult to estimate. The estimator that we shall describe works in two steps. In the first step, the economist estimates the reduced form of the model. In the second step, the economist estimates the structural parameters taking the reduced form as given. We shall sketch this estimator heuristically for

intuition. We will formalize the econometric details more precisely in the next section.

2.1.1 Reduced Form

The reduced form is the distribution of the dependent variable given the exogenous variables in our model. Formally, the reduced form can be viewed as the solution to equations (5) and (6). We can view this system as two equations in the two unknown entry probabilities. In general, the solution to this equation cannot be expressed in closed form and will depend on the exogenous variables $POP_m, DIST_{1m}, DIST_{2m}$. We shall let $\sigma_1(a_{1m}=1|POP_m, DIST_{1m}, DIST_{2m})$ and $\sigma_2(a_{2m}=1|POP_m, DIST_{1m}, DIST_{2m})$ denote the solution.

The reduced form is a "flexible" estimate of $\sigma_1(a_{1m}=1|POP_m, DIST_{1m}, DIST_{2m})$ and $\sigma_2(a_{2m}=1|POP_m, DIST_{1m}, DIST_{2m})$. We form this using the underlying data $(a_{1,m}, a_{2,m}, POP_m, DIST_{1m})$ for $m = 1, \dots, M$ and a suitably flexible estimation method, for example, a sieve logit (see Newey and Powell (2003) and Ai and Chen (2003)). Here we use M to denote the number of markets which can be pooled because the same equilibrium is played in these markets. Typically M denotes the number of time periods in a single geographical location as in Pesendorfer and Schmidt-Dengler (2003). Occasionally in rare cases it might also refer to the number of spatially heterogeneous markets. Let $z_k(POP_m, DIST_{1m}, DIST_{2m})$ denote the vector of terms in a k^{th} order polynomial in $POP_m, DIST_{1m}, DIST_{2m}$. We could model the reduced form choice probabilities as:

$$\sigma_i(a_i = 1|s, \beta) = \frac{\exp(z_k(s)'\theta)}{1 + \exp(z_k(s)'\theta)}$$

Where we let $k \rightarrow \infty$ as the number of markets $M \rightarrow \infty$, but not too fast, i.e. $\frac{k}{M} \rightarrow 0$. For any finite sample size, we can estimate the sieve logit using a standard software package, such as STATA. This is simply a method to model choice probabilities in a flexible way that exhausts the information in the data. Other flexible methods are also possible. In the next section, we shall demonstrate that in many cases, the choice of the method to estimate the first stage typically does not matter for the asymptotic distribution of the structural parameters. In our applied work we have found that with large sample sizes, results are reasonably robust to the specification of the first stage so long as it is sensibly and flexibly specified. Let $\hat{\sigma}_1(a_{1m}=1|POP_m, DIST_{1m}, DIST_{2m})$ and $\hat{\sigma}_2(a_{2m}=1|POP_m, DIST_{1m}, DIST_{2m})$ denote these first stage estimates.

If our problem is well behaved and the data across multiple markets are generated from a unique equilibrium, it will be the case that $\hat{\sigma}_i(a_{im}=1|POP_m, DIST_{1m}, DIST_{2m})$ will converge to

$$\sigma_i(a_{im}=1|POP_m, DIST_{1m}, DIST_{2m})$$

as the sample size becomes large. When multiple equilibria are present in different markets, such estimates can potentially diverge. Suppose that we replace σ_i with our consistent estimate $\hat{\sigma}_i$ in equation (2). Then

the agent's decision rule can be rewritten as:

$$a_i = 1 \iff \alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-im} = 1) + \varepsilon_{im} > 0$$

This implies that the probability that i chooses to enter is:

$$\sigma_i(a_{im}=1) = \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-im} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-im} = 1))} \quad (7)$$

If our problem is well behaved, the convergence of $\hat{\sigma}_{-i}$ to σ_i implies that (7) will converge to (3).

Taking our first step reduced form estimates of $\hat{\sigma}_1$ and $\hat{\sigma}_2$ as given, we can form the following pseudo likelihood function to estimate our structural parameters α, β, δ in a second step.

$$L(\alpha, \beta, \delta) = \prod_{m=1}^M \prod_{i=1}^2 \left(\frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-i} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-i} = 1))} \right)^{1_{\{a_{i,m}=1\}}} \left(1 - \frac{\exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-i} = 1))}{1 + \exp(\alpha \cdot POP_m + \beta \cdot DIST_{im} + \delta \cdot \hat{\sigma}_{-i}(a_{-i} = 1))} \right)^{1_{\{a_{i,m}=0\}}}$$

As we shall show in the next section, solving this pseudo MLE problem generates consistent estimates of the structural parameters α, β, δ . when the equilibria in different markets are the same. When multiple equilibria exist in the data, these parameter estimates are clearly inconsistent.

2.2 Discussion

A few points are worth discussing. First, note that the method that we propose is computationally very simple. In the first step, we simply estimate $\hat{\sigma}_i(a_{im}=1|POP_m, DIST_{1m}, DIST_{2m})$ flexibly using standard methods, such as a conditional logit model using a flexible basis function. In the second step, we run the conditional logit once again taking our first stage estimates as given. Using standard econometric packages, such as STATA, we can estimate this model very simply, often with less than 20 lines of code. A key advantage of this approach is that it will be very accurate computationally and robust to programming errors. In the literature, structural models are often estimated using the nested fixed point algorithm. While this is a very important and useful approach, it often requires considerable CPU time and numerical sophistication on the part of the econometrician.² In models with many firms, Weintraub, Benkard, and Van Roy (2008) propose the notion of oblivious equilibrium to ease the computational burden. The approach that we propose has the disadvantage of being inefficient as compared to the nested fixed point. However, it has the practical advantages of being much simpler and more accurate numerically. Aguirregabiria and Mira (2007) demonstrate how iterating the procedure proposed above increases efficiency while still retaining computational simplicity.

Second, because our approach builds on the well known conditional logit, this facilitates testing and specification. For example, in practice, the econometrician may not be certain how to specify payoffs

²For example, Dube, Fox, and Su (2009) demonstrate that the nested fixed point algorithm may fail to converge if the inner loop tolerance is not chosen very carefully.

in the game. The appropriate covariates and even strategies may not be evident based on a priori considerations. These choices can be made using the extensive literature on hypothesis and specification testing available for the conditional logit. For instance, the econometrician may experiment with multiple specifications of the model using different covariates and even strategies. Specifications that generate increased fit, covariates that have higher significance levels and models that survive standard specification tests (e.g. Pearson’s goodness of fit) may be preferred on econometric grounds. We note that all of these procedures are available in canned statistics packages. The simplicity of our framework allows the economist to test the robustness of substantive economic conclusions to alternative specifications of the game, similar to how one would report the sensitivity of results in standard discrete choice models. If panel data is available, the use of firm and market level fixed effects can be used to control for unobserved heterogeneity using standard methods for analyzing the conditional logit. In general, as pointed out in Pesendorfer and Schmidt-Dengler (2003), the logit assumption is not necessary. The researcher is free to substitute the logit distribution with other forms of distributions.

Third, the framework we propose strictly generalizes standard discrete choice models, which are single agent models where decisions are made in isolation. This corresponds to a specification in which $\delta = 0$. We note that this assumption is testable. Furthermore, as we shall discuss below, the presence of strategic interactions can be identified from the data under weak functional form assumptions so long as appropriate exclusion restrictions exist.

Fourth, key insight above generalizes to a much richer set of models which strictly nest the specification considered above. In the above, we simply replaced $\sigma_i(a_{im}=1)$ with its empirical analogue $\hat{\sigma}_i(a_{im}=1)$. Instead of solving for the equilibrium directly in estimation, we learn the equilibrium strategies from the data. This approach has practical advantages in face of game theoretic models that generate multiple equilibrium, which complicates estimation. This means that a random mechanism must exist which is used by the players to coordinate on which equilibrium to play, an assumption that is arguably strong and possibly ad hoc. The data generating process must specify how the equilibrium is chosen in order to form a likelihood function for the model. Also, the existence of multiple equilibrium complicated a straightforward application of maximum likelihood or the generalized method of moments. The complications frustrated efforts to derive point estimators for structural parameters in the earlier literature. The prior literature either used bounds estimation strategies (as in Tamer (2003), Ciliberto and Tamer (2007) or Pakes, Porter, Ho, and Ishii (2005) or restricted attention to specific games, using the properties of equilibrium to identify parameters (Bresnahan and Reiss (1991) and Berry (1992)).

In a sense, the two step approach ”solves” the multiplicity problem by letting the data tell us which equilibrium is played if the model generates multiple outcomes. Note that we place ”solves” in quotation marks because of potential disagreement with this approach that we are taking. One might consider this approach as assuming away the problem instead of ”solving” the problem. For example, if our entry model generates two solutions but agents coordinate on a single solution, a consistent estimator

of $\sigma_i(a_{im}=1)$ will allow us to learn the data from the equilibrium. However, this approach comes at a cost. First, if the agents flip back and forth across equilibrium at random in different markets m , it may not be possible to consistently estimate $\sigma_i(a_{im}=1)$, rendering our estimates of the structural parameters α, β, δ inconsistent. This can be true particularly when equilibria are endogenously chosen in each market m . Second, the ability to estimate our model in two steps depends crucially on our private information assumption. If we assumed instead that the shocks $\varepsilon_{i,m}$ are complete information, the equilibrium probabilities would depend on both ε_{1m} and ε_{2m} . This would destroy our estimation strategy since we do not directly observe these variables and cannot condition on them in a first step.

3 A General Static Model

In this section, we generalize the model of the previous section by allowing for a more general set of players and payoffs. We shall restrict attention to games with two strategies for notational simplicity. This assumption can be easily generalized and the next section allows for this complication as a special case.

In the model, there are $i = 1, \dots, n$ players in the game. Each player can simultaneously choose an action $a_i \in \{0, 1\}$. Let $A = \{0, 1\}^n$ denote the Cartesian product of the agents' actions and $a = (a_1, \dots, a_n)$ a generic element of A . Let $s \in S$ denote a vector of variables which can influence an agent's payoffs. In what follows, we shall assume that s is observed by both the agent and the econometrician. The issue of omitted variables will be treated later in the article. Let the payoffs to agent i be written as:

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

In the above, $\Pi_i(a_i, a_{-i}, s)$ is a general function that models how an agent's own actions, a_i , the actions of other agents, a_{-i} and the variables s influence payoffs. Our model will be general in that we will not impose any particular parametric restriction on this function. The term ϵ_i is an iid disturbance to agent i 's payoffs, which we assume to have an extreme value distribution as before. Alternative functional forms are possible, but this specification is the most popular due to its analytical convenience.

Similar to the analysis of discrete choice models, it is necessary to impose a normalization on the function $\Pi_i(a_i, a_{-i}, s)$ similar to normalizing the utility from the outside good to zero in a single agent discrete choice model.

$$\text{For all } a_{-i} \in A_{-i} \text{ and all } s, \Pi_i(a_i = 0, a_{-i}, s) = 0 \tag{8}$$

Note that this normalization must hold for any strategy a_{-i} that could be used by the other agents and any state. We implicitly imposed this normalization in (1) when we normalized the profits from not entering a market to zero. Note however, that any other normalization scheme can be used in place of the zero profit or utility normalization. Essentially one normalized value is needed for each value of the

state variable s . In a static game, normalizing one of the baseline utilities to zero is innocuous, although this is not longer the case in dynamic models.

3.1 Equilibrium

Having defined payoffs, we next turn to the problem of defining a Bayes-Nash equilibrium to our game. In our model, ϵ_i is agent i 's private information. Let $\sigma_i(a_i = 1, s)$ denote the probability that i chooses the action a_i conditional on the vector of payoff relevant variables s . In a Bayes-Nash equilibrium, player i has consistent beliefs about the actions of other agents in the model. That is, she knows the equilibrium probabilities $\sigma_{-i}(a_{-i}|s)$. We define the *choice specific value function* (or the choice specific expected profit function) as the expected value of $\Pi_i(a_i = 1, a_{-i}, s)$, the return to agent i from taking the action $a_i = 1$, given i 's beliefs about the equilibrium strategies of other agents in the model:

$$\Pi_i(a_i = 1, s) = \sum_{a_{-i} \in A_{-i}} \sigma_{-i}(a_{-i}|s) \Pi_i(a_i = 1, a_{-i}, s) \quad (9)$$

In the above, we sum over all actions $a_{-i} \in A_{-i}$ that could be taken by other agents in the model. The term $\sigma_{-i}(a_{-i}|s)$ is the equilibrium probability that the actions a_{-i} are observed. The above summation marginalizes out agent i 's uncertainty about the actions of other agents to compute the expected return from choosing $a_i = 1$ given s .

