NBER WORKING PAPER SERIES

ESTIMATING DYNAMIC MODELS OF IMPERFECT COMPETITION

Patrick Bajari
C. Lanier Benkard
Jonathan Levin

Working Paper 10450
http://www.nber.org/papers/w10450

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2004

We have had very helpful conversations with Victor Aguirregabiria, Jeremy Fox, Phil Haile, Igal Hendel, Guido Imbens, Phillip Leslie, Ariel Pakes, Peter Reiss, Azeem Shaikh, Elie Tamer, and Ed Vytlacil as well as seminar participants at University of Minnesota, University of Wisconsin, and the SITE Applied Microeconomics workshop. We are also indebted to Matthew Osborne for providing exemplary research assistance. Bajari and Benkard thank the NSF and the Bureau of Economic Analysis for supporting this research. The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

Estimating Dynamic Models of Imperfect Competition
Patrick Bajari, C. Lanier Benkard, and Jonathan Levin
NBER Working Paper No. 10450
April 2004
JEL No. L0, C5

ABSTRACT

We describe a two-step algorithm for estimating dynamic games under the assumption that behavior is consistent with Markov Perfect Equilibrium. In the first step, the policy functions and the law of motion for the state variables are estimated. In the second step, the remaining structural parameters are estimated using the optimality conditions for equilibrium. The second step estimator is a simple simulated minimum distance estimator. The algorithm applies to a broad class of models, including I.O. models with both discrete and continuous controls such as the Ericson and Pakes (1995) model. We test the algorithm on a class of dynamic discrete choice models with normally distributed errors, and a class of dynamic oligopoly models similar to that of Pakes and McGuire (1994).

Patrick Bajari
Department of Economics
Duke University
Durham, NC 27708
and NBER
bajari@econ.duke.edu

C. Lanier Benkard
Graduate School of Business
Stanford University
Stanford, CA 94305-5015
and NBER
lanierb@stanford.edu

Jonathan Levin
Department of Economics
Stanford University
Stanford, CA 94305
jdlevin@stanford.edu

1 Introduction

In many branches of applied economics, it is now common practice to estimate structural models of decision making and equilibrium. In most cases, however, attention has been focused on static environments. Estimating dynamic parameters has been seen as substantially more difficult, both conceptually and computationally. To the extent that previous work has succeeded in estimating dynamic models, it has, with a few notable exceptions described below, been limited to single-agent problems. Of course, many of the parameters at the heart of policy debates are inherently dynamic, such as entry and exit costs, the returns to advertising or R&D, or adjustment costs of investment. Dynamics are also of particular interest in industrial organization. For instance, models incorporating network effects, durable goods, experience goods, consumer learning, or firm learning-by-doing are inherently dynamic.

One reason that empirical work on dynamic competition has been limited is the perceived difficulty of incorporating information from a dynamic equilibrium into an estimation algorithm. The literature on dynamic oligopoly, including Ericson and Pakes (1995), Pakes and McGuire (1994, 2001), Gowrisankaran and Town (1997), Benkard (forthcoming), has shown that computing an equilibrium for even the simplest of industries is all but prohibitive. For models with the complexity usually required for empirical work, the situation is even more bleak. Even with advancing computer technology, computing equilibria over and over, as would be required in a typical estimation routine, seems out of the question. Moreover, dynamic games typically admit a vast multiplicity of equilibria. This multiplicity greatly complicates the application of estimators that require computing equilibria and then matching these equilibria to observed data.

This paper develops a method for estimating dynamic models of imperfect competition that is straightforward to apply, without requiring the ability to compute an equilibrium even once. The approach involves two steps. The first is to recover the agents' policy functions, as well as the probability distributions determining the evolution of the observed and unobserved state variables. In practice, this would typically involve running a nonparametric regression of observed actions (such as investment, quantity, price, entry or exit) on the observed state variables (such as aggregate demand/cost shifters, and firm and product characteristics). A feature of equilibrium models is that agents understand and have correct beliefs about their environment and the behavior of other agents. As a consequence, by recovering the probability distributions for actions and states, one is effectively recovering the agents' beliefs at each point in time.

The second step of the approach involves finding the set of parameters that makes the observed policies optimal. This parallels the second feature of equilibrium models, namely

that agents maximize expected discounted profits given their beliefs. These optimization conditions can be represented as a system of inequalities requiring that each agent's observed choice at each state be weakly preferred to all feasible alternatives. The structural parameters are estimated as the solution to this system of inequalities. In practice, this involves applying a simple simulated minimum distance estimator that minimizes violations of the optimality conditions.

Our approach relates most closely to that of Hotz and Miller (1993), and Hotz et al. (1993). Hotz and Miller (1993) showed, in a single agent dynamic discrete choice problem, that knowledge of the probability distribution over an agent's choices is sufficient to derive the agent's value function. They used this result to develop a two stage estimation approach in which the choice probabilities were estimated in the first stage, and then the structural parameters estimated in the second stage. Several recent papers similar to our own have greatly extended these ideas in the contexts of dynamic discrete games (Aguirregabiria and Mira (2002a)), entry games (Pakes, Ostrovsky, and Berry (2003), Pesendorfer and Schmidt-Dengler (2003)), and dynamic auction games (Jofre-Benet and Pesendorfer (forthcoming)).

This paper makes several contributions to this literature. The first is that the estimation algorithm is applicable to a wider class of models than previous methods. A limitation of the existing literature is that it requires discrete actions (such as entry or exit) with each alternative subject to an idiosyncratic profit/cost shock. This rules out many important dynamic problems in I.O. such as investment games (e.g., Ericson and Pakes (1995)), dynamic pricing games (durable goods, network effects, learning by doing, et al.), as well as entry games where data is also available on investment or prices. One of the primary contributions of this paper is that the algorithm described applies equally to all of these cases. For example, in section 5.1 we apply the estimator to a single agent dynamic discrete choice problem similar to that of Rust (1987) with normally distributed errors. In section 5.2 we apply the estimator to a version of the Pakes and McGuire (1994) model that has both discrete (entry and exit) and continuous (investment) controls. To our knowledge, none of the estimation algorithms in the previous literature can be applied to this second example.[1]

The other main contributions of the paper are computational. A key point that greatly reduces the computational burden of the algorithm is that many dynamic games are linear in the parameters of interest. The argument is as follows: assuming agents maximize the expected discounted value of future profits, if the period return function is linear in the parameters, then the value function can also be written as linear in the parameters.

[1]Note that Berry and Pakes (2000) also address the issue of continuous controls by using an euler-equation based estimator. However, it is not currently known how to incorporate information from discrete controls into this estimator.

This linearity can be exploited to achieve a substantial computational savings because it means that the parameters of interest essentially factor out of the value function. Thus, the simulated expected discounted value terms in the value function need only be computed once and then held fixed for all parameter values, rather than recomputed at every parameter value as would be typical. We also obtain computational savings over the previous literature by making heavier use of simulation.

One benefit of the two-stage approach in the context of dynamic games is that it somewhat mitigates the multiple equilibrium problem. Because equilibrium beliefs are recovered from the observed data, the researcher need not take a stand on which of many potential equilibria is being played. Instead, this determination is made by the observed data.

The algorithm also extends easily to models that are not point identified. In many models involving discrete controls, such as entry models, some parameters may not be point identifiable from the observed policy functions. Instead, even with an infinite amount of data, it might only be possible to place the parameters within some set. Examples can be found in Bresnahan and Reiss (1991), Ciliberto and Tamer (2003), Haile and Tamer (2003), and Pesendorfer and Schmidt-Dengler (2003). In other cases, an identification proof may simply be difficult to obtain. The estimation algorithm can be applied to these cases with little alteration while being agnostic on the issue of identification. In that case it produces bounds on the parameters of interest.

The primary cost to employing a two-stage approach is that the approach does not utilize all of the information available from the structural model. Some efficiency is lost because the first stage policy function estimates are typically performed nonparametrically, without being informed by the structure of the model. Aguirregabiria and Mira (2002b) show in the single agent context that two-step estimators may perform poorly relative to estimators that incorporate more information, for example, by iterating the Bellman equation. In many contexts that interest us, this inefficiency is inescapable because iterating the Bellman equation mapping even once would often be computationally infeasible. Furthermore, in the context of a dynamic game, the Bellman equation is not a contraction mapping so there is no guarantee that iterating it a small number of times would improve the efficiency of the estimates. Given this efficiency loss, one question of interest is whether the two stage approach is efficient enough to provide good estimates with reasonably sized data sets.

To answer this question, we evaluate the efficiency and computational burden of the estimators using monte carlo experiments on the two examples mentioned above. We find that the algorithm has very low computational burden. We also find that the algorithm works surprisingly well even for relatively small data sets, containing fewer observations

than one might reasonably find in real world applications. For example, in the dynamic oligopoly example, we find that with a moderately sized data set it is possible to recover the entry cost distribution nonparametrically (ignoring all prior information about its parametric form) with acceptable precision.

The paper proceeds as follows. The next section introduces the class of models that we are interested in, and provides two specific examples. Sections 3 and 4 outline the estimation algorithm and provide the relevant asymptotic theory. Section 5 details how the algorithm applies to the two examples and provides monte carlo evidence on the performance of the estimator. Section 6 concludes the paper.

2 The Model

This section outlines a model of dynamic competition between oligopolistic competitors that encompasses many applications in industrial organization. Examples that fit into the general framework include entry and exit decisions, dynamic pricing, and investments in capital stock, advertising, or research and development. The defining feature of the model is that actions taken in a given period affect future payoffs, and future strategic interaction, by influencing a set of commonly observed state variables. The model also includes single agent dynamic decision problems as a special case; the single agent model is just a game with a single player. We solve the model using the concept of Markov Perfect Equilibrium.

The model consists of N firms, denoted by $i = 1, ..., N$, who make decisions at times $t = 0, 1, 2, ..., \infty$. At a given time, t, the prevailing conditions are summarized by a vector of state variables, $\mathbf{s}_t \in S \subset \mathbb{R}^G$. Relevant state variables might include the firms' production capacities, their technological progress up to time t, the current market shares, stocks of consumer loyalty, or simply the set of firms that are incumbent in the market. We assume that these state variables are commonly observed by the firms.