In a Bayes-Nash equilibrium, agents maximize expected utility. From the definition of the choice specific value function, it follows that:

$$\begin{aligned} a_i = 1 &\iff \sum_{a_{-i} \in A_{-i}} \sigma_{-i}(a_{-i}|s) \Pi_i(a_i = 1, a_{-i}, s) + \epsilon_i > 0 \\ &\iff \Pi_i(a_i = 1, s) + \epsilon_i > 0 \end{aligned} \quad (10)$$

This is the generalization of (2) to our more general model. This equation states that the agent i chooses $a_i = 1$ when the choice specific value function plus the iid disturbance is greater than zero, which is the payoff from choosing $a_i = 0$.

Since ϵ_i is distributed extreme value, it follows that a Bayes Nash equilibrium is a collection of functions $\sigma_i(a_i = 1|s)$ for $i = 1, \dots, n$ such that

$$\text{for all } i \text{ and all } s, \sigma_i(a_i = 1|s) = \frac{\exp(\Pi_i(a_i = 1, s))}{1 + \exp(\Pi_i(a_i = 1, s))}. \quad (11)$$

Note that similar to (5) and (6), i 's decision rule has the convenient and familiar form of the conditional logit in this more general model. For a fixed s , the definition of equilibrium (11) and the definition of the choice specific value function(9) implies that the n equilibrium probabilities, $\sigma_i(a_i = 1|s)$ can be viewed

as the solution to following system of n equations:

$$\begin{aligned}
\sigma_1(a_1 = 1|s) &= \frac{\exp(\sum_{a_{-1} \in A_{-1}} \sigma_{-1}(a_{-1}|s) \Pi_1(a_1 = 1, a_{-1}, s))}{1 + \exp(\sum_{a_{-1} \in A_{-1}} \sigma_{-1}(a_{-1}|s) \Pi_1(a_1 = 1, a_{-1}, s))} \\
\sigma_2(a_2 = 1|s) &= \frac{\exp(\sum_{a_{-2} \in A_{-2}} \sigma_{-2}(a_{-2}|s) \Pi_2(a_2 = 1, a_{-2}, s))}{1 + \exp(\sum_{a_{-2} \in A_{-2}} \sigma_{-2}(a_{-2}|s) \Pi_2(a_2 = 1, a_{-2}, s))} \\
&\vdots \\
\sigma_n(a_n = 1|s) &= \frac{\exp(\sum_{a_{-n} \in A_{-n}} \sigma_{-n}(a_{-n}|s) \Pi_n(a_n = 1, a_{-n}, s))}{1 + \exp(\sum_{a_{-n} \in A_{-n}} \sigma_{-n}(a_{-n}|s) \Pi_n(a_n = 1, a_{-n}, s))}
\end{aligned} \tag{12}$$

As in the simple example of the previous section, many of the simple attractive features of the model are preserved in this more general set up. The equilibrium choice probabilities are defined by a smooth system of equations that can be simply characterized in closed form. This will facilitate both the identification and estimation of our model. We note that the above model appears to have been re-invented several times in the literature. McKelvey and Palfrey (1995) interpret the error terms as optimization error, rather than private information, and refer to their model as Quantile Response Equilibrium. This has been widely applied in the experimental economics literature. Manski (1993) and Brock and Durlauf (2001) consider a closely related model with a very large number of agents. This model has played an important role in the literature in peer effects. In industrial organization, this model was first proposed by Seim (2006) to the best of our knowledge. We take this re-invention of the model as a signal that it is an econometrically natural and useful way to characterize behavior in a surprisingly wide variety of settings including peer effects, experiments, games, and empirical industrial organization.

4 Identification

A central question in the analysis of any econometric model is what conditions are required for identification. The model is *identified* if we can reverse engineer the structural payoff parameters $\Pi_i(a_i, a_{-i}, s)$ if we observe the agent's strategies $\sigma_i(a|s)$ for $i = 1, \dots, n$. This question has been studied for related models by Bresnahan and Reiss (1991), Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2003), and Manski (1993), among others. In what follows, we limit attention to a set of results that are transparent and may be useful in empirical work. We maintain the assumption that $\epsilon_i(a_i)$ are distributed i.i.d. extreme value across actions a_i and agents i . Bajari, Hong, and Ryan (2004) demonstrate that if payoffs are allowed to be nonparametrically specified in a single agent discrete choice model, then the error term distribution needs to be parametrically specified. Since discrete choice models are less general than discrete games, assumptions at least as strong will be required to identify discrete game models. An alternative approach is to allow the error terms to have an unknown distribution and specify payoffs using a linear index. This approach is discussed in Tamer (2003) and Bajari, Chernozhukov, Hong, and

Nekipelov (2009). We shall also maintain the assumption that (8) holds.

Our approach to identification is constructive, and is only applicable when there is no multiple equilibria in the data. In the absence of the complications induced by the presence of multiple equilibria, we begin by taking the equilibrium choice probabilities $\sigma_i(a_i = 1|s)$ as given. We then ask whether it is possible to reverse engineer the structural parameters $\Pi_i(a_i, a_{-i}, s)$. We begin our analysis by working with equation (11). We perform the "Hotz and Miller (1993)" inversion and take logs of both sides of this equation. This yields the following result:

$$\Pi_i(a_i = 1, s) = \log(\sigma_i(a_i = 1|s)) - \log(\sigma_i(a_i = 0|s)) \quad (13)$$

This equation shows that we can learn the choice specific value function, for any s , from knowledge of the choice probabilities. Therefore, in what follows, we treat $\Pi_i(a_i = 1, s)$ as known.

Let us fix s . By the definition of the choice specific value function, recall that:

$$\Pi_i(a_i = 1, s) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s) \Pi_i(a_i = 1, a_{-i}, s), \text{ for } i = 1, \dots, n \quad (14)$$

In the above system, $\Pi_i(a_i = 1, s)$ and $\sigma_{-i}(a_{-i}|s)$ can be treated as known for all i . Identification requires us to find a unique set of $\Pi_i(a_i = 1, a_{-i}, s)$ that solves this system of equations. A necessary condition for identification is that there are at least as many equations as free parameters $\Pi_i(a_i = 1, s)$. For fixed s , there are $n \times 2^{n-1}$ unknowns corresponding to the $\Pi_i(a_i = 1, a_{-i}, s)$ after we impose the normalization (8). However, there are only n equations. This implies that with restrictions, the structural parameters are not identified.

Result If the number of players $n > 1$, then the structural parameters $\Pi_i(a_i = 1, s)$ are not identified.

Upon reflection, the lack of identification of the model should not be surprising. Our problem is similar to a standard linear simultaneous equations problem where the $\Pi_i(a_i = 1, a_{-i}, s)$ play the role of the unknown coefficients. Without restrictions, linear simultaneous system is not identified.

A standard approach to identification in simultaneous equation models is to impose exclusion restrictions. We follow that approach here. Let us 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 s_i . That is, it is possible to find a set of covariates that shift the payoffs of agent i independently of the payoffs of the other agents in the model. For example, in our simple example, $DIST_i$ played this role. We did not let the cost shifter for firm $-i$, $DIST_{-i}$, enter into firm i 's payoffs.

Many commonly used oligopoly models generate similar exclusion restrictions. In oligopoly models, the payoff of firm i depend on the actions that its competitors $-i$ take in the marketplace, such as their pricing, product offering or entry decisions. The profits of i typically exclude costs and other shocks experienced by competing firms which shift $-i$'s incentives to take actions in the marketplace. These restrictions are commonly imposed in empirical work as well. There are models in which such exclusions

would be violated, for example, if there are externalities in costs or other shocks between firms. However, the above result demonstrates that if we let the payoffs of all firms depend on all covariates, our model will be underidentified without further restrictions.

If the economist is willing to impose exclusion restrictions, based on theoretical considerations or institutional knowledge, then the choice specific value function can be written as:

$$\Pi_i(a_i, s_{-i}, s_i) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}|s_{-i}, s_i) \Pi_i(a_i, a_{-i}, s_i) \text{ for } i = 1, \dots, n \quad (15)$$

In the above equation, note that s_{-i} enters σ_{-i} but is excluded from Π_i . Holding s_i fixed, if we vary s_{-i} , we can increase the number of equations that $\Pi_i(a_i, a_{-i}, s_i)$ must satisfy. If there are at least 2^{n-1} points in the support of the conditional distribution of s_{-i} given s_i , we will generate more equations than free parameters.³ In models with a finite number of discrete state variables, this is a result developed in Pesendorfer and Schmidt-Dengler (2003).

Result. If there are at least 2^{n-1} points in the support of the conditional distribution of s_{-i} given s_i and the matrix $[\sigma_{-i}(a_{-i}|s_{-i}, s_i) \sigma_{-i}(a_{-i}|s_{-i}, s_i)']$ is not co-linear, then $\Pi_i(a_i, a_{-i}, s_i)$ is identified.

The intuition behind how exclusion restrictions generate identification is fairly straightforward. Let us return to the simple example of the previous section. We exclude $DIST_{-i}$ from the payoffs of firm i . If we wish to identify δ , we would hold $DIST_i$ fixed and perturb $DIST_{-i}$. If firm i is a Walmart store, this will involve tracing out a constant radius from the Walmart fulfillment center. Along this radius, the distance to Target fulfillment centers will vary. This will perturb Target's entry probability, but leave Walmart's payoffs unchanged. This variation shifts Target's equilibrium strategy, leaving the other components of Walmart's payoffs fixed.

Another point to notice about our model is that it is in fact over-identified. In many applications, s_{-i} may include continuous variables with a rich support. This means that there are a large number of overidentifying restrictions because there are many more equations than free parameters. This implies that our model in principal could be rejected by the data.

5 Nonparametric Estimation

Next, we discuss a simple method to estimate the structural parameters of our model non-parametrically. The estimation method that we propose is the empirical analogue of our identification argument and will take place in three fairly simple steps. We do not attempt to construct the most efficient estimator. Instead, our goal is to present the estimator that is the simplest to program and understand.

³Of course, our model must not be degenerate in the sense that the equations are co-linear with each other. However, that condition is testable if we know the population probabilities.

5.1 First Step: Estimation of Choice Probabilities.

Suppose that the econometrician has data on $t = 1, \dots, T$ repetitions of the game with actions and states $(a_{i,t}, s_{i,t})$, $i = 1, \dots, n$. In the first step, we form an estimate $\hat{\sigma}_i(a_i|s)$ of $\sigma_i(a_i|s)$ using a flexible method. A particularly convenient method is a sieve logit as in the previous section. We let $z_k(s)$ denote the vector of terms in a k^{th} order basis function. For example, $z_k(s)$ could correspond to splines, B-splines, orthogonal polynomial or other set of rich basis functions. Define

$$\sigma_i(a_i = 1|s, \beta) = \frac{\exp(z_k(s)' \theta)}{1 + \exp(z_k(s)' \theta)}$$

Let $k \rightarrow \infty$ as sample size $T \rightarrow \infty$, but not too fast, i.e. $\frac{k}{T} \rightarrow 0$. The consistency of this estimator depends obviously on the uniqueness of equilibrium in the data. Unless there is a unique equilibrium in the game, this estimator will generally be inconsistent. When the consistency condition holds, kernel methods or other smoothing methods, such as local regressions, can also be used to replace the sieve nonparametric estimation method. It is possible to show that in a semiparametric model, the asymptotic distribution of the second stage estimator does not depend on the form of the smoother chosen in the first stage.

In Bajari, Hong, Krainer, and Nekipelov (2006), we find that a sieve logit used to estimate the first stage and a straightforward parametric model generate almost identical estimates. The researcher should use standard tools from the analysis of discrete choice in specifying the first stage (e.g. tests of goodness of fit, insignificant regressors, etc.). Specifying the first stage is a simple curve fitting exercise.

5.2 Second Step: Inversion

In the second stage, we perform the empirical analogue of the Hotz-Miller inversion.

$$\hat{\Pi}_i(a_i = 1, s_t) = \log(\hat{\sigma}_i(a_i = 1|s_t)) - \log(\hat{\sigma}_i(a_i = 0|s_t)) \quad (16)$$

That is, we estimate the choice specific value function by evaluating the empirical analogue of (13) using our first stage estimates.

5.3 Third Step: Recovering The Structural Parameters

In our third step, we evaluate the empirical analogue of equation (14). If we substitute in the empirical analogue of $\hat{\Pi}_i(a_i, s_t)$ for $\Pi_i(a_i, s_t)$ recovered from the second step and $\hat{\sigma}_{-i}(a_{-i}|s_{-it}, s_{it})$ for $\sigma_{-i}(a_{-i}|s_{-it}, s_{it})$, this equation can be written as:

$$\hat{\Pi}_i(a_i, s_t) \approx \sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s_{-it}, s_{it}) \Pi_i(a_i, a_{-i}, s_{i,t}) \quad (17)$$

If the empirical analogues converge to their population values as T becomes large, this approximate equality will become exact.

Suppose that we wish to estimate agent i 's payoff at a particular value of the explanatory variables \tilde{s}_i . Let us denote the payoffs at this point by $\Pi_i(a_i, a_{-i}, \tilde{s}_i)$. At an abstract level, we would do this by inverting (14) at \tilde{s}_i by perturbing s_{-i} holding \tilde{s}_i fixed. However, in a finite sample, this will typically not be possible because there will be no value t where \tilde{s}_i is exactly equal to s_{it} and we can only evaluate the empirical analogue of (14).

However, we can estimate $\Pi_i(a_i, a_{-i}, \tilde{s}_i)$ through a technique called local linear regression (see Fan and Gijbels (1996)). The idea is that, while we cannot see values of s_{it} exactly equal to \tilde{s}_i , we will see values that are close in a large sample. We will overweight the values of s_{it} that are close to \tilde{s}_i using kernel weights. That is, for every observation t , we construct a weight $K\left(\frac{s_{it}-\tilde{s}_i}{h}\right)$ where K is a kernel and h is the bandwidth. This weighting scheme overweights observations t that are close to \tilde{s}_i .

We then estimate $\Pi_i(a_i, a_{-i}, \tilde{s}_i)$ by solving the following weighted least squares problem:

$$\arg \min_{\Pi_i(a_i, a_{-i}, \tilde{s}_i)} \frac{1}{T} \sum_{t=1}^T \left(\frac{\hat{\Pi}_i(a_i, s_t) - \sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i} | s_{-it}, s_{it}) \Pi_i(a_i, a_{-i}, \tilde{s}_i)}{\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i} | s_{-it}, s_{it}) \Pi_i(a_i, a_{-i}, \tilde{s}_i)} \right)^2 K\left(\frac{s_{it}-\tilde{s}_i}{h}\right)$$

Asymptotically, if we let the bandwidth h become small at an appropriate rate as the sample size T increases, this will generate a consistent estimate of $\Pi_i(a_i, a_{-i}, \tilde{s}_i)$ if the game has a unique equilibrium as in Seim (2006). The form of this last stage regression is suitable for the nonparametric specification of the primitive payoff function $\Pi_i(a_i, a_{-i}, \tilde{s}_i)$. In the semiparametric specification that is discussed below, there are two alternatives. A similar regression can be run in the space of primitive payoff functions. Alternatively, least squares can also be implemented in the choice probability spaces. Then the nonparametrically identified $\hat{\Pi}_i(a_i, s_t)$ can be projected on the fitted value of $\Pi_i(a_i, a_{-i}, \tilde{s}_i)$ which combines the parametric functional form and the least square estimates of the choice probabilities.

5.4 A Semiparametric Estimator

While nonparametric estimation of a game is certainly elegant, it may not be practical for empirical work, especially when the number of control variables in s is large, due to curse of dimensionality. However, this estimator is useful for pedagogical reasons: it shows how our identification arguments can be directly translated into an estimation strategy.

We next describe a semiparametric estimator that is much more practical for empirical work. Suppose that the econometrician is willing to specify payoffs using a parametric model. For example, we could approximate payoffs using a set of basis function, (e.g. linear index, polynomial, etc...)

$$\Pi_i(a_i, a_{-i}, s_i) = \Phi_i(a_i, a_{-i}, s)' \theta$$

For example, in our simple example, payoffs were specified using a linear index, which is very common in empirical work. If we made this assumption, it would be possible to estimate θ by solving the following minimization problem:

$$\arg \min_{\theta} \frac{1}{T} \sum_{t=1}^T \left(\frac{\hat{\Pi}_i(a_i, s_t) - \sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s_{-it}, s_{it}) \Phi_i(a_i, a_{-i}, s)' \theta}{\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s_{-it}, s_{it}) \Phi_i(a_i, a_{-i}, s)' \theta} \right)^2$$

Note that this is simply a least squares procedure where the regressors are defined using the nonparametric first step, $\hat{\sigma}_{-i}(a_{-i}|s_{-it}, s_{it})$, and the basis functions, $\Phi_i(a_i, a_{-i}, s)$. The dependent variable is simply the empirical analogue of the choice specific value function found in our second step. The point estimate $\hat{\theta}$ is then simply found by solving the above least squares problem.

An alternative estimation strategy is to form a pseudo-likelihood as in our simple example. Suppose that we simply substitute $\hat{\sigma}_{-i}(a_{-i}|s_{-it}, s_{it})$ in place of $\sigma_{-i}(a_{-i}|s_{-it}, s_{it})$ in equation (12). Then the probability that $a_i = 1$ conditional on s is equal to:

$$\begin{aligned} & \frac{\exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Pi_i(a_i = 1, a_{-i}, s))}{1 + \exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Pi_i(a_i = 1, a_{-i}, s))} \\ &= \frac{\exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Phi_i(a_i, a_{-i}, s)' \theta)}{1 + \exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Phi_i(a_i, a_{-i}, s)' \theta)} \end{aligned}$$

The pseudo log-likelihood function is:

$$L(\theta) = \frac{1}{T} \sum_{i,t} \left[\log \left(\frac{\exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Phi_i(a_i, a_{-i}, s)' \theta)}{1 + \exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Phi_i(a_i, a_{-i}, s)' \theta)} \right) \mathbf{1}\{a_{i,t} = 1\} + \log \left(\frac{1}{1 + \exp(\sum_{a_{-i}} \hat{\sigma}_{-i}(a_{-i}|s) \Phi_i(a_i, a_{-i}, s)' \theta)} \right) \mathbf{1}\{a_{i,t} = 0\} \right]$$

We note that this is simply a multinomial logit model where the regressors are defined using $\hat{\sigma}_{-i}(a_{-i}|s) \Phi_i(a_i, a_{-i}, s)$. As in our simple example, this can be estimated using conventional statistical packages.

Bajari, Hong, Krainer, and Nekipelov (2007) develop the asymptotic theory for this regressor. The key result is that both the least squares and pseudo-MLE consistently estimate θ , unless there are multiple equilibria issues in the data. Moreover, $\hat{\theta}$ converges at a $T^{1/2}$ rate and has normal asymptotics. This implies that we can simply bootstrap our regressions or pseudo-MLE for computing confidence intervals or hypothesis testing.

The estimator is computationally simple. Bajari, Hong, Krainer, and Nekipelov (2007) were able to estimate a discrete game model with hundred of players and fixed effects (including the standard errors) using a simple STATA programming written with less than a page of code.

6 Dynamic Discrete Games

In this section, we demonstrate that the principles exposited earlier can be extended to a more general dynamic model. The key insight of estimation is the same as that of the static case. We again break our estimation into two steps. In the first step, we estimate the model's reduced forms using flexible methods. In the second step, we recover the structural parameters rationalize the reduced forms as the outcome of optimal decisions. In what follows, we exposit the model that is the easiest to estimate, since

we intend to present our method to the most general possible audience. We shall discuss generalizations in the next section.

6.1 The Model

We consider a stationary dynamic environment. Let $t = 1, \dots, \infty$ denote time; let $s \in S$ denote a payoff relevant state variable. We assume that s is a real valued vector and S is a possibly uncountable subset of \mathfrak{R} . Let a_{it} and s_{it} denote the choice and state respectively of agent i at time t . We shall assume that $a_{it} \in \{0, 1, \dots, K\}$ for all i and t . The model could be easily generalized to allow for different choice sets across agents or choice sets that depend on the state variable s_{it} .

In the model, i 's period return function is:

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

In the above, $\Pi_i(a_{it}, a_{-it}, s_t; \theta)$ is a deterministic function which influences agent i 's payoffs through her actions, the actions of her competitors, and the state. The term $\epsilon_{it}(a_{it})$ represents an iid choice specific random shock to i 's payoffs at time t . As in the previous section, we assume that s_{it} is observed by both the agents and the econometrician, while $\epsilon_{it}(a_{it})$ is private information to agent i . We shall let $g(s'|s, a_i, a_{-i})$ be the law of motion for state s_{it} . Note that because ϵ_{it} private information and not public information, we can restrict the law of motion to be Markov and hence do not include ϵ_{it} in g . While such restrictions can be weakend, they are extremely common in applied work. See Rust (1994) for a discussion.

We shall restrict attention to Stationary Markov Perfect equilibrium. This means that the probability that agent i taking action a_i can be written as a function $\sigma_i(a_i|s)$. Note that this rules out "punishment" schemes that would allow for the wide multiplicity of equilibrium as in the folk theorem. While this is a restriction of the model, it is not as harsh as it might first sound. First, the model has overidentifying restrictions that will allow us to test our framework. These restrictions would allow us to detect at least some punishment schemes. For instance, note that after we condition on s_t and its lagged values, the choices of the agents should be independent. This means that there should be no correlation between a_{it} and lagged choices a_{-it} of the other agents. We can test for this correlation using fairly straightforward hypothesis testing. (Navarro and Takahashi (2010b)) presents a more elaborate framework). If agents were using time dependent punishments, such a test might likely fail. This can be an important point that is worthy of further development. Second, restricting attention to Markov Perfect equilibrium forces us to be parsimonious in our modeling choices. The folk theorem leads to few, if any, useful implications for dynamic oligopoly models because it predicts that a very large space of outcomes is possible. From an applied econometrics point of view, taking a model that generates a folk theorem rest to data is akin to estimating an underidentified model. A more practical and scientific approach would specify a

priori plausible restrictions and then test their validity by specification testing and by making refutable predictions.

As in the previous section, define $\Pi_i(a_{it}, s)$ as

$$\Pi_i(a_i, s) = \sum_{a_{-i} \in A_{-i}} \Pi_i(a_i, a_{-i}, s) \sigma_{-i}(a_{-i}|s). \quad (18)$$

Once again, this is the period return that agent i receives from choosing the strategy a_i , marginalizing out the actions of the other agents in the game.

In the model, agents maximize the present value of expected discounted utility using a fixed discount factor β . While we confide the following discussion to models in which β is a fixed parameter, the discount rate can also be heterogeneous across agents, especially in models with unobserved heterogeneity. The value function for our model can be written as:

$$W_i(s, \epsilon_i; \sigma) = \max_{a_i \in A_i} \left[\Pi_i(a_i, s) + \epsilon_i(a_i) + \beta \int \sum_{a_{-i}} W_i(s', \epsilon'_i; \sigma) g(s'|s, a_i, a_{-i}) \sigma_{-i}(a_{-i}|s) f(\epsilon'_i) d\epsilon'_i ds' \right] \quad (19)$$

In the above, the period return is defined by (18) plus an iid preference shock. If an agent chooses action a_i , the probability that the state in the next period is s' is equal to $\sum_{a_{-i}} g(s'|s, a_i, a_{-i}) \sigma_{-i}(a_{-i}|s)$. Note that this formula integrates out the actions of all the other agents by summing over a_{-i} .