Given the state $\mathbf{s}_t$ at date t, the firms simultaneously choose actions. Depending on the application, the firms' actions could include decisions about whether to enter or exit the market, investment or advertising levels, or choices about prices and quantities. We denote firm i's action at date t as $a_{it} \in A_i$. We also assume that before choosing its action each firm, i, observes a *private* shock $\nu_{it} \in \mathbb{R}$, drawn independently from a distribution $G(\cdot|\mathbf{s}_t)$.[2] Private information might derive from variability in marginal costs

[2]We assume firm i's private shock is a single scalar variable, as would be natural if we wanted to ultimately estimate i's decision rule with a logit or probit model. There is no conceptual difficulty in allowing the shock to be multi-dimensional, which would be convenient if, for instance, we ultimately

of production, or the need for plant maintenance, or from variability in sunk costs of entry or exit. We denote the vector of private shocks as $\nu_t = (\nu_{1t}, ..., \nu_{Nt})$.

In each period, each firm earns profits equal to $\pi_i(\mathbf{a}_t, \mathbf{s}_t, \nu_{it})$. Each firm, i, is interested in maximizing its discounted sum of profits:

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \pi_i(\mathbf{a}_t, \mathbf{s}_t, \nu_{it}). \tag{1}$$

We assume firms have a common discount factor β.

The final aspect of the model is to specify the transitions between states. We assume that the state at date $t+1$, denoted $\mathbf{s}_{t+1}$, is drawn from a probability distribution $P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$. The dependence of $P(\cdot|\mathbf{s}_t, \mathbf{a}_t)$ on the current period actions $\mathbf{a}_t$ reflects the fact that some time t decisions may affect future payoffs, as is clearly the case if the relevant decision being modeled is an entry/exit decision or a long-term investment. Of course, not all the state variables necessarily depend on past actions; for instance, one component of the state could be a transitory *iid* shock that affects only the current payoffs — for instance, *iid* shocks to market demand.

We are interested in equilibrium behavior for the model we have just outlined. Because the firms interact repeatedly, and the horizon is infinite, there are likely to be very many Nash, and even subgame perfect equilibria, possibly involving complex behavioral rules. We focus instead on pure strategy Markov Perfect Equilibria (MPE).

A Markov strategy for firm i describes the firm's behavior at time t as a function of the commonly observed state variables and firm i's private information at time t. Formally, it is a map, $\sigma_i : S \times \mathbb{R} \to A_i$. A profile of Markov strategies is a vector, $\sigma = (\sigma_1, ..., \sigma_n)$, where $\sigma : S \times \mathbb{R}^n \to A$.

If behavior is given by a Markov profile σ, firm i's present discounted profits from the start of a period can be written in recursive form:

$$V_i(\mathbf{s}|\sigma) = \mathbb{E}_\nu \left[\pi_i(\sigma(\mathbf{s}, \nu), \mathbf{s}_t, \nu_i) + \beta \int V_i(\mathbf{s}'|\sigma(\mathbf{s}, \nu)) dP(\mathbf{s}'|\sigma(\mathbf{s}, \nu), \mathbf{s}) \right].$$

A Markov strategy profile, σ, is a Markov Perfect Equilibrium if there is no firm, i, and alternative Markov strategy, σ_i', such that firm i prefers the strategy σ_i' to the strategy σ_i given its opponents use the strategy profile σ_{-i}. That is, σ is a MPE if for all firms, i, all states, $\mathbf{s}$, and all Markov strategies, σ_i':

$$V_i(\mathbf{s}|\sigma) \geq V_i(\mathbf{s}|\sigma_i', \sigma_{-i}).$$

wanted to estimate i's decision rule with a multinomial logit model.

A useful simplification to verify that each firm's strategy is indeed a best-response to the strategies of the other firms is to make use of Bellman's principle of optimality. By principle of optimality, to check that σ_i is an optimal strategy given opponent strategies, σ_{-i}, it suffices to check that, for all pairs, ($\mathbf{s}$,ν_i), the action $\sigma_i(\mathbf{s},\nu_i)$ leads to higher expected discounted profits than any other action $a_i \in A_i$ given that firm i will continue to use its strategy σ_i for all future decisions. This simplification will prove useful in the estimation procedure below.

We will make extensive use of the set of inequalities defining equilibrium in the estimation procedure. The key idea is that if one knows β, and is able to recover $\sigma(\mathbf{s},\nu)$ and $P(\mathbf{s}'|\mathbf{a},\mathbf{s})$ through some first-stage estimation, we can use the equilibrium inequalities to recover at least some information, and perhaps complete information, about the unknown per-period profit function $\pi_i(\mathbf{a},\mathbf{s},\nu_i)$.

Before turning to the econometric methods, however, we first introduce two simple examples to illustrate the model. We return to these examples in Section 5 where we test the estimation procedure using simulated data.

2.1 Example 1: Dynamic Discrete Choice

Perhaps the simplest example of the model is dynamic discrete choice problems.[3] In dynamic discrete choice problems, there is a single agent who chooses a single action, a, out of a finite set of actions, A. The state includes variables, s, that are observed by the economist, and a vector, ν, of unobserved state variables. The period profits that agent i receives from choice j is

$$\pi(a_j, s, \nu) = \tilde{\pi}(a_j, s; \theta) + \nu(a_j), \tag{2}$$

where $\tilde{\pi}$ is a known (up to θ) parametric function that depends on the choice and the vector of state variables, and $\nu(a_j)$ is a stochastic shock to preferences.

For example, consider the machine replacement problem in Rust (1987). In Rust (1987) there is a single manager who owns a machine that operates in every period. The state variable is the machine's age $s_t \in \{1, ..., M\}$. At each date the manager chooses whether to replace the machine ($a_t = 1$) or maintain it ($a_t = 0$), in which case the machine ages by one period until it reaches age M, after which it remains M years old until it is replaced.

[3]Dynamic discrete choice has received considerable attention in the past literature, including Rust (1987), Hotz and Miller (1993), Hotz et al. (1993), and Keane and Wolpin (1997), and Aguirregabiria and Mira (2002b).

The manager's per-period payoff is:

$$\pi(a_t, s_t, \nu_t) = \begin{cases} -\theta s_t - \nu_t & \text{if } a_t = 0 \\ -R & \text{if } a_t = 1 \end{cases},$$

where $\theta s_t + \nu_t$ represents the (random) cost of maintaining the machine and the parameter R is the cost of machine replacement. In this example, a Markov policy is a function, $\sigma(s, \nu)$, specifying whether the machine should be maintained or replaced. It is straightforward to show that an optimal policy takes the form of a cut-off rule: $\sigma(s, \nu) = 1$ if $\nu \geq \nu^*(s)$, where $\nu^*(\cdot)$ is decreasing in the machine age s, consistent with the idea that older machines are more likely to be replaced.

2.2 Example 2: Dynamic Oligopoly

The next example represents a class of dynamic oligopoly problems similar to that of Ericson and Pakes (1995) and Pakes and McGuire (1994).

In this class of models, there are a set of incumbent firms competing in a market. Firms are heterogeneous, with differences across firms described by their state variables, s_{it}. Each period, firms choose their investment levels, $I_{it} \geq 0$, so as to improve their state the next period. Investment outcomes are random, and a firm's own investment is assumed not to influence the investment outcomes of other firms.

Examples of models that are consistent with this framework include:

(i) Firms' state variables could represent product quality, where investment stochastically improves product quality. In this example, investment could either be technological or it could be in the form of advertising.

(ii) Firms' state variables could represent capital stock, where investment stochastically increases a firm's capital stock.

Firms earn profits by competing in a spot market. Because quantity and price are assumed not to influence the evolution of the state variables, they are determined in static equilibrium conditional on the current state. The period return function is given by

$$\pi_i(\mathbf{a}_t, \mathbf{s}_t, \nu_{it}; \theta) = q_{it}(\mathbf{s}_t, \mathbf{p}_t; \theta_1) \left(p_{it} - \text{mc}(s_{it}, q_{it}; \theta_2)\right) - C(I_{it}, \nu_{it}; \theta_3), \tag{3}$$

where q_{it} is quantity produced by firm i in period t, $\mathbf{p}_t$ is the vector of prices, mc is the marginal cost of production, ν_{it} represents a private shock to the cost of investment,

$\theta = (\theta_1, \theta_2, \theta_3)$ is a parameter vector to be estimated, and we have assumed that the spot market equilibrium is Nash in prices.[4]

The model also allows for entry and exit. Each period, each incumbent firm has the option of exiting the market and receiving a scrap value, Ψ, which is the same for all firms.[5] There is also one potential entrant each period with entry cost, x^e. The entrant enters if the expected discounted value of entering exceeds the entry cost. The entry cost is private information. However, the distribution of entry costs, $F(x^e)$, is commonly known.

A detailed example follows in section 5.2.3.

3 Estimation: Overview

3.1 First Step

The estimation approach has two steps. The first step is to estimate the policy functions, $\sigma_i : S \times \mathbb{R} \to A_i$, and the state transition probabilities, $P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$, using past observations of play in the game. In practice, the best methods for doing this depend on the details of the application, making a general treatment difficult. We therefore reserve a detailed discussion of the first step estimation to a discussion of the two examples in section 5, and provide only an overview here.

The first step estimation would typically involve a variety of econometric techniques. Because the state transition probabilities are model primitives, they would typically be parameterized and estimated using parametric methods such as maximum likelihood. The policy functions, on the other hand, result from equilibrium play, and would typically be estimated nonparametrically, even if the parametric form of the profit function were known. For example, in the dynamic oligopoly model presented above, the policy functions could be estimated by running nonparametric regressions of investment, entry, and exit on the state variables. In many cases it would also be possible to estimate some of the profit function parameters in the first stage using techniques such as demand estimation or production function estimation. For example, in the dynamic oligopoly model, the demand and cost functions would typically be estimated this way.

Note, however, that the ability to estimate the first stage is by no means guaranteed. In

[4] It would not complicate the model to assume instead Nash in quantities.

[5] It would not complicate the model to have the exit cost be random.

practice, there are three main difficulties that might arise. Consider first the simple case where all the variables, $(\mathbf{a}, \mathbf{s})$, are observed and equilibrium is unique. Even under these strict conditions it may not be possible to completely recover the first stage functions because not all states, $\mathbf{s} \in S$, may be visited infinitely often under the Markov chain of states, $P(\mathbf{s}_{t+1}|\mathbf{s}_t)$, that is generated by equilibrium play. At best, the policy functions can only be recovered for states visited recurrently on the equilibrium path, or on the recurrent class of states, $R \subseteq S$. As a simple example of this, in an industry such that equilibrium play always leads to at least two incumbent firms, it is impossible to use the data to learn nonparametrically how a firm would behave if it were to be an incumbent monopolist. We will revisit this issue further below because it can affect the identifiability of certain model parameters.

Next, suppose that all variables are observed but there are multiple equilibria to the game. In that case, so long as the data comes from only one market observed over time, the existence of multiple equilibria to the game does not affect the ability to estimate the first stage. The reason is that, assuming $\mathbf{s}$ is a complete description of the state variables of the model, agents can only form beliefs based on $\mathbf{s}$. In equilibrium, agents' beliefs must be correct. Therefore, by standard results in nonparametrics, their beliefs must be recoverable from observed play. The time when multiple equilibria might interfere with estimation of the first stage would be when pooling together data for agents in disconnected markets, because it might be that different equilibria were being played in each market. In that case the policy functions would be different in each market, so pooling the data across markets would lead to incorrect inferences.

Probably the biggest hurdle to the first stage estimation is the observability of the state variables, $\mathbf{s}$. Estimation of the first stage does not require that all states be observed. In many cases unobserved states are recoverable in the first stage estimation. Leading examples of this include unobserved product characteristics in demand models (e.g. Berry (1994), Berry, Levinsohn, and Pakes (1995)), and unobserved productivity shocks in production functions (e.g. Olley and Pakes (1996)). In cases like these, the values of the unobserved states are uniquely determined from the observed states and actions, and so it is as if the unobserved states were observed. Once the unobserved states are recovered, the first stage policy functions can be estimated as before and even serial correlation in the unobserved states is easily handled. In other examples, such as dynamic discrete choice (e.g. Rust (1987), Hotz and Miller (1993), Hotz et al. (1993)) or dynamic discrete games (e.g. Aguirregabiria and Mira (2002a)), a parametric form is assumed for the distribution of the unobserved states and therefore, even though the values of the unobserved states are not uniquely determined by the observed variables, the first stage is still easily recovered.

Still, the ability to estimate the first stage clearly places strong restrictions on the extent

to which it is possible to allow for unobserved states, as well as the extent to which it is possible to handle serially correlated unobserved states. However, characterizing these restrictions is difficult without the context of a detailed model so we instead discuss this issue further in the examples in section 5.

3.2 Second Step

The first step estimation typically recovers many, but not all, of the parameters of interest. It is necessary to impose the dynamic equilibrium conditions in order to recover parameters such as investment costs, adjustment costs, fixed costs, entry and exit costs, etc. This is the purpose of the second step estimation.

For notational simplicity, we assume that the set of actions as well as the set of possible state variables are discrete. Our methods easily generalize to continuous problems (indeed, the dynamic oligopoly example involves continuous investment policies). We also assume that the discount factor, β, is known. Assuming it is identified, there is some scope for estimating the discount factor. However, the analysis is more straightforward with this assumption.

Consider the optimality conditions for the strategy, σ_i, generated by the objective function defined in (1). Optimality requires that for all i, all alternative policies σ_i', and all initial states $\mathbf{s}_0$:[6]

$$\mathbb{E}_{\sigma_i,\sigma_{-i}|\mathbf{s}_0}\sum_{t=0}^{\infty}\beta^t\pi_i(\mathbf{a}_t,\mathbf{s}_t,\nu_{it}) \geq \mathbb{E}_{\sigma_i',\sigma_{-i}|\mathbf{s}_0}\sum_{t=0}^{\infty}\beta^t\pi_i(\mathbf{a}_t,\mathbf{s}_t,\nu_{it}), \tag{4}$$

where the expectations operator $\mathbb{E}_{\sigma_i,\sigma_{-i}|\mathbf{s}_0}$ refers to the expectation with respect to equilibrium policies, starting from initial state $\mathbf{s}_0$, and the expectations operator $\mathbb{E}_{\sigma_i',\sigma_{-i}|\mathbf{s}_0}$ refers to the same expectation only where agent i's strategy is changed to σ_i' rather than σ_i.

The system of inequalities given by (4) represents the information contained in the model's equilibrium assumption. Because the policy functions and transition probabilities have been recovered at the first step, the only unknowns in (4) are the period profits. A key point is that this system of inequalities is linear in these profits.

Our main additional assumption is that the period profit functions are known up to some finite vector of parameters, $\theta = (\theta_1, ..., \theta_M)$, and moreover, that each profit function is

[6]We remind readers that in practice, by the principle of dynamic programs, it is only necessary to consider one-step deviations.

linear in these parameters (we discuss this assumption in detail in the section below):

$$\pi_i(\mathbf{a}, \mathbf{s}, \nu_i) := \Phi_i(\mathbf{a}, \mathbf{s}, \nu_i) \cdot \theta. \tag{5}$$

Here $\Phi_i(\mathbf{a}, \mathbf{s}, \nu_i)$ is a M-dimensional vector of "basis functions" in $\mathbf{a}$, $\mathbf{s}$, and ν_i, ϕ_1,...,ϕ_M. If $\pi_i(\mathbf{a}, \mathbf{s}, \nu_i)$ is linear, then the functions ϕ_j are simply the elements of $(\mathbf{a}, \mathbf{s}, \nu_i)$. More generally, the functions ϕ_j could be polynomial, or other nonlinear, basis functions in $(\mathbf{a}, \mathbf{s}, \nu_i)$.

Given this restriction, we can express the optimization inequalities (4), as a system that is linear in the unknown parameters, θ. To do this, define:

$$W(\mathbf{s}_0; \sigma_i, \sigma_{-i}) := \mathbb{E}_{\sigma_i, \sigma_{-i}|\mathbf{s}_0} \sum_{t=0}^{\infty} \beta^t \Phi_i(\mathbf{a}_t, \mathbf{s}_t, \nu_{it}).$$

Then the conditions for equilibrium are that for every i, initial state $\mathbf{s}_0$, and alternative policy σ_i':

$$W(\mathbf{s}_0; \sigma_i, \sigma_{-i}) \cdot \theta \geq W(\mathbf{s}_0; \sigma_i', \sigma_{-i}) \cdot \theta \tag{6}$$

The primary difference between this set of equilibrium conditions and (4) is that the parameter vector, θ, factors out linearly. This factorization results in a computational savings in estimation because, as we will see, it means that the W terms need only be computed once, rather than over and over for every value of θ. Computing the W terms is the primary computational burden of the estimator so this reduction is substantial.[7]

3.3 The Linearity Assumption

The primary assumption reducing the computational burden of the estimator is the linearity assumption in (5). First, note that the assumed linear representation in (5) is stronger than required because many of the parameters of the period return function can typically be estimated in the first step along with the policy functions. Thus, the model need not be linear in all the parameters, just those parameters that are to be estimated in the second step.

[7] If the state and action spaces are both discrete and there are a small number of points in the state space, then there is also a useful analytic matrix representation of the system (6): Let B be a matrix of ones and zeros that picks out, for each (s, ν_i), the optimal action for agent i, $a_i = \sigma_i(s, \nu_i)$, and $\tilde{B}$ be the same for the alternative policy σ'. Let the vector, $\Phi'\theta$, represent the period returns at each (s, a, ν_i) combination. Then the system (6) is,

$$\left[I - \delta \mathbb{E}_{\sigma_i, \sigma_{-i}}\right]^{-1} B\Phi'\theta \geq \left[I - \delta \mathbb{E}_{\sigma_i', \sigma_{-i}}\right]^{-1} \tilde{B}\Phi'\theta. \tag{7}$$

With this distinction, many dynamic oligopoly models in the literature satisfy the linearity assumption. For example, entry, exit, and fixed costs parameters enter additively into a firm's profit function and so are naturally linear. Thus, dynamic entry models typically satisfy this assumption. Investment cost and marginal cost parameters would also typically enter profits linearly. The dynamic oligopoly model is an example of a class of models that is linear in the second step parameters (though highly nonlinear in the first step parameters).

Of course, it is also easy to write down models that are not linear in parameters. For example, many dynamic pricing models would be nonlinear in the parameters. Additionally, if treated as a parameter to be estimated, the discount factor enters the system nonlinearly. The inequalities (4) must still hold in these cases, and the estimation algorithm described in the next section could still be used. However, the computational burden of the algorithm would be higher because the parameters no longer factor out as in (6).

In the event that the computational burden proved too high, an alternative approach for nonlinear models would be to use a flexible approximation to π_i that is linear in the second-step parameters. For example, a polynomial approximation could be used. Alternatively, an Mth order Taylor expansion of the desired nonlinear profit function could be used. Note that while we believe that such an approach is likely to work well, we have not attempted to explicitly derive the properties of an approximation estimator.

3.4 Identification

Given observed behavior, σ, and transitions, P, the equilibrium inequalities, (6), define a set of parameters that *rationalize* the data in the sense that the strategies, σ, are a Markov Perfect Equilibrium of the game defined by P and payoff parameters θ. Let Θ_0 be the set of parameters that rationalize the observed data,

$$\Theta_0(\sigma,P) := \{\theta : \theta, \sigma, P \text{ satisfy (6) for all } \mathbf{s}_0, i, \sigma_i'\}. \tag{8}$$

The goal of estimation is to learn this set. A natural way to proceed with estimation, then, would be to construct empirical counterparts to the inequalities, (6), and then find the value(s) of θ that minimize the squared distance in violations of these inequalities. This is the approach followed below.