Definition A Markov perfect equilibrium is a collection of policy functions $a_i^* = \delta_i^*(s, \epsilon_i)$ and corresponding conditional choice probabilities $\sigma_i^*(a_i|s)$ such that, for all i , all s , and all ϵ_i , $a_i^* = \delta_i^*(s, \epsilon_i)$ maximizes

$$\begin{aligned} \delta_i^*(s, \epsilon_i) &= \arg \max_{a_i \in A_i} \left[\Pi_i(a_i, s) + \epsilon_i(a_i) + \beta \int \sum_{a_{-i}} W_i(s', \epsilon'_i; \sigma) g(s'|s, a_i, a_{-i}) \sigma_{-i}^*(a_{-i}|s) f(\epsilon'_i) d\epsilon'_i ds' \right] \\ \sigma_i^*(a_i|s) &= \int \delta_i^*(s, \epsilon_i) f(\epsilon_i) d\epsilon_i \end{aligned}$$

In practice, it is common to parametrically specify the period utility function. As before, we shall let $\Pi_i(a_i, s; \theta)$ denote a parametric specification of payoffs.⁴ The key idea of estimation is the same as that in the static case. Suppose that the econometrician has access to $t = 1, \dots, T$ observations of actions and strategies $(a_{1t}, \dots, a_{nt}, s_t)$ from repeated plays of the game. In the first step, we flexibly estimate the choice probabilities $\hat{\sigma}_i(a|s)$, which is consistent when only one equilibrium outcome is observed. One should choose a flexible parametric specification or a more flexible method such as a sieve logit, as discussed previously.

⁴In principal, it is possible to non-parametrically estimate the model. See Bajari, Chernozhukov, Hong, and Nekipelov (2009).

As in the static model, we substitute $\widehat{\sigma}_i(a|s)$ for $\sigma_i(a|s)$ in the agent's dynamic programming problem. Agent i solves the following DP problem:

$$\widehat{W}_i(s, \epsilon_i; \theta, \widehat{\sigma}) = \max_{a_i \in A_i} \left[\Pi_i(a_i, s; \theta) + \epsilon_i(a_i) + \beta \int \sum_{a_{-i}} \widehat{W}_i(s', \epsilon'_i; \theta, \widehat{\sigma}) g(s' | s, a_i, a_{-i}) \widehat{\sigma}_{-i}(a_{-i} | s) f(\epsilon'_i) d\epsilon'_i ds' \right]. \quad (20)$$

Note that conditional on the first stage estimate $\widehat{\sigma}_i(a|s)$, this is a single agent dynamic programming problem! This is an extremely well studied set of problems and can be estimated using a variety of standard techniques, such as the nested fixed point algorithm or the approach of Hotz and Miller (1993) (It can also be handled using the methods exposted in the excellent article by Rust (1997)). Using the method of Hotz and Miller (1993), for example, it is possible to quickly and accurately estimate θ for problems even with several state variables.

The asymptotic distribution theory for the above model is slightly more complicated, since it involves the treatment of how $\widehat{\sigma}$ converges to σ as the sample size becomes large. This issue is treated more carefully in various papers including Pesendorfer, Schmidt-Dengler, and Street (2008), Aguirregabiria and Mira (2007), Berry, Pakes, and Ostrovsky (summer 2007) and Bajari, Chernozhukov, Hong, and Nekipelov (2009). See Rust (1994) for a survey.

6.2 A framework for identification

In practice, it is often desirable to modify the standard single agent methods described in Rust (1994) to accomodate the special features encountered in dynamic games. For example, Pesendorfer, Schmidt-Dengler, and Street (2008) describe how to choose efficient weighting matrices to generate more efficient estimators. Aguirregabiria and Mira (2002) propose a method for efficient estimation based on iterating on multistep estimators. That is, after they estimate θ , they find the decision rule that is consistent with this parameter and maximize the likelihood function again. Bajari, Chernozhukov, Hong, and Nekipelov (2009) propose a sieve estimation strategy that allows $\widehat{\sigma}_i(a|s)$ to be specified as a sieve estimator. Following Ai and Chen (2003), they construct a three stage least squares like procedure that is semiparametrically efficient.

It is also possible to estimate dynamic game models with a mix of continuous and discrete strategies. This is useful in setting where choice variables, such as investment, are not well modeled as discrete decisions.

In the rest of this section we describe a framework for identification and estimation that follow the footsteps of Pesendorfer, Schmidt-Dengler, and Street (2008) and Aguirregabiria and Mira (2007) and that was described in Bajari, Chernozhukov, Hong, and Nekipelov (2009). This framework treats single agent dynamic discrete choice models, static and dynamic discrete games within the same model, under the incomplete information assumption. For dynamic discrete games, Pesendorfer and Schmidt-Dengler

(2003) demonstrate existence of equilibrium for discrete state variables, while Jenkins, Liu, McFadden, and Matzkin (2004) show existence of equilibrium for continuous state variables.

As section 3, in each of these models, the reduced form identified feature of the data is the set of choice probability functions $\sigma_i(a_i|s)$, for all i , a_i and s , together with the transition probabilities $g(s'|s, a)$ in a dynamic setting. A variety of estimators $\hat{\sigma}_i(a_i|s)$ can be obtained by nonparametric or flexible parametric methods. In general the data reveals the joint distribution of $a_i, i = 1, \dots, n$ conditional on s . However, the (clearly testable) assumption that the shocks are private information and are independently distributed implies the conditional independence of the actions across i given the state variables. The conditional correlation structure between the a_i 's given s can have identifying information in the presence of unobserved heterogeneity or under the private information assumption.

In the rest of this section we will assume that the joint distribution of the error term $f(\epsilon)$ is known and takes the standard i.i.d type I extreme value form in a multinomial logit model. We make this assumption because it is not possible to identify both the payoff function and the error distribution entirely nonparametrically. However, ongoing work by Nekipelov et al (2010) shows that if one is willing to impose a parametric assumption, such as a linear index structure on the mean payoff function, the joint distribution of the error terms can also potentially be identified nonparametrically.

In both static and dynamic single agent discrete models, the object of identification is the mean expected payoff function $\Pi_i(a_i, s)$. In discrete games however, the object of identification is the structural payoff function $\Pi_i(a_i, a_{-i}, s_i)$ which is a function of the entire action profile for all players. Note that the structural payoff function for player i depends only on a subset s_i of all the observed state variables s , because it can not be identified without exclusion restrictions, as pointed out by Pesendorfer and Schmidt-Dengler (2003). Identification analysis seeks a mapping from set of the reduced form choice probabilities $\sigma_i(a_i, s)$ and the observed transition probabilities and the presumed discount rate in the dynamic setting to the set of structural payoff functions $\Pi_i(a_i, a_{-i}, s_i)$.

$$(\sigma_i(a_i, s), g(s'|s, a) \forall a_i, \forall i, \forall s, \beta) \mapsto \Pi_i(a_i, a_{-i}, s_i) \forall a_i, \forall i, \forall s.$$

The discount rate β is assumed to be known for now. An insight from Pesendorfer and Schmidt-Dengler (2003) and Aguirregabiria and Mira (2002) shows that if there are variables that affect the transition process but are excluded from the static mean utility functions. Intuitively, individuals with the same static utility but different transition probabilities should behave identically when they are both myopic, but will behave differently when they are both forward-looking.

In static models, the Hotz-Miller inversion in (13) identifies $\Pi_i(a_i, s) - \Pi_i(0, s)$, which in turn identifies $\Pi_i(a_i, s)$ because $\Pi_i(0, s)$ is normalized to zero. The exclusion restriction allows for the identification and estimation of $\Pi_i(a_i, a_{-i}, s_i)$ as described in the previous sections. In a dynamic model, the Hotz-Miller inversion in (13) instead identifies the difference in the choice specific value functions:

$$\hat{V}_i(a_i, s_t) - \hat{V}_i(a_i = 0, s_t) = \log(\hat{\sigma}_i(a_i|s_t)) - \log(\hat{\sigma}_i(a_i = 0|s_t)). \quad (21)$$

In the above, the choice specific value function, which is defined as,

$$V_i(a_i, s) \equiv \Pi_i(a_i, s) + \beta \int \sum_{a_{-i}} W_i(s', \epsilon'_i; \sigma) g(s' | s, a_i, a_{-i}) \sigma_{-i}(a_{-i} | s) f(\epsilon'_i) d\epsilon'_i ds',$$

represents the expected utility for player i from choosing the action a_i excluding the current period error term $\epsilon_i(a_i)$. Player i takes into account the uncertainties in both the transition probabilities and the actions of competing players when forming this expectation.

In contrast to a static model, it may not be a reasonable assumption to normalize $V_i(a_i = 0, s_t)$. In fact we will impose the normalization assumption $\Pi_i(a_i = 0, a_{-i}, s) = 0$ for all a_{-i} instead, which implies $\Pi_i(a_i, s) = 0$. But under this assumption, $V_i(a_i = 0, s_t) \neq 0$.

Therefore identification needs to proceed in several steps. In the first step, we need to identify the collection of $V_i(a_i, s_t)$. Given the Hotz-Miller inversion, it will be sufficient to identify $V_i(a_i = 0, s_t)$. In the second step, the collection of choice specific value functions $V_i(a_i, s_t)$ can be further inverted to obtain the static expected utilities $\Pi_i(a_i, s)$. These two first steps are applicable to both single agent dynamic discrete choice models and dynamic discrete games. Additional step is needed for discrete games to identify the structural payoff functions $\Pi_i(a_i, a_{-i}, s_i)$. This will follow from the previous sections on static models.

Consider step one, define the ex ante value function (also called McFadden's social surplus function)

$$\begin{aligned} V_i(s) &= \int W_i(s, \epsilon_i; \sigma) f(\epsilon_i) d\epsilon_i \\ &= \log \sum_{k=0}^K \exp(V_i(k, s)) = \log \sum_{k=0}^K \exp(V_i(k, s) - V_i(a_i, s)) + V_i(a_i, s). \end{aligned}$$

The choice specific value function can also be written as

$$V_i(a_i, s) \equiv \Pi_i(a_i, s) + \beta E[V_i(s') | s, a_i]. \quad (22)$$

If we take $a_i = 0$, this relation identifies $V_i(a_i = 0, s_t)$ because we normalize $\Pi_i(a_i, s) = 0$:

$$\begin{aligned} V_i(a_i = 0, s) &\equiv \beta E \left[\log \left(\sum_{k=0}^K \exp(V_i(k, s') - V_i(0, s')) \right) | s, a_i = 0 \right] \\ &\quad + \beta E[V_i(0, s') | s, a_i = 0]. \end{aligned} \quad (23)$$

The first term on the right hand side is identifiable and estimable from the data. Given knowledge of the first term, the above relation represents a contraction mapping that uniquely recovers the baseline choice specific value function $V_i(a_i = 0, s)$. Subsequently, $V_i(a_i, s)$ for all a_i can be identified and estimated. In step two, the collection of expected static payoff functions $\Pi_i(a_i, s)$ can then be identified and estimated as

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

This recovers the primitive static mean utilities nonparametrically for single agent discrete choice models. For dynamic discrete games, a final additional step recovers the primitive structural mean utilities $\Pi_i(a_i, a_{-i}, s_i)$ from the static expected mean utilities $\Pi_i(a_i, s)$ using same techniques based on an exclusion restriction in the static models.

In the above, the equation that relates the choice specific value functions to the social surplus functions, $V_i(s) = \log \sum_{k=0}^K \exp(V_i(k, s))$, depends on the Type I extreme value parametric assumption of the error terms. If another known error term distribution is used in place of the Type I extreme value value assumption, then as pointed out in Pesendorfer and Schmidt-Dengler (2003), a different function $V_i(s) = G(V_i(k, s), k = 0, \dots, K)$ needs to be used. Conditional on the use of the correct function $G(\cdot)$ which depends on the error distribution, the rest of the identification arguments goes through without modifications. Importantly, identification does not depend on distributional assumptions.

A multi-stage nonparametric estimation procedure follows straightwardly from the above identification arguments. Nonparametric estimators typically suffer from the curse of dimensionality when the number of continuous state variables is relatively large. With a small sample of the data the researcher might wish to specify the static payoff function parametrically, as $\Pi_i(a, s_i, \gamma)$ with finite dimensional parameter γ . For example, we can specify $\Pi_i(a, s_i) = \Phi_i(a, s_i)' \gamma_{i,a}$, where $\Phi_i(a, s_i)$ is a vector of known basis functions. The parameters $\gamma_{i,a}$ in this specification can be estimated by projecting the nonparametric estimates of $\Pi_i(a, s_i)$ on the basis functions $\Phi_i(a, s_i)$ using for example least square regressions. Because the transition probability $g(s'|s, a)$ is nonparametrically specified, a model with parametric static utility functions $\Pi_i(a, s_i) = \Phi_i(a, s_i)' \gamma_{i,a}$ is a semiparametric model.