The problem of identification concerns whether or not the set Θ_0 is a singleton. While identification can be shown for some dynamic decision problems (e.g. Rust (1994), Aguirregabiria and Mira (2002a,b), Pesendorfer and Schmidt-Dengler (2003)), other models

such as some entry models may not yield point identification (see also Pesendorfer and Schmidt-Dengler (2003), Ciliberto and Tamer (2003)). Even in non-identified models, however, knowledge of the set Θ_0 may convey useful information about the underlying parameters (Manski and Tamer (2002), Haile and Tamer (2003)). We therefore proceed by outlining two alternative approaches. The first estimator requires that the model be identified and yields standard point estimates as well as standard errors. The second estimator does not rely on identification and instead yields consistent estimates of bounds on the parameters. These approaches are discussed in detail in the next section.

Another issue with respect to identification is the fact that agents' beliefs can only be learned on the recurrent class of states, R. In practice, this means that it is only possible to implement the optimality inequalities for initial states that are in R. This limitation would affect the identifiability of profit parameters that are only realized outside of R. An example of parameters of this type would be parameters determining firm profits in a punishment regime that is never triggered by equilibrium play. Typically, knowledge of the fact that the punishment regime is never triggered by equilibrium play would only provide enough information to place bounds on such parameters.

4 Estimation and Asymptotic Theory

This section outlines an intuitive minimum distance (MD) approach designed to minimize the squared distance in violations of the optimization inequalities in (6).

4.1 Notation and Assumptions

We maintain the assumption that the policy function and the transition probabilities can be summarized by a parameter vector, α. That is, the policy function takes a known parametric form, $\sigma_i(\mathbf{s}, \nu_i; \alpha)$, where σ_i is twice continuously differentiable in α. We also assume that the transition probabilities have known parametric form, $P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t; \alpha)$, and that the transition probabilities are twice continuously differentiable in α when evaluated at the policy function $a_{it} = \sigma_i(\mathbf{s}_t, \nu_i; \alpha)$. We denote the true policy function parameters α_0. This assumption holds in any model with a discrete state space and also in continuous state space models where the parametric forms of the policy function and transition probabilities are known.[8]

[8] We also believe that the estimation algorithm could be shown to work if the assumption were relaxed to allow for a nonparametric first stage on a continuous state space, but we have not proved this.

Given an initial state, $\mathbf{s}_0$, the first stage estimates of the policy function and the transition probabilities provide all of the information needed to simulate equilibrium paths of states and controls. The expected discounted value terms, W, can be estimated by simulating many such paths, computing the discounted value of each basis function along each path, and then averaging. The simulated W terms provide the basis for the MD estimator.

Note that this "forward" simulation procedure is similar to those of Hotz et al. (1993) and Rust (1994) for single agent models. However, one difference is that Hotz et al. (1993) and Rust (1994) require that the simulations be performed for every possible action at every observed state. In our context there may be a large number of observed states and infinitely many alternative actions per agent per state. Therefore, such a requirement would be computationally prohibitive. We instead perform the simulations at only a small number of states and a small number of actions per state as follows.

Let x denote a particular $(i, \mathbf{s}_0, \sigma')$ combination, such that each value of x refers to one optimality inequality. For some value of α, let

$$g(x, \theta; \alpha) = \left[W(\mathbf{s}; \sigma_i(\alpha), \sigma_{-i}(\alpha)) - W(\mathbf{s}; \sigma_i', \sigma_{-i}(\alpha))\right] \cdot \theta,$$

and $1\{g(x, \theta; \alpha) < 0\}$ represent the indicator function for the event that $g(\cdot)$ is less than zero, so that the optimality inequality defined by x is not satisfied at θ.

4.2 Case 1: Estimation of Identified Models

The estimator is computed by sampling $k = 1, ..., n_I$ inequalities (defined by states and alternative actions), denoted X_k, according to some known distribution, F, that is chosen by the researcher. While this sampling distribution will have implications to the efficiency of the estimator, the only requirement for consistency is that the distribution used must have a support that yields identification of the model. That is, this distribution must place positive density on a set of optimality inequalities that makes Θ_0 a singleton.

The true parameter vector, θ_0, solves

$$\min_{\theta \in \Theta^*} Q(\theta; \alpha_0) \tag{9}$$

where Θ^* is a compact subset of $\mathbb{R}^M$ containing θ_0 and

$$Q(\theta; \alpha) \equiv \int 1\{g(X_k, \theta; \alpha) < 0\} g(X_k, \theta; \alpha)^2 dF(X_k).$$

Furthermore, by assumption, no $\theta \neq \theta_0$ solves this problem. Therefore, a natural way to estimate θ_0 would be to use a sample analog of (9).

Let $g_{n_s}(x, \theta; \alpha)$ be an unbiased and smooth simulator for $g(x, \theta; \alpha)$, where n_s is the number of simulation draws. Define the two-step minimum distance (MD) estimator as follows,

$$Q_n(\widehat{\theta}, \widehat{\alpha}_n) = \inf_{\theta \in \Theta^*} Q_n(\theta, \widehat{\alpha}_n). \tag{10}$$

where

$$Q_n(\theta, \alpha) = \frac{1}{n_I} \sum_{i=1}^{n_I} 1\{g_{n_s}(X_i, \theta; \alpha) < 0\} g_{n_s}(X_k, \theta; \alpha)^2.$$

Conditional on the first stage estimates, the MD estimator minimizes the squared distance of the violations to the simulated optimality inequalities.[9]

In addition to the assumptions above, assume that

(i) θ is known to lie in a compact set $\Theta^* \subset \mathbb{R}^M$.

(ii) $\widehat{\alpha}_n \xrightarrow{p} \alpha_0$ and $\sqrt{n}(\widehat{\alpha}_n - \alpha_0) \xrightarrow{d} N(0, V_\alpha)$.

(iii) $g_{n_s} \xrightarrow{p} g$ with probability 1 for every (θ, α). Additionally, $E|g_{n_s}(x, \theta; \alpha_0)| < \infty$ for all x and all $\theta \in \Theta^*$.

(iv) $\lim_{n\to\infty} \frac{n}{n_I} = r < \infty$ and $\lim_{n\to\infty} \frac{n}{n_s} = 0$.

The first assumption is standard and is largely technical. The second assumption says that the first stage estimator is consistent and asymptotically normal. This would typically be satisfied by any first stage estimator under consideration. The third assumption says that the simulator is consistent and has finite first moment. Again, this assumption would typically be satisfied by any simulator under consideration. The last assumption says that the number of inequalities (n_I) and the number of simulation draws (n_s) must increase at least as fast as the number of first stage observations. Since n_I and n_s are chosen by the researcher and since the estimator has a low computational burden, these assumptions are easily implemented.

Proposition 1. *Under these assumptions,*

$$\widehat{\theta} \xrightarrow{p} \theta_0$$

[9]The objective function in 10 puts equal weight on each inequality. Presumably it would be more efficient to weight the inequalities differently to account for their differing variances, as well as for covariances between inequalities. The econometrics literature has yet to solve this problem for inequality moment conditions and pursuing it further is beyond the scope of this paper.

and

$$\sqrt{n}(\widehat{\theta} - \theta_0) \xrightarrow{d} N(0, H_0^{-1}\Lambda_0 V_\alpha \Lambda_0' H_0^{-1}).$$

where,

$$H(\theta) \equiv -\mathbb{E}\frac{\partial^2}{\partial\theta\partial\theta'}h(X_k, \theta; \alpha_0),$$

$H_0 = H(\theta_0)$*, and*

$$\Lambda_0 \equiv \mathbb{E}\frac{\partial^2\{g(X_k, \theta_0; \alpha_0) < 0\}g^2(X_k, \theta_0; \alpha_0)}{\partial\theta\partial\alpha'}.$$

Proof. See appendix A. □

Consistency requires both the number of inequalities sampled (n_I) and the number of simulation draws per inequality (n_s) to go to infinity at least as fast as the number of first stage observations (n). As a result of these assumptions, the simulation error does not contribute to the asymptotic distribution of the estimator. The standard errors are thus determined by first stage sampling error, adjusted for its effect on the second stage estimates. Efficiency of the estimator does, however, depend on the distribution used to sample over inequalities through the expectations operators in Λ_0 and H_0.

Derivation of an expression for Λ_0 in terms of model primitives is difficult because of the complex way in which the first stage parameters enter into the W terms. Therefore, in practice we believe it will typically be easiest to use subsampling or the bootstrap to estimate standard errors.[10]

4.3 Case 2: Bounds Estimation

In this section we show how to extend the estimation method to models where point identification is not known to hold. This might arise either because the model is known to be underidentified, or simply because an identification proof is difficult to obtain. We include this section because there have been several recent papers on non-identified models in I.O. (Haile and Tamer (2003), Ciliberto and Tamer (2003)), and because the estimation method extends naturally to this case.

[10]The proposition guarantees that subsampling will work. We do not have a proof that the bootstrap works, but in monte carlo experiments it appears to work.

The estimator used in this section is nearly identical to that described above. The estimator converges to bounds in cases where the model is not identified, and converges to a point if the model is identified. The primary cost of dropping the identification assumption is that the bounds estimator adds a small amount to the computational burden of the estimator.

We maintain all of the assumptions from the last section, *except* we now assume n_I is fixed. Let

$$\widetilde{\Theta}_0 \equiv \{\theta : \theta, \sigma, P \text{ satisfy (5) for all } x_k,\ k = 1, ..., n_I\}.$$

If $x_1, ..., x_{n_I}$ represents all of the optimality inequalities implied by the model, then $\widetilde{\Theta}_0 = \Theta_0$. More generally, $\widetilde{\Theta}_0$ is a superset of Θ_0. By fixing n_I, we consider only estimation of $\widetilde{\Theta}_0$.

Our estimator and the consistency theorem, closely resemble those for the MMD estimator of Manski and Tamer (2002), as well as that of the semiparametric estimator of Haile and Tamer (2003). The distribution theory for our estimator comes from Chernozhukov et al. (2004).

In this section we employ the same estimator as in (10) except that now it is understood that the optimality inequalities are fixed at some chosen values rather than random:

$$Q_n(\widehat{\theta}, \widehat{\alpha}_n) = \inf_{\theta \in \Theta^*} Q_n(\theta, \widehat{\alpha}_n) \tag{11}$$

where

$$Q_n(\theta, \alpha) \equiv \frac{1}{n_I} \sum_{k=1}^{n_I} 1\{g_{n_s}(x_k, \theta; \alpha) < 0\} g_{n_s}(x_k, \theta; \alpha)^2.$$

Following Manski and Tamer (2002), we employ a slightly weakened version of the estimator:

$$\widehat{\Theta}_n \equiv \{\theta :\ Q_n(\theta, \widehat{\alpha}_n) \leq \min_{\gamma \in \Theta^*} Q_n(\gamma, \widehat{\alpha}_n) + \mu_n\} \tag{12}$$

for some $\mu_n > 0$, where $\mu_n \xrightarrow{p} 0$. The term, μ_n, selects level sets of the objective function.

Let $\rho(\widehat{\Theta}_n, \widetilde{\Theta}_0)$ measure the distance from $\widehat{\Theta}_n$ to $\widetilde{\Theta}_0$ and let $\rho(\widetilde{\Theta}_0, \widehat{\Theta}_n)$ measure the distance from $\widetilde{\Theta}_0$ to $\widehat{\Theta}_n$, as follows:

$$\rho(\widehat{\Theta}_n, \widetilde{\Theta}_0) \equiv \sup_{\theta_1 \in \widehat{\Theta}_n} \inf_{\theta_2 \in \widetilde{\Theta}_0} |\theta_1 - \theta_2| \tag{13}$$

$$\rho(\widetilde{\Theta}_0, \widehat{\Theta}_n) \equiv \sup_{\theta_1 \in \widetilde{\Theta}_0} \inf_{\theta_2 \in \widehat{\Theta}_n} |\theta_1 - \theta_2| \tag{14}$$

Proposition 2. *Under the assumptions listed above, (a)*

$$\rho(\widehat{\Theta}_n, \widetilde{\Theta}_0) \xrightarrow{p} 0.$$

Moreover, (b) if

$$\frac{sup_{\theta \in \Theta^*} |Q_n(\theta, \widehat{\alpha}_n) - Q(\theta, \alpha_0)|}{\mu_n} \xrightarrow{p} 0$$

then

$$\rho(\widetilde{\Theta}_0, \widehat{\Theta}_n) \xrightarrow{p} 0.$$

Proof. See appendix A. □

Part (a) of the consistency theorem says that, for large n, every point in $\widehat{\Theta}_n$ is close to a point in $\widetilde{\Theta}_0$. This is guaranteed under conditions similar to those required for standard estimators, and holds even if $\mu_n = 0$ for all n. Part (b) of the consistency theorem says that, for large n, every point in $\widetilde{\Theta}_0$ is close to a point in $\widehat{\Theta}_n$. Essentially this means that all of the set $\widetilde{\Theta}_0$ is eventually captured by the estimator. In order for part (b) to hold, it is necessary that μ_n go to zero slowly enough, with the required rate provided above.

The primary difficulty in estimation is computation of the set $\widehat{\Theta}_n$, which may potentially be defined by a very large number of inequalities. Fortunately, the inequalities are linear in the parameters, and a variety of methods are available for the purpose of solving such systems. One set of techniques in operations research, dating back to Motzkin et al. (1954), uses the fact that $\widehat{\Theta}$ is a convex polyhedron to describe it in terms of its vertices. Computer code for solving for the vertices of the set, $\widehat{\Theta}$, is widely available. Alternatively, Bajari and Benkard (2003) show how to construct a Gibbs sampling procedure that produces simulation draws from a uniform distribution over the set $\widehat{\Theta}$, at very low computational cost. These simulation draws can be used to estimate bounds on the parameters, θ. Their method is particularly efficient at handling large numbers of inequalities, making it ideally suited to this problem. Finally, Manski and Tamer (2002) and Haile and Tamer (2003) use simulated annealing to sample the objective function and then construct one-dimensional bounds on the parameters.

Our model is similar to Example 1 in Chernozhukov et al. (2004) and therefore their subsampling algorithm (Algorithm 2.1) can be used to estimate confidence intervals for the true parameter, θ_0. This procedure is designed to produce a set, $\widehat{\Theta}_{0.95}$ such that

$$Pr(\theta_0 \in \widehat{\Theta}_{0.95}) \geq 0.95.$$

In the context of our model, their procedure can be described as follows:

1. Begin with an initial estimate of $\widehat{\Theta}_0$ representing a level set of the objective function. This estimate would typically be obtained by obtaining an approximate cutoff value such as that implied by the chi-squared (which falsely assumes that the model is identified).

2. For all $\theta \in \widehat{\Theta}_{0.95}$, use subsampling to estimate the distribution of

$$n * Q_n(\theta, \widehat{\alpha}_n).$$

3. Using the subsampling estimates, estimate the cutoff value, $\widehat{c}_n$, such that

$$\sup_{\theta \in \widehat{\Theta}_{0.95}} Pr(n * Q_n(\theta, \widehat{\alpha}_n) < \widehat{c}_n) = 0.95.$$

4. Use this new cutoff value to estimate $\widehat{\Theta}_{0.95}$.

Chernozhukov et al. (2004) shows that this procedure leads to a consistent estimate of the set $\widehat{\Theta}_{0.95}$.

5 Examples

5.1 Dynamic Discrete Choice

In this section we show how the estimation algorithm applies to the dynamic discrete choice example presented earlier. Note that in this simple context the algorithm is very similar to that of Hotz et al. (1993), with the main difference being the second step procedure. The primary advantage of second step procedure in this paper is that it extends easily to more general models such as the dynamic oligopoly example below. We begin with the general dynamic discrete choice model presented in (2) and continue with the Rust (1987) example below.

Much of the past literature (Keane and Wolpin (1997) is an exception) has made the assumption that error terms, $\nu(a)$, are *iid* and have the extreme value distribution. Here, because it generalizes the model without adding to the computational burden of the second stage estimator, we instead assume that $\nu(a)$ is mean zero, *iid* over time, and has a joint normal distribution with variance matrix, V_ν. In that case, the policy function $P(a_j|s)$ takes the form,

$$P(a_j|s) = Pr(v(a_j, s) + \nu(a_j) \geq v(a_k, s) + \nu(a_k), \text{ for all alternatives } k) \tag{15}$$

where $v(a_j, s)$ satisfies

$$v(a_j, s) = u(a_j, s) + \beta \int \int \max_{j'} \left[v(a_{j'}, s') + \nu(a_{j'}) \right] P(d\nu_{j'}) P(ds'|s, a = a_j). \quad (16)$$

The function $v(a_j, s)$ is called the *choice specific value function*. It represents the agent's expected discounted utility from choosing the action j today, excluding today's values of the preference shock, $\nu(a_j)$).

Equation (15) can be used to estimate the policy function associated with this model. Let $f^{(M)}(a, s; \theta)$ be a flexible function in a and s. Applying a probit model to the observed conditional choice data provides flexible estimates of the choice specific value functions, $\widehat{f}^{M}(a, s)$, as well as an estimate of the joint distribution of the preference shockes, $\widehat{V}_\nu$. Given these estimates, it is straightforward to construct the policy function,

$$\sigma_i(s, \nu(a)) = a_j \text{ if } \widehat{f}^{(M)}(a_j, s) + \nu(a_j) \geq \widehat{f}^{(M)}(a_k, s) + \nu(a_k) \text{ for all } k.$$

Suppose the period return function is linear in the parameters,

$$\pi_{ij}(a_j, s; \theta) = \Phi(a_j, s) \cdot \theta.$$

Then, the optimality inequalities are, for all s, σ_i' :

$$\left[\widehat{\mathbb{E}}_{\sigma_i, s} \sum_{t=0}^{\infty} \beta^t \Phi(a_j, s) - \widehat{\mathbb{E}}_{\sigma_i', s} \sum_{t=0}^{\infty} \beta^t \Phi(a_j, s) \right] \cdot \theta + \left[\widehat{\mathbb{E}}_{\sigma_i, s} \sum_{t=0}^{\infty} \beta^t \nu(a_t) - \widehat{\mathbb{E}}_{\sigma_i', s} \sum_{t=0}^{\infty} \beta^t \nu(a_t) \right] \geq 0 \quad (17)$$

The system of inequalities above is the system, (6), in the context of the dynamic discrete choice model. Note that the two systems are essentially identical except that here the term representing the presented discounted value of the unobserved state variables factors out as a separate term.

The policy functions, (16), along with the distributions for $\nu(a)$ and $P(s'|a, s)$ can be used to simulate the expectation terms in (17). Starting at a particular state, $\mathbf{s}_0$, a vector of ν's is drawn from a joint normal with variance matrix, $\widehat{V}_\nu$. The policy, (16), is then implemented to find a_0 and s_1 (depending on the specific model, it may also be necessary to draw s_1 from the distribution $P(\mathbf{s}_1|\mathbf{s}_0, a_0)$). This process is repeated to construct one time series draw on a path of states and actions. Many such draws are then used to simulate the expectations terms in (17).

Note that in this problem there are a continuum of alternative policies, given by alternative functions, $\widetilde{f}(a_j, s)$. In addition, if the discount factor is known, then this model is typically identified (see Rust (1994)).

5.1.1 A Simple Monte Carlo

In this section we provide Monte Carlo evidence for the dynamic discrete choice model using the simple machine replacement model of Rust (1987) presented earlier. In the Monte Carlo exercise we set the maximum age of the machine to $M = 5$, with the remaining parameters set to $\theta = 1$, $R = 4$, and $\beta = 0.9$.

The first stage estimates for this model are the average replacement probabilities in the simulated data for the five states (age $\in \{1, 2, 3, 4, 5\}$). In the second stage estimates we assumed that identification held and implemented the MD estimator. Alternative policies consisted of random draws from a normal distribution centered at the estimated policy from the first stage and with a standard deviation of 0.5. We computed the estimator and standard errors 500 times for various values of n, n_s and n_I. One goal of the experiment is to show that the estimator has practical value so we generally chose very small values for n: 50, 100, 200, 400. The standard error estimates were obtained using 25 subsamples of size $n/2$. Results are shown in Table 1.