However, this identification based estimation approach inherits the disadvantages of many multi-stage estimation techniques. First of all, the standard errors are hard to compute because of propagation of errors from the previous steps of the procedure which will depend on the degree of smoothness of the unknown functions of the model. Second, this multistage estimation procedure is not semiparametrically efficient. It is well known that it is difficult to design multistage estimation procedures that can achieve semiparametric efficiency bounds, because at each subsequent step has to compensate the estimation errors that will arise from previous estimation errors.

Before we further discuss issues related to estimation, we note that both identification and the optimal estimation strategy depends crucially on the type of data that is available, and on what assumptions are being made on the data generation process in the presence of multiple equilibrium. Both the static interaction model and dynamic games under the incomplete information assumption can generate multiple equilibrium. While a number of papers have investigated the cardinality of multiple equilibrium in the static interaction model, each under specific assumptions, the cardinality of multiple equilibria in dynamic discrete games appears to be a wide open issue.

In public finance and labor economics (e.g. Brock and Durlauf (2001)), typically a data set (such as a school) consists of a large number of players. In this case despite the possible of multiple equilibria, the

equilibrium that generates the data is also uniquely identified and estimatable from the data. In this case, the parameters in the individual utility functions can easily be estimated through the two step procedure by a standard textbook multinomial regression with the estimated equilibrium choice probability as one of the regressors on the right hand side. A nested maximum likelihood method typically will need to solve for all equilibria, compute the likelihood under all equilibria, and use the one that generates the highest likelihood to estimate the parameters. The two step approach is obviously a lot more straightforward.

These issues become a lot more difficult in empirical IO, one typically has access to either a cross-section of independent markets each with a smaller number of firms as players. In some applications (as pointed out in Pesendorfer and Schmidt-Dengler (2003)) a stationary time series of firms can also be available. In order for the multi-step identification and estimation strategy to be viable, it is necessary to assume that the same equilibrium is being played out in each of the markets in the data. However, the number of equilibrium is potentially endogenous and in general one can not make an assumption on an endogenous outcome. A proper approach is to formulate a hypothesis of unique equilibrium and to test it using the data. Without further testing, in general it is only possible to estimate the model using the time dimension of each market. If a model generates multiple equilibria, and each market (especially in spatially heterogeneous markets) plays a different equilibria, then this strategy breaks down. In this case, one has to either postulate or estimate an equilibrium selection mechanism, or make use of inference methods, such as those that derive inequality bounds on the structural parameters, that are agnostic to the process through which multiple equilibria are being selected. There is a vast literature on bounds and inequality based models that are applicable to discrete dynamic games, which can also potentially be applicable to dynamic games. In this article we will focus mostly on methods that point identifies and estimates the parameters of interest. We refer the readers to the insightful articles by Tamer (2003) and Berry and Tamer (2006), and the related literature discussed therein. If a panel data with large cross sectional and time series dimensions are available, it is known that the same equilibrium is played out in each market over time, and different equilibria can be played out in different markets. In this case, one can obviously estimate the model market by market, and then try to find a way to improve estimation efficiency by pooling information about the parameters across the markets.

In the following discussion about estimation, we confine ourselves to the case with a large number of markets that are best interpreted in the time dimension, each of which has a small number of players.

6.3 Efficient method of estimation

If the period payoff functions and the transition probability distribution are both parametrically specified, the nested fixed point maximum likelihood estimator of Rust (1987) will automatically deliver an efficient estimator that achieves the asymptotic Cramér-Rao lower bound. The nested fixed point maximum likelihood estimator is however computationally intensive if not infeasible. Aguirregabiria and Mira (2007)

and Pesendorfer, Schmidt-Dengler, and Street (2008) show that a nested pseudo maximum likelihood method and an asymptotic least square method can both be used to estimate the parameters flexibly. In particular, Pesendorfer, Schmidt-Dengler, and Street (2008) show that an asymptotic least square method achieves the parametric efficient information bound but the nested pseudo maximum likelihood estimator is not efficient. In a semiparametric or nonparametric model, either or both of the static period utility function and the transition probability densities are nonparametrically specified. One would conjecture that a flexible sieve approach, which uses a flexible parametric approach to approximate the underlying nonparametric function, will also achieve the properly defined semiparametric or nonparametric efficiency bound if certain regularity conditions are met and if the degree of flexibility of the parametric approximation increases as the sample size increases at appropriate rates. However, this needs to be proven.

The approach that we are taking to obtain a flexible one-step efficient estimator for both semiparametric and nonparametric specifications of the model is motivated by recent developments in the literature of sieve estimation and conditional moment models. Intuitively, if we can write down a set of conditional moment restrictions that incorporates all the information provided by the model, then we can appeal to the vast literature of conditional moment models with finite and infinite dimensional parameters to obtain estimators that are efficient in the sense of achieve the semiparametric efficient bound. More precisely, this means that there is no other asymptotically regular estimator that will have a smaller asymptotic variance and that does not impose a stronger parameter assumption beyond those that are already in the model. The finite dimensional parameters are the coefficients of the linear mean utility index. The conditional choice probabilities and the value functions can both be considered infinite dimensional parameters that can be approximated by a flexible sieve structure. When the support of the state variables is discretely finite, Pesendorfer and Schmidt-Dengler (2003) is the first to develop the insight that the collection of values in the support of the discrete state variables generates a set of moment conditions that can be used to efficiently estimate the structural parameters in a generalized method of moment framework. We further extend this insight to allow for continuously distributed state variables, which can naturally be casted into a sieve conditional moment framework.

The main observation is that the value function iteration itself can be viewed as a conditional moment restriction. Note that $\Pi_i(a_i, s) = E(\Pi_i(a, s_i; \gamma) | a_i, s)$, and therefore

$$V_i(a_i, s) = E \left[\Pi_i(a, s_i; \gamma) + \beta \log \left[\sum_{l=0}^K \exp(V_i(l, s')) \right] \middle| a_i, s \right]. \quad (24)$$

This in addition to the second set of conditional moment restrictions

$$E(1(a_i = k) | s_{m,t}) = \sigma_i(k | s_{m,t}) \quad (25)$$

exhausts all the information contained in the model. To reduce the total number of parameters, we can make use of the relation $\sigma_i(a_i = k | s) = \frac{\exp(V_i(k, s))}{\sum_{l=0}^K \exp(V_i(l, s))}$, for $i = 1, \dots, n$ and $k = 0, \dots, K$ to rewrite (24)

and (27) as

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

and

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

In the above $m = 1, \dots, M$ indexes markets and $t = 1, \dots, T$ indexes time. $d_{m,t}^{a_i,k}$ is a dummy variable that takes value one if a_i takes value k in the m th market at time t , and takes zero otherwise. The moments (26) and (27) implicitly assume a balanced panel of data. (26) holds for $t = 1, \dots, T - 1$ while (27) holds for all T . In principle, $T = 2$ is sufficient for identification and estimation in an infinite horizon stationary model. The second period data is only needed for identifying the transition process through (26). Once the transition density is estimated, only a cross section is required for estimating the structural parameters. The unknown parameters in (26) and (27) include the finite dimensional structural parameters γ and the unknown infinite dimensional functions $V_i(0, s)$ and $\sigma_i(k|s)$.

The literature of sieve estimation and conditional moment models (e.g. Newey and Powell (2003), Ai and Chen (2003), Chen and Pouzo (2009) suggest that the unknown functions can be approximated by flexible parameter functions such that the number of parameters increase to infinity at appropriate rates when the sample size increases. Given this approximation, (26) and (27) essentially constitute of a parametric conditional moment model, which can be estimated consistently and efficiently using for example an increasing set of instrumental variables which are functions of the state variables. In particular, in practical terms, the efficient estimators proposed in this literature essentially amount to nonlinear three stage least square estimators, which are straightforward to implement on most numerical softwares. It is possible to show that while the infinite dimensional function estimates converge at slower nonparametric rates, the finite dimensional parameter estimates are still \sqrt{n} consistent and achieve the semiparametric efficiency bound. In particular, recent work by Akerberg, Chen, and Hahn (2009) demonstrates that an estimator for the asymptotic variance of a semiparametric two step estimator that is based on assuming that the model is purely parametric is in fact consistent for the asymptotic variance of the semiparametric model. Yet Hong, Mahajan, and Nekipelov (2009) shows that if numerical methods are used to estimate the asymptotic variance, one has to be careful with the step size that is used in numerical differentiation. They give a set of weak conditions of the step size as a function of the sample size in order for the asymptotic variance estimates to be consistent, and suggest adaptive methods to obtain the optimal step size.

While this procedure is efficient when the transition density is nonparametrically specified, it can be improved when the transition density takes a parametric form. In this case, a parametric likelihood

estimator is likely to be more efficient. The moment condition in (26) can also be more efficiently computed by integrating against the parametric transition density.

As discussed above, we have been working under the assumption of either a unique equilibrium in the model or the assumption that the data is generated by only one equilibrium even when the model can generate multiple equilibria. A recent approach suggested by Aradillas-Lopez (2005) overcomes this problem by focusing on the subset of observations that admits only one equilibrium for all possible value of the parameters to estimate the structural payoff parameters. This approach requires that the conditions for a unique equilibrium can be checked for each value of the covariates and parameters, and that the set of covariates for which the unique equilibrium condition is satisfied is not empty.

In complete information static games, Bajari, Hong, and Ryan (2004) and the literature cited therein postulates and estimates an equilibrium selection mechanism. This appears to be relatively easier for static complete information games because equilibria there are determined by polynomial equalities and inequalities. We are not aware of papers that take a similar approach for incomplete information discrete games or dynamic discrete games. Part of the difficulty of applying this approach to incomplete information games is the difficulty of characterize the cardinality of multiple equilibria in incomplete information games. While the number of equilibrium is generically finite in these models, it is difficult to derive an explicit upper bound.

Even if a researcher is willing to make the assumption that only one equilibrium is being played out in the data or in a market, the computational ability to solve for multiple equilibria of the model is still important for conducting global post estimation counterfactual policy evaluations. One approach to equilibrium computation is the constrained optimization method of Doraszelski, Judd, and Center (2008) and Judd and Su (2006). The all-solution homotopy method, discussed by Bajari, Hong, Krainer, and Nekipelov (2009) and Borkovsky, Doraszelski, and Kryukov (2009) among others, can also be used to compute asymptotic approximations to the sets of equilibria with relatively low computational costs.

7 Extensions of the incomplete information discrete game

7.1 Nonstationary data and finite horizon

The assumptions of stationary and infinite horizon are often plausible in empirical industrial organization, such as in entry models where firms are assumed to compete period by period for the indefinite future. These assumptions are less plausible in other areas, such as in labor economics, public finance and health economics when modelling individual behaviors. Formally, a finite horizon model can be both stationary, where the period utility functions are the same, or nonstationary, where the utility functions are different from period to period. In a finite horizon model, because the end period choice probabilities only identify the difference in the utilities across choices, it is necessary to normalize the utility of one of the choices

in the last period. Without bequest incentives the end of period utilities can often be normalized to zero. It is difficult to identify bequest incentives in a finite horizon model nonparametrically without imposing strong parametric assumptions. The stationary infinite horizon model can be considered a limiting case of the stationary finite horizon model where the number of time periods increases to infinity. When a panel data set is available, as shown in Taber (2000) and Magnac and Thesmar (2002), a nonstationary model where the static utilities are different from period to period is exactly identified when there is no bequest incentives and when the utility of one of the choices in the end period is normalized to zero. A stationary model is clearly over-identified in the presence of panel data, because the data from each period can be used to estimate the same static utility function. Furthermore, while normalizing the static profit of not entering the market is a sensible assumption for entry games, normalizing the utility of any one of the choices to zero may not be reasonable in some other applications. In static discrete choice models, it is known that only the differences in the utilities across the choices are identified. In dynamic models, it can be shown that only a linear combination of the utilities under different choices across the state variables can be identified. The exact form of the linear combination depends on the transition probability matrixes. For these models, a parametric specification of the static period utility functions that is ex ante plausible can be a more reasonable approach than a nonparametric specification in which the utility of one of the choices is forced to be equal to zero. Of course, to improve the degree of nonparametric identification, one can also consider exclusion restrictions in which some state variables only enter a subset of the choice utility functions.