The results in the table show several things. First, in this example the estimators are close to unbiased for all but the smallest sample sizes. This need not be the case, but since no smoothing was used in the first stage it is not surprising. Second, the subsampled standard errors are close to, though generally slightly smaller than, the true standard errors. The differences are most likely due to the small sample sizes used, such that the asymptotic approximations are imperfect.

We are also interested in the ability of the estimation approach to work on real world data sets. Standard errors in all three cases were generally quite small. For $n = 400$, the t-statistics were on the order of 6-8, while for $n = 100$ they fell to about 3-4. Even the $n = 50$ case led to t-statistics on the order of 2. These results are quite encouraging as in practice data sets would typically number more than 100 observations and often much more than 400 observations.

5.2 Dynamic Oligopoly

In this section we show how the algorithm applies to the dynamic oligopoly example. We first cover the general model described earlier, and then proceed to a specific example.

5.2.1 Step 1: Estimating the Policy Functions

Standard techniques for estimating demand and supply, production functions, etc., can be applied to estimating the demand function ($q_{it}(\mathbf{s}, \mathbf{p}; \theta_1)$) and cost function ($mc(s_{it}, q_{it}; \theta_2)$). The spot market equilibrium can then be solved for numerically, yielding the static portion of profits at every state (excluding the cost of investment). We denote the static portion of profits as $\tilde{\pi}(\mathbf{a}_t, \mathbf{s}_t)$ below.

The transition probabilities, $P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{I}_t)$, would typically be estimated parametrically using maximum likelihood (see below for more details). We also assume that the first stage estimation provides estimates of (at least) the transition distributions of any unobserved state variables, such that they can be simulated.

It is also necessary to estimate the investment policy function $I(\mathbf{s}, \nu_i)$, as well as the entry probability at each state, $\chi^e(\mathbf{s})$, and the exit policy function, $\chi(\mathbf{s}, \nu_i)$. As these functions are determined in equilibrium, this would typically be done using some nonparametric technique. Again, for details in a specific example, see the next subsection.

5.2.2 Step 2: Estimating the Dynamic Parameters

Step one provides estimates of many of the structural parameters. The remaining parameters to be estimated in step two are the parameters of the cost of investment function, $C(I, \nu_i; \theta_3)$, the scrap value, Ψ, and the entry cost distribution, $F(x^e)$. Our method requires linearity in these parameters. In general, the entry and exit costs enter the period return linearly. We assume that the cost of investment function is also linear in a vector of parameters, θ_3,

$$C(I, \nu_i) = \theta_3 \cdot \Phi(I, \nu_i),$$

for some vector of basis functions, $\Phi(\cdot)$. For example, one function that satisfies this condition is,

$$C(I, \nu_i; \theta_3) = \theta_{3,0} + \theta_{3,1} * I + \theta_{3,2} * I^2 + \nu_i * I.$$

Note that this function accomodates a fixed cost of investment, as well as convex investment costs, and thus is quite liberal.

If an incumbent firm is making its investment and exit choices optimally, then it must be the case that it would not be made better off by choosing an alternative policy, holding the actions of rival firms fixed at equilibrium levels. This implies that for every initial

state $\mathbf{s}$ and every alternative policy, $\sigma'(\mathbf{s}) = (I'(\mathbf{s}), \chi'(\mathbf{s}))$,

$$\begin{aligned}
&\left[\widehat{\mathbb{E}}_{\sigma_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \tilde{\pi}_i(\mathbf{a}_t, \mathbf{s}_t) - \widehat{\mathbb{E}}_{\sigma_i',\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \tilde{\pi}_i(\mathbf{a}_t, \mathbf{s}_t)\right] \\
&+\theta_3 \cdot \left[\widehat{\mathbb{E}}_{\sigma_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \Phi(I_{it}, \nu_{it}) - \widehat{\mathbb{E}}_{\sigma_i',\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \Phi(I_{it}, \nu it)\right] \\
&+\Psi \left[\widehat{\mathbb{E}}_{\sigma_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \{\chi(\mathbf{s}_t) = 1\} - \widehat{\mathbb{E}}_{\sigma_i',\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \{\chi'(\mathbf{s}_t) = 1\}\right] \quad \geq 0
\end{aligned} \tag{18}$$

This system of inequalities can be used to estimate the investment cost function parameters, θ_3, and the scrap value, Ψ.

In order to estimate the entry cost distribution note that, for a potential entrant, i,

$$Pr(i \text{ enters}|\mathbf{s}) = Pr\left(x^e < \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \left(\tilde{\pi}_i(\mathbf{a}_t, \mathbf{s}_t) - C(I_i, \nu_i; \theta_3) + \{\chi = 1\}\Psi\right)\right). \tag{19}$$

Once θ_3 and Ψ have been estimated, the expectations term on the right hand side can be simulated for any state, $\mathbf{s}$. For many states, the left hand side can be estimated directly from the data. Thus, the distribution $F(x^E)$ can be estimated nonparametrically by regressing the left hand side observations on the ride hand side simulated values in order to obtain an estimate of $F(x^e)$.

If, in addition, the parametric form of the entry cost distribution is known, then equation (19) can be used to form a likelihood. Essentially, such an approach would be a straightforward extension of Hotz and Miller (1993). Additionally, when the parametric form for the entry distribution is known, (19) contains some overidentifying information that could be used to help estimate the investment and exit costs. In practice, this would be implemented by estimating the investment and exit costs parameters jointly with the entry cost distribution parameters. We implement both versions of the estimator below.

5.2.3 A Simple Monte Carlo for a Differentiated Products Oligopoly

In this section we estimate a simple version of the dynamic oligopoly model above. The model here is a slight modification of the model computed in Pakes and McGuire (1994). In the model, each firm has one state variable, s_j, representing the quality of its product relative to that of the outside good. Demand is given by a logit demand system where the utility that consumer r derives from good j is,

$$U_{rj} = \gamma h(s_j) + \alpha \ln(y_r - p_j) + \epsilon_{rj}$$

where y_r is income, γ and α are parameters, and ϵ_{rj} is an *iid* logit error term. The function $h(\cdot)$ is a concave function so that there are declining returns to investment.

This ensures that the Markov chain of states generated by the model is stationary. For simplicity, we also assume that all consumers have the same income, $y_r = y$. The total mass of consumers in the market is denoted M.

Each period, firms choose investment levels, $I_j \in \mathbb{R}^+$, optimally to increase their product quality the next period. Note that the choice of investment is continuous. The specification we use for the evolution of product quality is identical to that of Pakes and McGuire (1994). Firm j's investment is successful, in which case its quality moves up by one, with probability

$$aI_j/(1 + aI_j),$$

where a is a parameter. If the firm's investment is not successful, product quality remains unchanged. There is also an outside good, whose quality moves up with probability δ each period. The cost of investment function is given by[11]

$$c(I) = \theta_{3,1} * I.$$

The scrap value is denoted Ψ, and the random entry cost distribution is specified as $U[x^l, x^h]$. Marginal costs are constant at mc.

For the monte carlo experiments, the parameters were set at the values shown in Table 2. Implementing the estimation algorithm is straightforward. However, it is difficult to properly test an estimation algorithm that is designed to be used in cases where equilibrium computation is infeasible. Therefore, in order to keep the equilibrium computation simple, we considered a model in which a maximum of three firms could be active in each period. We then generated data sets of varying numbers of periods, from 100-400. Typically, real world data sets would have more firms than this and fewer periods. However, computing the equilibrium for models with large numbers of firms is prohibitive. Thus, we instead chose period lengths for which the number of total firm-year observations is comparable to available data sets.

We assumed that quantities, prices, product quality, income, and market size were observed, and that the discount rate was known. It would be a straightforward extension to make product quality unobserved. However, in that case it would be necessary to have an instrument, such as an observed cost shifter, in order to estimate the demand side of the model. This would be easy to handle in estimation (and probably would not lead to different results since it would only impact the first stage), but would complicate computation of the equilibrium of the model. All other parameters, including marginal cost, are estimated.

[11] As far as the estimation is concerned, it would also be straightforward to include a shock to the marginal cost of investment as above. Our reason for not doing so here is that it would somewhat complicate the equilibrium computation required to generate the simulated data.

All of the static parameters of the model can be estimated in the first stage using maximum likelihood. The quantity, price, and product quality data was used to estimate the demand parameters in a logit model. Marginal cost was estimated jointly with demand using the static markup formula. The product quality data was also used (separately) to estimate the investment evolution parameters. In general, these procedures resulted in the demand parameters being recovered very precisely, and the investment evolution parameters being recovered somewhat less precisely. These estimates are inputs to the second stage of estimation.

The policy functions (investment, entry, exit) were estimated using local linear regression with a normal kernel.[12] The bandwidths were chosen by eye once and then held constant across all of the monte carlos. The remaining parameters $(\theta_{3,1}, \Psi)$ were then estimated using the MD estimator on the system of inequalities (18). Alternative investment and exit policies were obtained by adding a mean zero normally distributed error term to the estimated first stage investment and exit policies.[13] In all of the monte carlos we used $n_s = 2000$ simulated paths, with each path being of length 80. Standard errors were computed using 20 subsamples of size $n/2$. To show the computational burden of the estimator, for these parameters, the computational burden of the second stage was about one third that of the first stage. Together they took less than one minute for one estimation (not including subsampling), with the difference across sample sizes not substantial enough to comment on. The results are shown in Table 3.

The first thing that we noticed in the results is that, for these small sample sizes, there is a slight bias in the estimates of the exit value. The bias is reduced with more observations, and indeed we found that it goes away entirely if the true first stage functions are used instead of the estimated first stage. We thus conclude that the second stage bias is generated by bias in the first stage local linear estimates. Note, however, that the investment cost parameter estimates are essentially unbiased even for the smallest sample size ($n = 100$). The difference between the two most likely reflects the difference in the number of first stage observations. For the $n = 100$ case there were typically about 250 investment observations, but only about 25 entry or exit observations. Thus, for any given sample size, the investment policy function is substantially better estimated than the entry and exit policy functions.

Similarly to the last example, the subsampled standard errors are on average slightly smaller than the true standard errors. This is again likely due to the small sample sizes used.

[12] We also tried using polynomials in the first stage and found that they did not perform as well.

[13] The alternative investment policies had standard deviation 0.3 and the alternative exit policies had standard deviation 0.5.