7.2 Discrete observations and continuous time models

The previous model of dynamic discrete models and dynamic discrete games that we discussed operates in equal-spaced discrete time intervals. Often times individual behaviors are reported in prespecified time intervals in each data set. This can be daily, monthly, quarterly or annually, etc. This requires that a researcher adapts the dynamic discrete choice model to the observed time intervals reported in the data set. Hence the discount rate that is used and that is estimated has to be based on the same time intervals of individual behavior that is observed in the data, which can be daily, monthly, quarterly, etc. While the presumed or estimated discount rates can be converted into other time interval lengths, it is not clear that doing so is internally coherent within the model. An econometric model of dynamic discrete choice is necessarily misspecified when the time interval reported in the data and used in formulating the dynamic model is inherently different from the actually time interval at which individuals make discrete decisions.

For example, if in the true model, firms decide whether to enter the market at the beginning of every year. However, researchers only observe the entry behavior of firms once every two years, and typically will estimate a dynamic discrete game with the two year time interval and a discount rate that is ex ante suitable for the two year interval. The transition process is estimated for the two year interval. This

obviously is misspecified, and will not give a consistent estimate of the structural payoff functions.

There is a possible solution in the context of a stationary infinite horizon model. Consider, for example, the case when the researcher only observes the data once every T period of the individual choice behavior. Note that the entire nonparametric identification argument only depends on estimating the reduced form choice probabilities and the transition process consistently. Given the stationary assumption, the choice probabilities can be still consistently estimated from the observed data. While the single period transition process is not directly observed, it can be recovered from the observed T -th period transition probabilities. When the state variables are discrete, the single period transition probability matrix is closely related to taking the T th-root of the T th period transition probability matrix. If the state variables are continuously distributed, the single period transition density is related to the T th period transition density through an integral equation that has a unique solution. A nested fixed point maximum likelihood method, if the likelihood function for parameters in the T th period transition process is correctly specified, is obviously consistent. As in single period observation data, the likelihood can also be factored into the product of two pieces. One piece is the likelihood for the transition process, now correctly specified to reflect the T -th period observation feature of the data. The second piece is the likelihood for the observed choice behavior which is identical to single period models.

A "correct" two step nonparametric or semiparametric procedure for estimating a dynamic discrete choice in which the time interval of the observed individual behavior is a multiple of the actual length of time interval at which point individuals make decisions is the following. First the T -th period transition density or probability matrix is estimated from the data, either nonparametrically or using a flexible parametric model. Then an estimate of the single period transition density or probability matrix can be obtained by inverting the T -th period transition density.

From this point on there are two choices, a nonparametric one and a parametric one. In the first choice, we can follow the nonparametric identification and estimation procedure. First the reduced choice probability is estimated from the data, then it is used in combination with the single period transition process to recover the static structural payoff functions based on the presumed single period discount rate. In the second choice, one might wish to specify the static utility functions parametrically and adopt an estimation approach that has more of a parametric flavor. To do that, a researcher can simulate data from the estimated single period transition process and the estimated choice probabilities, and then apply one of the several possible two step estimation methods, such as Hotz and Miller (1993), Aguirregabiria and Mira (2007) and Pesendorfer, Schmidt-Dengler, and Street (2008) to the simulated data.

The drawback of discrete time models can be overcome in the continuous time dynamic games developed in Doraszelski and Judd (2007) and Bayer, Arcidiacono, Blevins, and Ellickson (2010). In the continuously time discounting game models developed by these authors, players move sequentially because of the nature of the continuous time game where the probability of simultaneous moves is zero. When a player gets to make a move is randomly determined by nature. Bayer, Arcidiacono, Blevins, and

Ellickson (2010) shows that two step estimation methods can also be used to estimate continuous time games flexibly. Since firms move randomly and sequentially, the decision problem faced by each firm is more straightforward than the need to compute equilibrium fixed points in simultaneous move games. The arguments for nonparametric identification in discrete time games also extends easily to the continuous time. The Hotz-Miller inversion in both cases recovers the difference in the choice specific value functions between a given choice and a baseline choice. The difference is that while in discrete time, the expectation operator in the contraction mapping in (23) that maps the difference in the choice specific value function to a baseline choice specific value integrates over the transition density of future state variables conditional on the current choice and current state variables, a related expectation operator for such a contraction mapping also has to integrate with respect to the uncertainty with respect to the next even time. In other words, in continuous time, (22) only needs to be replaced by

$$V_i(a_i, s) \equiv \Pi_i(a_i, s) + \int e^{-\rho t} E[V_i(s') | s, a_i, t] f(t) dt, \quad (28)$$

where $f(t)$ is nature's distribution of the random arrival time of the next move. Both the conditional expectation above and $f(t)$ can be identified and presumably estimated from the data. Subsequently, if one is willing to assume a normalization, a similar replacement of the conditional expectation operator in (23) identifies the baseline choice specific value function, which can be plugged back into (28) to back out all the flow utilities. Obviously there is a remaining issue of how to compute or estimate the expectation and integral in (28) efficiently and precisely. Without the need of solving for the fixed point that defines the equilibria in simultaneous move games, the sequential move continuous time model of Doraszelski and Judd (2007) and Bayer, Arcidiacono, Blevins, and Ellickson (2010) is also far easier to simulating counterfactual policy outcomes, and does not suffer from the curse of multiple equilibria.

7.3 Unobserved heterogeneity

If some of the state variables s that are observed by all the players in the model are not observable by the econometrician, an estimation procedure that does not take into account the presence of unobserved heterogeneity will not be consistent. To clarify notation, we decompose s into (s^o, s^u) where now s^o denotes the set of state variables that are observed by both the players in the model and the econometricians, but s^u denotes the state variables that are only observed by the players but not the econometricians. We will focus on the cases where the econometricians always have less information than the players in the market. Given that econometric analyses are often conducted many years after the market interaction takes place, it can also be conceivable that the econometricians might have more information at the time of analysis than the players at the time of market interaction. Such a scenario is considered for example in Aradillas-Lopez (2005). Both or either of the period mean utility functions $\Pi_i(a_i, a_{-i}, s^o, s^u)$ and the transition process $g(s^{o'}, s^{u'} | s^o, a, s^u)$ can potentially be contaminated by the unobserved heterogeneity component s^u . While the discount rate can also depend on observed and unobserved heterogeneity, the

current literature has typically considered the discount rate as a constant. Unobserved heterogeneity raises substantial challenges for identification and estimation in many models, is a burgeoning area of exciting research in the context of dynamic discrete models.

Identification A key insight in Kasahara and Shimotsu (2008) and Hu and Shum (2008) is that the identification arguments (when all state variables are observed) in Pesendorfer and Schmidt-Dengler (2003) and Bajari, Chernozhukov, Hong, and Nekipelov (2009) only depends on the reduced form choice probabilities $\sigma_i(a|s)$ and the transition process $g(s'|s, a)$. Indeed identification can be considered a mapping from $\sigma_i(a|s)$ and $g(s'|s, a)$ and the discount rate to the structural mean payoff functions $\Pi_i(a_i, a_{-i}, s)$ when the discount rate is known, or when the discount rate is estimated, a mapping from $\sigma_i(a|s)$ and $g(s'|s, a)$ to $\Pi_i(a_i, a_{-i}, s, u)$ and the discount rates. Hence, analogously, in presence of unobserved heterogeneity variable s^u , if we find a way to identify or flexibly estimate $\sigma_i(a|s^o, s^u)$ and $g(s^{o'}, s^{u'}|s^o, a, s^u)$, then we can follow the identification procedure to recover or nonparametrically estimate $\Pi_i(a_i, a_{-i}, s^o, s^u)$ and possibly the discount rate parameters. Intuitively, we can simulate data repeatedly from $\sigma_i(a|s^o, s^u)$ and $g(s^{o'}, s^{u'}|s^o, a, s^u)$, and apply a variety of parametric or nonparametric single step or multi-step estimators to the simulated data to recover the structural parameters. Hence both Kasahara and Shimotsu (2008) and Hu and Shum (2008) focus on nonparametric identification of the reduced form parameters $\sigma_i(a|s^o, s^u)$ and $g(s^{o'}, s^{u'}|s^o, a, s^u)$. The general results by Hu and Shum (2008) can allow for general unobserved heterogeneity that can be serially correlated and vary over time, and can be both discrete and continuously distributed.

Estimation Bayes methods can be particularly useful for estimating fully parametric dynamic discrete models with unobserved observed state variables with or without serial correlations. Norets (2009) and Imai, Jain, and Ching (2009) pioneered the use of Bayesian estimators for dynamic discrete models. Both make use of Monte Carlo Markov Chain simulations to simulate the posterior distributions of the parameters. Even in a model where the unobserved state variables are i.i.d. over time, Imai, Jain, and Ching (2009) show that Bayesian methods can be a viable alternative to maximum likelihood estimation, and provide a procedure in which the value function computation is only once for each draw of the parameter value along the MCMC chain. Norets (2009) makes use of Gibbs Samplers to incorporate serial correlated unobserved state variables in the model, and improve the computation of the value function iteration using an importance sampler. It has been shown in Chernozhukov and Hong (2003) that MCMC methods are applicable regarding whether one takes a frequentist or a Bayesian viewpoint to inference, that the posterior mean and median of the parameters are as efficient as maximum likelihood estimators as point estimators, and that the posterior percentile intervals are also asymptotically correct confidence intervals when the model is correctly specified.

The likelihoods of these models, which are proportional to the posterior distribution of the parameters

and unobserved state variables multiplicatively up to the prior density, are factored into the product of the transition density (of both observed and unobserved state variables) and the probability of the choices conditional on contemporaneous state variables. It can therefore be used in a Metropolis-Hastings in Gibbs sampler to make draws from the posterior distribution of the parameters conditional on the data and the unobserved state variables, and the posterior distribution of the latent state variables conditional on the parameters and the data. As noted in Norets (2009), when an element of the unobserved state variables is i.i.d over time with a given parametric distribution that can be integrated analytically to give an explicit formula for the choice probability conditional on all other remaining state variables, such as in the Rust model with a Type I extreme value i.i.d shock, it is only necessary to use the conditional choice probabilities given all the other remaining state variables in the likelihood factorization.

The computational difficulty of Bayesian estimators in dynamic discrete choice models comes in three parts. The first part is computing the expectation of the next period social surplus function given the state variables and actions in this period during the value function iteration process. This conditional expectation estimation can be implemented using kernel methods (Imai, Jain, and Ching (2009)), with nearest neighborhood and importance sampler (Norets (2009)), or given that the model is fully parametric, with numerical quadrature routines (Gallant, Hong, and Khwaja (2008)). Computing the value function is needed to form the part of the likelihood for the choice probabilities given the parameters and the vector of state variables. Both Imai, Jain, and Ching (2009) and Norets (2009) consider only single agent dynamic discrete models. The extension of their methods to dynamic games requires a solution algorithm for the Markov perfect dynamic equilibrium in combination with the value function iteration at each simulation draw for the parameters and the state variables. A dynamic equilibrium selection mechanism can be potentially far more computationally challenging than value function iteration, and also has a potential issue with finding more than one solution leading to multiple equilibria.

The other two parts of computation challenge are simulating from the conditional posterior distributions of the parameters and latent state variables. While both Imai, Jain, and Ching (2009) and Norets (2009) use MCMC methods for both of these components, sequential importance sampling methods can also be applied. There are at least four possible combinations of choices. For example, the first choice is to use nonlinear sequential importance sampler, which is often called particle filter, to sample from the conditional posterior distribution of the latent state variables, while MCMC methods are used to sample from the conditional posterior distribution of the parameters. A second choice is to use sequential importance samplers for both the parameter and the latent state variables. The third choice is to use MCMC for latent state variables and importance sampler for the parameters. The fourth choice is to use MCMC for both, as in Imai, Jain, and Ching (2009) and Norets (2009). The second and fourth choices are natural because in a Bayesian setting, both the "parameters" and the latent state variables are random and are not fundamentally different. Therefore it is natural to apply the same sampling methods to both of them. In implementations, however, it is possible that the first one might have some statistical and

computational advantage. It can be very difficult to simulate from the parameter space using an importance sampler, especially when the sample size is reasonably large so that the posterior distribution of the parameters is concentrated on a small area that can be difficult to find. In this case, it is difficult for an importance sampler to find the area of positive posterior densities for parameters. By walking through the parameter space in small increments and by only using the difference in the log likelihood between adjacent trial parameter values, MCMC methods are more likely to be able to visualize the shape of the positive density area of the posterior distribution of the parameter values. The posterior distribution of the latent state variables, however, are more likely to be amenable to a sequential importance sampler. This is because the dimension of the latent state variable distribution is likely to be much larger and more dispersed than the parameters, and does not tend to concentrate even when the sample size increases without bound. A priori, it seems that the third choice is unlikely to be used in implementations.

The ultimate goal of inference is the posterior distribution of the model parameters conditional on the observed data of state variables and choice outcomes. The Gibbs sampler that iterates between the parameters and the latent state variables. Instead of using the Gibbs sampler to iterate between two conditional posterior distribution, another method is to integrate out the latent state variables conditional on the parameters and the observed data, to obtain the marginal conditional likelihood of the parameters given the observed data. Integration with respect to the latent state variables can also be done by either an importance sampling algorithm or by MCMC, or by some combination of both sampling schemes. The current state of the literature, however, seems to be focusing on using sequential importance sampling algorithms, or particle filters, to integrate out the latent state variables in forming the marginal conditional likelihood for the parameters.

If a researcher opts for using the integration method to obtain the likelihood for the parameters based on observed data, there are also several choices regarding how to process the likelihood. First the likelihood can be run through a MCMC chain, or an importance sampler to obtain the posterior distribution of the parameters. Any notion of a location of the posterior distribution, such as the mean or the median, can be used as a point estimator for the parameters. MCMC methods are widely used in practice. In contrast, importance samplers do not seem to be the method of choice for sampling from the parameter space. One can obviously maximize the likelihood instead of simulating from it. This will give rise to a frequentist maximum likelihood estimator.

Particle filtering methods are introduced by Fernández-Villaverde and Rubio-Ramírez (2007) to economics for estimating nonlinear dynamic stochastic general equilibrium models in macroeconomics. They show that this overcomes the errors from approximating a complex DSGE model by linearization. They use the particle filters to form the likelihood which they maximize to obtain the maximum likelihood parameter estimates. A recent paper by Blevins (2008) proposes using particle filters to estimate dynamic discrete games under the incomplete information assumption. Blevins (2008) also proposes to incorporate the particle filters with Hotz-Miller style multiple step estimation methods that are based on inverting

the conditional choice probabilities. A closely related approach in the frequentist maximum likelihood domain is the EM (expectation-maximization) algorithm developed by Arcidiacono and Miller (2006). The EM algorithm iterates between estimating the parameter given a distribution of the latent state variables and computing the expected distribution of the latent state variables given each iteration of the parameter estimates. In the likelihood setting, it is known that each iteration of the EM algorithm will not decrease the likelihood value and contains the MLE as a stationary point. Arcidiacono and Miller (2006) show that the EM algorithm also has superior properties for Hotz-Miller style two step estimators.

Even if alternative non-Bayesian methods, such as MLE, GMM or a variety of multi-stage estimators, are used to parameters, the particle filtering algorithm can still be useful when the distribution of the latent variables itself is of interest. This might not be useful per se in some applications, such as in Heckman and Singer (1984), where the latent unobserved heterogeneity variables are mostly treated as nuisance parameters that are to be integrated out when the likelihood for the key structural parameters is being constructed. But in some other examples, for example in some marketing applications, the latent variables can be unobserved consumer taste parameters. It should be interesting to use the sequential importance sampler for filtering out the distribution of the unobserved taste variables conditional on knowledge of the parameters and observe data set, particularly if one is interested in making out of sample forecasts.

For a given parameter value, and a serially correlation set of observations of length T , the particle filter updates through $t = 1, \dots, T$ stages. At each stage t the particle filter has a particle generation step and a particle reweighting step. The particle generation step corresponds to the transition density of the state variables given the current state variables and the actions observed in the data up to stage t . The weights are the choice probabilities of the observed actions given the state variables. At the end of each stage t , the distribution of the sequentially generated and reweighted particles represents the posterior conditional distribution of the sequence of latent state variables up to time t conditional on the state variables and the choice outcomes up to time t , at the given parameter value. The usual setup of a particle filtering algorithm consists of a state variable transition equation which represents the evolution of the state variables over time and a measurement equation which relates distribution of the observed outcome or dependent variable to the realized state variable. In some applications, such as in some GARCH style stochastic volatility models in financial econometrics, the state variables (the latent unobserved volatility) evolves autonomously over time. The observed outcome (price) represents the measurement equation. In both the DSGE models of Fernández-Villaverde and Rubio-Ramírez (2007) and the dynamic discrete models of Blevins (2008) and Gallant, Hong, and Khwaja (2008), however, the evolution of the state variables is not autonomous but instead depends on the feedback from the observed actions. A particle filter can be properly setup to handle both situations.

Semiparametric Estimation As each of the papers reviewed above shows, incorporating unobserved heterogeneity can be challenging. With the exception of Arcidiacono and Miller (2006) and Blevins (2008), the main computational challenge comes from calculating the dynamic equilibrium. The insight of the identification results of Hu and Shum (2008), however, implies a semiparametric approach for incorporating serially correlated unobserved heterogeneity in the estimation procedure. Their insight is based on the sequential factorization of the likelihood function of the complete data into the choice probabilities $\sigma_i(a_i|s, u)$ and the transition densities $g(s', u'|a, s, u)$. Consider a flexible sieve parameterization of $\sigma_i(a_i|s, u, \theta_n)$ and $g(s', u'|a, s, u, \gamma_n)$. If the dimensions of both θ_n and γ_n increase to infinity at appropriate rates when the sample size increases, as dictated by the theory of sieve estimation, then both $\sigma_i(a_i|s, u, \theta_n)$ and $g(s', u'|a, s, u)$ should be consistently estimated by $\sigma_i(a_i|s, u, \hat{\theta}_n)$ and $g(s', u'|a, s, u, \hat{\gamma}_n)$. Once consistent flexible nonparametric or parametric estimates of $\sigma_i(a_i|s, u)$ and $g(s', u'|a, s, u)$ are available, they can be passed through the inversion mapping of the identification argument to obtain the structural mean payoff functions $\Pi_i(a|s, u)$, and possibly the discount rate parameter β . The "reduced form" sieve parameters θ_n and γ_n , the dimension of which now depends on the sample size, can be flexibly estimated using the variety of estimation algorithm, including MCMC, data augmentation, sequential importance sampler, or the EM algorithm. In the next step, one can then either nonparametrically or parametrically invert out $\hat{\sigma}_i(a_i|s, u)$ and $\hat{g}(s', u'|a, s, u)$ to estimate the mean structural payoff function $\Pi_i(a|s, u)$; or one can simulate data from $\hat{\sigma}_i(a_i|s, u)$ and $\hat{g}(s', u'|a, s, u)$ and then apply any of the single step or multiple step estimators discussed previously to estimate the structural mean payoff functions. In the very special case where unobserved heterogeneity is only in the period payoff function but not in the transition density, $g(s'|s, a)$ can obviously be estimated in a first step either nonparametrically or semiparametrically without any complexities. In another case, with time-invariant unobserved heterogeneity captured by the transition density $g(s'|s, u, a, \hat{\gamma})$, then the transition density, integrated against the unobserved heterogeneity distribution, can be used in the first stage to make inference on the transition density parameter $\hat{\gamma}$ and the mixing distribution $f(u)$ using a variety of estimation approaches.

In addition to estimating the choice probabilities $\sigma_i(a_i|s)$ and $g(s'|s, a)$ either parametrically or nonparametrically from objectively observed choice and transition data, another possibility is to estimate them from subjective self-reports in survey data sets, which are becoming increasingly more available. See e.g. Wang (2009). Pantano and Zheng (2009) show how to make use of subjective expectations data to incorporate unobserved heterogeneity in dynamic discrete choice models.

In dynamic auction models with continuous choice variables, Jofre-Bonet and Pesendorfer (2003) are the first to estimate the reduced form observable bid distribution using a flexible parametric model, and use it to invert out the underlying structural distribution of the evaluations.

Under the incomplete information assumption, equilibrium computation can only be more daunting in dynamic games than in static games. However, it is possible that the issue of multiple equilibria is

less prevalent in dynamic games than in static games. Even in static games, not much work has been done to handle multiple equilibria in the data in the incomplete information setting and allow for point identification. The private shocks are often assumed to be independent across players in the incomplete information models. Models in which the joint distributions of the private shocks are correlated across players necessarily pose more challenges for both identification and estimation, even in the presence of unobserved state variables.

8 Complete information models

There are two types of variables unobservable to the economist in incomplete information games with unobserved heterogeneity, those that are observed by all the players in the market and those that are only observed by a single player which represents the private information of this player. A complete information model only contains unobserved heterogeneity that are observed by all the players. For particular applications, perhaps a realistic situation contains unobserved heterogeneity that are observed by a subset of the players but not the others. But we are not aware of such applications yet. Complete information normal form games are widely used to model entry behaviors of competing firms.

While multiple equilibria is a generic feature for both complete information and private information games, it has been studied more extensively in the context of normal form games. The literature has suggested three different approach to address multiple equilibria issues. The first approach is to introduce an equilibrium selection mechanism that specifies which equilibrium is picked as part of the econometric model. Examples include random equilibrium selection in Bjorn and Vuong (1984) and the selection of an extremal equilibrium, as in Jia (2008b). The second approach is to restrict attention to a particular class of games, such as entry games, and search for an estimator which allows for identification of payoff parameters even if there are multiple equilibria. For example, Bresnahan and Reiss (1990; 1991) and Berry (1992) study models in which the number of firms is unique even though there may be multiple Nash equilibria. They propose estimators in which the number of firms, rather than the entry decisions of individual agents, is treated as the dependent variable. A third method which we will not further discuss, pioneered by Tamer (2003), uses bounds to estimate an entry model. Yet another approach, forcefully developed in Pesendorfer and Schmidt-Dengler (2003), is to use the time dimension of each market.

Bajari, Hong, and Ryan (2004) makes an attempt to estimate an equilibrium selection in a static complete information discrete game that also allows for mixed strategy equilibrium, using a simulated method of moment estimator which simulated the conditional choice probabilities given the parameter values for each observation. They consider a data generating process in which the observed choice probabilities in the data is rationalized as a mixture over the choice probabilities predicted by each of the multiple equilibria. The mixing distribution is the equilibrium selection mechanism and is parametrically specified and estimated in Bajari, Hong, and Ryan (2004). This postulated data generating process does

imply the existence of a coordination mechanism that a researcher is agnostic about, because players in each market are assumed to observe the particular realization of the equilibrium that is dictated by the mixing distribution of the equilibrium selection mechanism and are able to act on it collectively.

Modeling and estimating equilibrium selection mechanisms in this sense requires computing all equilibria, including pure strategy and partially and fully mixed strategy ones, at each parameter value and for each draws of the latent utilities in the simulation procedure. Recent algorithms that compute all equilibria in normal form games are developed in McKelvy and McLennan (1996) and are computationally intensive in all except the simplest possible games. An importance sampling algorithm, related to those suggested by Akerberg and Gowrisankaran (2001) and Keane and Wolpin (1997; 2001), avoids recomputing all equilibria for each trial parameter value and reduces the computational burden of estimation. It can also be easily implemented as a parallel process. In a private information game, either static or dynamic, or in a single agent dynamic discrete choice model, the importance sampler works best in the case of a full random coefficient specification for all parameters. In this case, the set of equilibria or the value functions do not need to be recomputed when the simulated likelihood or simulated moment is calculated for the hyperparameters of the random coefficient distribution. Only the density of the draws of the random coefficients under different values of the hyperparameters needs to be recomputed each time. Otherwise, the recomputation of equilibrium or value function typically is needed for the non-random coefficients in each round of their iterations. The importance sampler in the static complete information game in Bajari, Hong, and Ryan (2004), however, does not require a random coefficient specification for the parameters in order to avoid recomputing the equilibria for each trial parameter values. This is because in this model, the equilibria are determined once the utilities for each player for all the action profiles are known. The mean utility function parameters only changes the latent distribution of the vector of utility profiles, and the parameters in the equilibrium selection mechanism only shifts the weights on the different equilibria.