We also wanted to get an idea for how well the estimators performed with these small samples in order to evaluate their potential for use on real world data sets. For $n = 400$, the results are extremely precise, with t-statistics on the order of 12-20. For the $n = 100$ case, in which there are very few observations of entry and/or exit, the standard errors are still surprisingly small, with t-statistics averaging between 7 and 10. For such a small data set it is clear that the first stage estimates can not possibly be very accurate pointwise. Thus, the estimation algorithm must be averaging across points in such a way as to come up with precise parameter estimates anyway. Since real world data sets often contain more observations than this last case, we believe that these results support the method's potential in applications.

In order to evaluate the potential for estimating the sunk cost of entry distribution, we chose to ignore the prior knowledge that the entry costs were uniformly distributed and instead estimated the entry cost distribution nonparametrically using a local linear regression on equation (19).[14] The results are shown in Figures 1-3. The figures show that in all three cases the entry cost distribution is recovered surprisingly well. These experiments show that it is easily possible to recover the distribution of entry costs with reasonably sized data sets.

Finally, we also reestimated the model incorporating the correct parametric form for the entry distribution. The most notable feature of the results (see Table 3) was that joint estimation did not improve the estimates of the investment and exit costs parameters. We conclude that in this model entry behavior contains very little information about the investment and exit cost parameters.

6 Conclusion

This paper describes two new estimators for a large class of dynamic environments. The estimators exploit the assumption that observed behavior is consistent with Markov Perfect Equilibrium. In that case, agents' beliefs can be recovered from observations of equilibrium play. Once agents' beliefs are known, the structural parameters can be solved for using the optimality conditions for equilibrium.

The biggest advantage of the approach is that it avoids the need for equilibrium computation. Avoiding equilibrium computation solves two problems. First, computing an equilibrium even just once, for even the simplest of empirical dynamic oligopoly models, can be computationally prohibitive. In contrast, in the monte carlo experiments

[14] Note that these regressions were done without imposing that the fact that the distribution function must be weakly monotonically increasing.

we found that the overall computational burden of the new estimators tends to be no more than that of many commonly used static estimation methods. Second, because equilibrium beliefs are obtained from the data, there is no need for the researcher to make assumptions about which of many potential equilibria is being played.

The primary cost of the approach is that, in reducing the computational burden, some efficiency is compromised. However, the monte carlo experiments show that the approach still works quite well for fairly small data sets. Furthermore, in the context of a game, it is not clear that intermediate approaches, such as those of Aguirregabiria and Mira (2002b), would necessarily improve efficiency.

Both estimators are also conceptually straightforward and relatively easy to program in standard statistical packages. As there are currently many options in the literature for estimating single agent dynamic models, we expect that the estimators will be most useful in the context of dynamic games, a topic that has proven more difficult. Our hope is that the estimators will facilitate future empirical work on applications of dynamic oligopoly.

A Proofs of Propositions 1 and 2

For ease of notation, let

$$h(x, \theta; \alpha) = 1\{g(x, \theta; \alpha) < 0\} g^2(x, \theta; \alpha).$$

The asymptotic objective function is given by

$$Q(\theta; \alpha_0) = Eh(x, \theta; \alpha_0).$$

Lemma 1.

$$sup_{\theta \in \Theta^*} |Q_n(\theta, \widehat{\alpha}_n) - Q(\theta, \alpha_0)| \xrightarrow{p} 0.$$

Proof. Consider first the convergence of the simulator, g_{n_s}. If a smooth simulator is used, g_{n_s} is continuous in θ. Under the assumptions in the text, a standard ULLN applies, giving

$$sup_{\theta \in \Theta^*} |g_{n_s}(x_k, \theta, \alpha_0) - g(x_k, \theta, \alpha_0)| \xrightarrow{p} 0 \text{ for all } k.$$

Since $Q_n(\theta, \alpha_0)$ is continuously differentiable in θ, this guarantees that

$$sup_{\theta \in \Theta^*} |Q_n(\theta, \alpha_0) - Q(\theta, \alpha_0)| \xrightarrow{p} 0.$$

Finally, since $Q_n(\theta, \alpha)$ is continuously differentiable in α and $\widehat{\alpha}_n \xrightarrow{p} \alpha_0$, using a mean value expansion it is easy to show that

$$sup_{\theta \in \Theta^*} |Q_n(\theta, \widehat{\alpha}_n) - Q(\theta, \alpha_0)| \xrightarrow{p} 0.$$

□

Proof of Proposition 1a

Proof. We have assumed identification and we have also assumed that the true parameter lies in a compact set, and the lemma above shows uniform convergence. These are sufficient to show that

$$\widehat{\theta}_n \xrightarrow{p} \theta_0.$$

□

Proof of Proposition 1b

Proof. We will do the asymptotics in the number of first stage observations, n. The reason for this is that the number of inequalities sampled, n_I, and the number of simulation draws per inequality, n_s, are both under the researcher's control.

Differentiating $Q_n(\theta)$, evaluating at $\theta = \theta_0$, and premultiplying by $\sqrt{n}$ gives:

$$\begin{aligned}
&\sqrt{n}\frac{\partial}{\partial\theta}Q_n(\theta_0;\widehat{\alpha}_n) = \\
&\sqrt{n}\frac{1}{n_I}\sum_{k=1}^{n_I}\frac{\partial}{\partial\theta}h(X_k,\theta_0;\alpha_0)+ \\
&\sqrt{n}\frac{1}{n_I}\sum_{k=1}^{n_I}\left(\mathbb{E}\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0) - \frac{\partial}{\partial\theta}h(X_k,\theta_0;\alpha_0)\right)+ \\
&\sqrt{n}\frac{1}{n_I}\sum_{k=1}^{n_I}\left(\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0) - \mathbb{E}\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0)\right)+ \\
&\qquad\qquad \sqrt{n}\frac{1}{n_I}\sum_{k=1}^{n_I}\left(\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\widehat{\alpha}_n) - \frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0)\right) \quad (20)
\end{aligned}$$

The first term is the standard term in the first order expansion. Note that in this case it is always zero (since h is always zero at the true values of the parameters) and thus drops out completely. This is because of the lack of sampling error in the second stage.

The second term is the simulation bias term. Even if an unbiased simulator is used for g, in general h_{n_s} is not unbiased for h, and therefore the derivative term in the expansion is not unbiased either. Doing a second order mean value expansion of one element in the sum for the second term gives:

$$\begin{aligned}
&\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0) - \frac{\partial}{\partial\theta}h_(X_k,\theta_0;\alpha_0) \\
&= \{g(X_k,\theta_0;\alpha_0) < 0\} * 2 * g(X_k,\theta_0;\alpha_0) * [g_{n_s}(X_k,\theta_0;\alpha_0) - g(X_k,\theta_0;\alpha_0)] + \\
&\qquad\qquad \{g^*_{n_s}(X_k,\theta_0;\alpha_0) < 0\} * [g_{n_s}(X_k,\theta_0;\alpha_0) - g(X_k,\theta_0;\alpha_0)]^2 \quad (21)
\end{aligned}$$

where $g^*_{n_s}$ lies between g_{n_s} and g. Taking the expectations with respect to the simulation error gives:

$$\mathbb{E}\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0) - \frac{\partial}{\partial\theta}h_(X_k,\theta_0;\alpha_0) = \{g^*_{n_s}(X_k,\theta_0;\alpha_0) < 0\} * Var(g_{n_s}(X_k,\theta_0;\alpha_0)) \quad (22)$$

Since $g^*_{n_s} \to g$, this term goes to zero at rate n_s. Therefore, so long as n_s goes to infinity faster than $\sqrt{n}$ the second term contributes nothing to the asymptotic variance.

The third term is the simulation variance term. It is mean zero by construction and, if independent draws are used for each inequality, it is the sum of independent terms.

Therefore, a CLT applies and the second term is asymptotically normal with rate $\sqrt{n_I}$ and variance matrix that disappears with n_s. Therefore, under the assumption above that n_s goes to infinity faster than $\sqrt{n}$, this term also contributes nothing to the asymptotic variance.

The fourth term is the first stage sampling error term. Doing a mean value expansion of the fourth term gives the following expression,

$$\sqrt{n}\frac{1}{n_I}\sum_{k=1}^{n_I}\left(\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\widehat{\alpha}_n)-\frac{\partial}{\partial\theta}h_{n_s}(X_k,\theta_0;\alpha_0)\right)= \left(\frac{1}{n_I}\sum_{k=1}^{n_I}\frac{\partial^2 h_{n_s}(X_k,\theta_0;\alpha_n^*)}{\partial\theta\partial\alpha'}\right)\sqrt{n}(\widehat{\alpha}_n-\alpha_0) \quad (23)$$

By a WLLN, consistency of $\widehat{\alpha}_n$, and consistency of h_{n_s},

$$\left(\frac{1}{n_I}\sum_{k=1}^{n_I}\frac{\partial^2 h_{n_s}(X_k,\theta_0;\alpha_n^*)}{\partial\theta\partial\alpha'}\right)\xrightarrow{p}\Lambda_0\equiv\mathbb{E}\frac{\partial^2 h(X_k,\theta_0;\alpha_0)}{\partial\theta\partial\alpha'} \quad (24)$$

Therefore, so long as n_I goes to infinity with n,, the fourth term has limiting distribution,

$$N(0,\Lambda_0 V_\alpha \Lambda_0).$$

Putting all four terms together gives,

$$\sqrt{n}\frac{\partial}{\partial\theta}Q_n(\theta_0,\widehat{\alpha}_n)\xrightarrow{d}N(0,\Lambda_0 V_\alpha\Lambda_0). \quad (25)$$

Let

$$H(\theta)=-\mathbb{E}\frac{\partial}{\partial\theta\partial\theta'}h(X_k,\theta;\alpha_0) \quad (26)$$

and $H_0=H(\theta_0)$. Note that while the hessian of the objective function is discontinuous, we will assume that H_0 exists and is positive definite (i.e., the asymptotic criterion function is twice continuously differentiable at θ_0). This assumption is justified as it would typically be the case for any model that is locally identified.

Under some additional regularity conditions (stochastic differentiability or stochastic equicontinuity), a standard expansion gives:

$$\sqrt{n}(\widehat{\theta}-\theta_0)\xrightarrow{d}N(0,H^{-1}(\Lambda_0 V_\alpha\Lambda_0')H^{-1}). \quad (27)$$

We will not directly show these additional regularity conditions here but instead note that we believe that they are likely to hold. Each term in the objective function is twice differentiable at all points except those where the indicator function is discontinuous. Furthermore, each term is discontinuous at a different set of points. Finally, each term eventually has a negligible impact on the function, and the asymptotic criterion function is assumed to be twice continuously differentiable. □

Proof of Proposition 2

Proof. (a) Part a of the proof is similar to that of Manski and Tamer (2002), Proposition 5a. Since the parameter space, Θ^*, is compact and $\mu_n \xrightarrow{p} 0$, to prove the result it suffices to show that $Q_n(\theta, \widehat{\alpha}_n)$ converges to $Q(\theta, \alpha_0)$ uniformly on Θ^*. This has been established in the lemma above.

(b) The proof of part b is identical to that of Manski and Tamer, Proposition 5b. □

B References

Aguirregabiria, V. and P. Mira (2002a), "Sequential Simulated-Based Estimation of Dynamic Discrete Games", mimeo, Boston University.

Aguirregabiria, V. and P. Mira (2002b), "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models", Econometrica, 70:4, 1519-1543.

Bajari, P., and C. L. Benkard (2003), "Demand Estimation With Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach", Working Paper, Stanford University.

Benkard, C. L. (forthcoming), "A Dynamic Analysis of the Market for Widebodied Commercial Aircraft", *Review of Economic Studies*.

Berry, S. T., and A. Pakes (2000), "Estimation from the Optimality Conditions for Dynamic Controls,", mimeo Yale University.

Berry, S. T., Levinsohn, J., and A. Pakes (1995), "Automobile Prices in Market Equilibrium", Econometrica, 63:4, 841-90.

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

Chernozhukov, V., Hong, H., and E. Tamer (2004), "Parameter Set Inference in a Class of Econometric Models," Working Paper, Princeton University.

Ciliberto, F. and E. Tamer (2004), "Market Structure and Multiple Equilibria in Airline Markets," Working Paper, Princeton University.

Ericson, R., and A. Pakes (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," Review of Economic Studies, 62:1, 53-83.

Gowrisankaran, G. and R. Town (1997), "Dynamic Equilibrium in the Hospital Industry," *Journal of Economics and Management Strategy*, **6:1**, 45-74.

Haile, P. A., and E. Tamer (2003), "Inference with an Incomplete Model of English Auctions," Journal of Political Economy, 111:1, 1-51.

Hotz, V. J., and R. A. Miller (1993), "Conditional Choice Probabilities and the Estimation of Dynamic Models", Review of Economic Studies, 60:3, 497-529.

Hotz, V. J., Miller, R. A., Sanders, S., and J. Smith (1993), "A Simulation Estimator for Dynamic Models of Discrete Choice," Review of Economic Studies, 60, 397-429.

Jofre-Bonet, M. and M. Pesendorfer, "Estimation of a Dynamic Auction Game," Econometrica, (forthcoming).

Keane, M. P. and K. I. Wolpin (1997), "The Career Decisions of Young Men", Journal of Political Economy, 105:3, 473-522.

Manski, C. F., and E. Tamer (2002), "Inference in Regressions with Interval Data on a Regressor or Outcome," *Econometrica,* 70:2, 519-546.

Motzkin, T. S., Raiffa, H., Thompson, G. L. and R.M. Thrall (1953), "The Double Description Method", in Kuhn, H. W. and A. W. Tucker (eds.), *Contributions to Theory of Games*, Volume 2, Princeton University Press.

Olley, G. S., and A. Pakes (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64:6, 1263-1298.

Pakes, A., M. Ostrovsky, and S. Berry (2003), "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)", mimeo Harvard University.

Pakes, A. and P. McGuire (1994), "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model", Rand Journal of Economics, 25:4, 555-89.

Pakes, A., and P. McGuire (2001), "Stochastic Approximation for Dynamic Models: Markov Perfect Equilibrium and the 'Curse' of Dimensionality," *Econometrica*, forthcoming.

Pesendorfer, M. and P. Schmidt-Dengler (2003), "Identification and Estimation of Dynamic Games," mimeo, London School of Economics.

Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", Econometrica, 55:5, 999-1033.

Rust, J. (1994), "Structural Estimation of Markov Decision Processes," in Engle, R. F. and D. L. McFadden (eds.), Handbook of Econometrics, Volume IV, Amsterdam: Elsevier Science.

A Tables and Figures

Table 1: DDC Monte Carlo, 500 Monte Carlo runs, 25 subsamples of size $n/2$

	Mean	SE(Real)	5%(Real)	95%(Real)	SE(Subsampling)
$n = 400$, $n_I = 200$, $n_s = 1000$					
θ	1.00	0.14	0.79	1.24	0.10
R	4.02	0.53	3.24	4.96	0.39
$n = 200$, $n_I = 200$, $n_s = 500$					
θ	0.99	0.18	0.72	1.37	0.17
R	4.00	0.78	2.94	5.95	0.86
$n = 100$, $n_I = 200$, $n_s = 250$					
θ	0.94	0.32	0.47	1.48	0.35
R	3.75	1.26	1.92	5.70	1.15
$n = 50$, $n_I = 200$, $n_s = 150$					
θ	0.89	0.54	0.11	2.03	0.47
R	3.57	2.35	0.60	8.16	2.27

Table 2: Dynamic Oligopoly Monte Carlo Parameters

Parameter	Value	Parameter	Value
Demand:		Investment Cost:	
α	1.5	$\theta_{3,1}$	1
γ	0.1		
M	5	Marginal Cost:	
y	6	mc	3
Investment Evolution		Entry Cost Distribution	
δ	0.7	x^l	7
a	1.25	x^h	11
Discount Factor		Scrap Value:	
β	0.925	Ψ	6

Table 3: Dynamic Oligopoly With Nonparametric Entry Distribution

	Mean	SE(Real)	5%(Real)	95%(Real)	SE(Subsampling)
$n = 400$, $n_I = 500$					
$\theta_{3,1}$	1.01	0.05	0.91	1.10	0.03
Ψ	5.38	0.43	4.70	6.06	0.39
$n = 200$, $n_I = 500$					
$\theta_{3,1}$	1.01	0.08	0.89	1.14	0.05
Ψ	5.32	0.56	4.45	6.33	0.53
$n = 100$, $n_I = 300$					
$\theta_{3,1}$	1.01	0.10	0.84	1.17	0.06
Ψ	5.30	0.72	4.15	6.48	0.72

Table 4: Dynamic Oligopoly With Parametric Entry Distribution

	Mean	SE(Real)	5%(Real)	95%(Real)	SE(Subsampling)
$n = 400$, $n_I = 500$					
$\theta_{3,1}$	1.01	0.06	0.92	1.10	0.04
Ψ	5.38	0.42	4.68	6.03	0.41
x^l	6.21	1.00	4.22	7.38	0.26
x^h	11.2	0.67	10.2	12.4	0.30
$n = 200$, $n_I = 500$					
$\theta_{3,1}$	1.01	0.07	0.89	1.13	0.05
Ψ	5.28	0.66	4.18	6.48	0.53
x^l	6.20	1.16	3.73	7.69	0.34
x^h	11.2	0.88	9.99	12.9	0.40
$n = 100$, $n_I = 300$					
$\theta_{3,1}$	1.01	0.10	0.84	1.17	0.06
Ψ	5.43	0.81	4.26	6.74	0.75
x^l	6.38	1.42	3.65	8.43	0.51
x^h	11.4	1.14	9.70	13.3	0.58

Figure 1: Entry Cost Distribution for $n = 400$

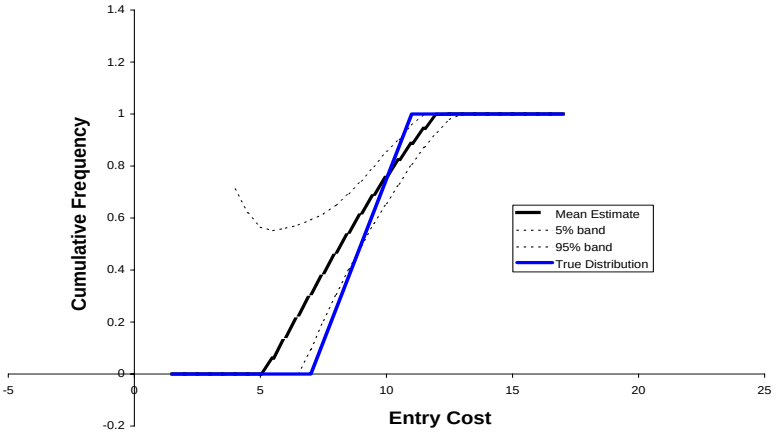

Figure 2: Entry Cost Distribution for $n = 200$

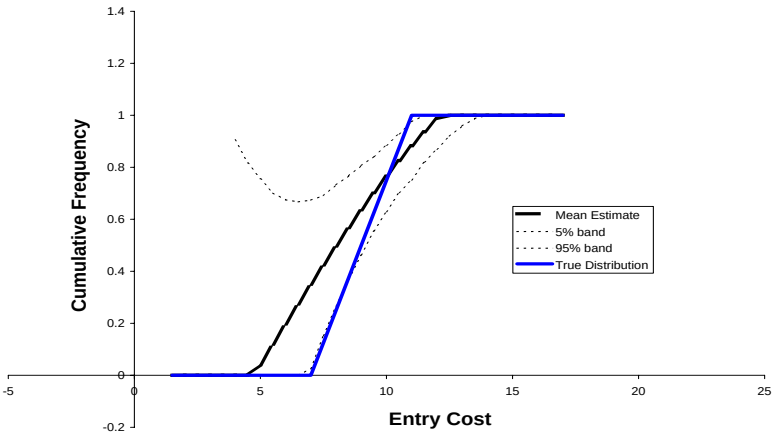

Figure 3: Entry Cost Distribution for $n = 100$