Gallant, Hong, and Khwaja (2008) computationally experiments with a dynamic version of the complete information discrete game, using MCMC to explore the parameter space and a sequential importance sampler to integrate out the serially correlated unobserved state variables, which are the production costs of pharmaceutical firms in their application. Many theoretical properties of this model still have to be studied. They bystep the issue with multiple equilibrium by imposing an selection rule of using the equilibrium with the lowest total cost among the entering firms. The most attractive feature of the importance sampler for the static complete information game does not seem to carry over to the dynamic extension. Similar to incomplete information games, this is because even conditional an draw of all latent variables in the model, the value function and therefore the equilibrium prediction that relates the unobserved cost variables to entry profiles still depend on the parameters. Unless a random coefficient specification is imposed on all parameters and the hyperparameters of the random coefficient distribution are the goal of the estimation procedure, the value function and equilibrium still need to be recomputed for the

nonrandom coefficient parameters. The advantage of the importance sampler for avoiding equilibrium recomputation may be visible if an importance sampler is also used to compute the posterior distribution and related point estimates of the parameters. However, in general, MCMC appears to do a better job than an importance sampler for exploring the parameter space.

9 Conclusion

Estimating models that are consistent with Nash equilibrium behavior has many important empirical applications. They pose exciting challenges for topics ranging from nonparametric identification, parametric and semiparametric estimation and computational algorithms for iterating value functions and finding static and dynamic equilibria. Attempts to meet these challenges have generated an extensive and burgeoning literature in both empirical industrial organization and in econometrics, and in other empirical fields of economics. A brief survey as the one given here is necessary incomplete. We only hope to at least give the readers a very broad idea of the structure of the literature, which may be helpful in identifying the remaining key issues and in bringing the advances in estimation methodology into applications.

References

- ACKERBERG, D., X. CHEN, AND J. HAHN (2009): “A Practical Asymptotic Variance Estimator for Two-Step Semiparametric Estimators,” unpublished manuscript, UCLA and Yale University.
- ACKERBERG, D., AND G. GOWRISANKARAN (2001): “A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation,” Working paper, UCLA.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the nested fixed point algorithm: a class of estimators for discrete Markov decision models,” *Econometrica*, 70(4), 1519–1543.
- (2007): “Sequential estimation of dynamic discrete games,” *Econometrica*, 75(1), 1.
- AI, C., AND X. CHEN (2003): “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, 71(6), 1795–1843.
- ARADILLAS-LOPEZ, A. (2005): “Semiparametric estimation of a simultaneous game with incomplete information,” manuscript, University of Wisconsin at Madison, forthcoming, *Journal of Econometrics*.
- ARCIDIACONO, P., AND R. MILLER (2006): “CCP estimation of dynamic discrete choice models with unobserved heterogeneity,” *Manuscript, Duke University*.
- BAJARI, P., C. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75(5), 1331–1370.

- BAJARI, P., V. CHERNOZHUKOV, H. HONG, AND D. NEKIPELOV (2009): “Nonparametric and semi-parametric analysis of a dynamic game model,” .
- BAJARI, P., H. HONG, J. KRAINER, AND D. NEKIPELOV (2006): “Estimating Static Models of Strategic Interactions,” working paper, forthcoming, *Journal of Business and Economic Statistics*.
- BAJARI, P., H. HONG, J. KRAINER, AND D. NEKIPELOV (2009): “Computing Equilibria in Static Games of Incomplete Information Using the All-Solution Homotopy,” .
- BAJARI, P., H. HONG, AND S. RYAN (2004): “Identification and estimation of discrete games of complete information,” *NBER Working paper*, forthcoming, *Econometrica*.
- BAYER, P., P. ARCIDIACONO, J. BLEVINS, AND P. ELLICKSON (2010): “Estimation of Dynamic Discrete Choice Models in Continuous Time,” *Economic Research*.
- BENKARD, C. (2004): “A dynamic analysis of the market for wide-bodied commercial aircraft,” *Review of Economic Studies*, 71(3), 581–611.
- BERRY, S. (1992): “Estimation of a model of entry in the airline industry,” *Econometrica*, 60(4), 889–917.
- BERRY, S., A. PAKES, AND M. OSTROVSKY (summer 2007): “Simple estimators for the parameters of dynamic games (with entry/exit examples),” *Rand Journal of Economics*, 38(2), 373–399.
- BERRY, S., AND E. TAMER (2006): “Identification in Models of Oligopoly Entry,” in *Advances in Economics and Econometrics, Vol 2, Ninth World Congress, 46-85*, ed. by Blundell, Newey, and Persson. Cambridge University Press.
- BJORN, P. A., AND Q. VUONG (1984): “Simultaneous Equations Models for Dummy Endogenous Variables: A Game Theoretic Formulation with an Application to Labor Force Participation,” SSWP No. 537, Caltech.
- BLEVINS, J. (2008): “Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models,” Duke University Working Paper.
- BORKOVSKY, R., U. DORASZELSKI, AND Y. KRYUKOV (2009): “A User’s Guide to Solving Dynamic Stochastic Games Using the Homotopy Method,” .
- BRESNAHAN, T., AND P. REISS (1990): “Entry in Monopoly Markets,” *Review of Economic Studies*, 57, 531–553.
- (1991): “Empirical Models of Discrete Games,” *Journal of Econometrics*, 48, 57–81.
- BROCK, W., AND S. DURLAUF (2001): “Discrete Choice with Social Interactions,” *Review of Economic Studies*, 68, 235–260.

- CHEN, X., AND D. POUZO (2009): “Efficient estimation of semiparametric conditional moment models with possibly nonsmooth residuals,” *Journal of Econometrics*, 152(1), 46–60.
- CHERNOZHUKOV, V., AND H. HONG (2003): “A MCMC Approach to Classical Estimation,” *Journal of Econometrics*, 115(2), 293–346.
- CILIBERTO, F., AND E. TAMER (2007): “Market Structure and Multiple Equilibria in Airline Markets,” University of Virginia and Northwestern University, forthcoming, *Econometrica*.
- DORASZELSKI, U., AND K. JUDD (2007): “Dynamic stochastic games with sequential state-to-state transitions,” Discussion paper, Citeseer.
- DORASZELSKI, U., K. JUDD, AND L. CENTER (2008): “Avoiding the curse of dimensionality in dynamic stochastic games,” .
- DUBE, J., J. FOX, AND C. SU (2009): “Improving the Numerical Performance of BLP Static and Dynamic Discrete Choice Random Coefficients Demand Estimation,” working paper, the University of Chicago.
- ERICSON, R., AND A. PAKES (1995): “Markov-perfect industry dynamics: A framework for empirical work,” *The Review of Economic Studies*, 62(1), 53–82.
- FAN, J., AND I. GIJBELS (1996): *Local polynomial modelling and its applications*. Chapman & Hall/CRC.
- FERNÁNDEZ-VILLAYERDE, J., AND J. RUBIO-RAMÍREZ (2007): “Estimating macroeconomic models: a likelihood approach,” *Review of Economic Studies*, 74(4), 1059–1087.
- GALLANT, A., H. HONG, AND A. KHWAJA (2008): “Estimating a dynamic oligopolistic game with serially correlated unobserved production costs,” Discussion paper, Working paper, Fuqua School of Business, Duke University.
- HECKMAN, J., AND B. SINGER (1984): “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data,” *Econometrica*.
- HOLMES, T. (2008): “The diffusion of Wal-Mart and economies of density,” *NBER Working Paper*, forthcoming *Econometrica*.
- HONG, H., A. MAHAJAN, AND D. NEKIPELOV (2009): “Consistent numerical asymptotic variation estimation,” UC Berkeley and Stanford University.
- HOTZ, J., AND R. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.

- HU, Y., AND M. SHUM (2008): “Nonparametric identification of dynamic models with unobserved state variables,” *Jonhs Hopkins University, Dept. of Economics working paper*, 543, 2009.
- IMAI, S., N. JAIN, AND A. CHING (2009): “Bayesian estimation of dynamic discrete choice models,” *Econometrica*, 77(6), 1865–1899.
- ISHII, J. (2005): “Compatibility, competition, and investment in network industries: ATM networks in the banking industry,” *Unpublished working paper*.
- 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.
- JIA, P. (2008a): “What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry,” *Econometrica*, 76(6), 1263–1316.
- (2008b): “What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry,” *Econometrica*, 76(6), 1263–1316.
- JOFRE-BONET, M., AND M. PESENDORFER (2003): “Estimation of a dynamic auction game,” .
- JUDD, K., AND C. SU (2006): “A New Optimization Approach to Maximum Likelihood Estimation of Structural Models,” *Computing in Economics and Finance 2006*.
- KASAHARA, H., AND K. SHIMOTSU (2008): “Nonparametric identification of finite mixture models of dynamic discrete choices,” *Econometrica*, 77(1), 135–176.
- KEANE, M., AND K. WOLPIN (1997): “The Career Decisions of Young Men,” *The Journal of Political Economy*, 105(3), 473–522.
- (2001): “The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment,” *International Economic Review*, 42(4), 1051–1103.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying dynamic discrete decision processes,” *Econometrica*, 70(2), 801–816.
- MANSKI, C. (1993): “Identification of Endogenous Social Effects: The Reflection Problem,” *Review of Economic Studies*, 60(3), 531–42.
- MCKELVEY, R., AND T. PALFREY (1995): “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 10(1), 6–38.
- MCKELVY, R., AND A. MCLENNAN (1996): “The Computation of Equilibrium in Finite Games,” in *The Handbook of Computational Economics, Vol I*, ed. by H. Amman, D. A. Kendrick, and J. Rust, pp. 87–142. Elsevier.

- MCLENNAN, A. (2005): “The Expected Number of Nash Equilibria of a Normal Form Game,” *Econometrica*, 73(1), 141–174.
- NAVARRO, S., AND Y. TAKAHASHI (2010a): “A semiparametric test of agent’s information sets for games of incomplete information,” Working paper, the University of Wisconsin at Madison.
- (2010b): “A semiparametric test of agent’s information sets for games of incomplete information,” Working paper, Department of Economics, University of Wisconsin-Madison.
- NEWHEY, W., AND J. POWELL (2003): “Instrumental variable estimation of nonparametric models,” *Econometrica*, 71(5), 1565–1578.
- NORETS, A. (2009): “Inference in Dynamic Discrete Choice Models With Serially correlated Unobserved State Variables,” *Econometrica*, 77(5), 1665–1682.
- PAKES, A., J. PORTER, K. HO, AND J. ISHII (2005): “Moment Inequalities and Their Application,” working paper, Harvard University.
- PANTANO, J., AND Y. ZHENG (2009): “Using Subjective Expectations Data to Allow for Unobserved Heterogeneity in Hotz-Miller Estimation Strategies,” working paper, Washington University in St. Louis.
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2003): “Identification and Estimation of Dynamic Games,” NBER working paper No. w9726.
- PESENDORFER, M., P. SCHMIDT-DENGLER, AND H. STREET (2008): “Asymptotic least squares estimators for dynamic games,” *Review of Economic Studies*, 75(3), 901–928.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- (1994): “Structural Estimation of Markov Decision Processes,” in *Handbook of Econometrics*, Vol. 4, ed. by R. Engle, and D. McFadden, pp. 3082–146. North Holland.
- (1997): “Using randomization to break the curse of dimensionality,” *Econometrica: Journal of the Econometric Society*, 65(3), 487–516.
- SEIM, K. (2006): “An empirical model of firm entry with endogenous product-type choices,” *The RAND Journal of Economics*, 37(3), 619–640.
- SWEETING, A. (2009): “The strategic timing incentives of commercial radio stations: An empirical analysis using multiple equilibria,” *The RAND Journal of Economics*, 40(4), 710–742.

- TABER, C. (2000): “Semiparametric identification and heterogeneity in discrete choice dynamic programming models,” *Journal of Econometrics*, 96(2), 201–229.
- TAMER, E. (2003): “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria,” *Review of Economic Studies*, 70(1), 147–167.
- TIOLE, J. (1988): *The Theory of Industrial Organization*. MIT Press.
- WANG, Y. (2009): “Subjective Expectations: Test for Bias and Implications for Choices,” PhD dissertation paper, Duke University and Lafayette College.
- WEINTRAUB, G., C. BENKARD, AND B. VAN ROY (2008): “Markov perfect industry dynamics with many firms,” *Econometrica*, 76(6), 1375–1411.