

Estimation of Dynamic Discrete Choice Models in Continuous Time

Peter Arcidiacono¹, Patrick Bayer¹,
Jason R. Blevins², and Paul B. Ellickson²

¹*Duke University and NBER*

²*Duke University*

June 17, 2009

Abstract

This paper provides a method of estimating dynamic discrete choice models (in both single- and multi-agent settings) in which time is a continuous process. The advantage of working in continuous time is that state changes occur sequentially, rather than simultaneously, eliminating a substantial curse of dimensionality that arises in multi-agent settings. Eliminating this computational bottleneck is the key to providing a seamless link between estimating the model and performing post-estimation counterfactuals. In the case of complex discrete games, the models that applied researchers typically estimate (where the curse of dimensionality is broken by using two-step approaches in which agent's beliefs—conditional choice probabilities (CCPs)—are estimated in a first stage) often do not match the models that are then used to perform counterfactuals. Building on the theoretical framework developed by [Doraszelski and Judd \(2008\)](#), we propose an estimation strategy for continuous time discrete choice models that can be implemented either via a full-solution nested fixed point algorithm or using a CCP-based approach. We also consider estimation in situations with imperfectly sampled data, such as when there is an unobserved choice, for example a passive decision to not invest, or when data is aggregated over time, such as when only discrete-time data are available at regularly-spaced intervals.

Keywords: dynamic discrete choice, discrete dynamic games, continuous time.

JEL Classification: C13, C35, L11, L13.

1 Introduction

Empirical models of single-agent dynamic discrete choice (DDC) problems have a rich history in structural applied microeconometrics, starting with the pioneering work of [Miller \(1984\)](#), [Pakes \(1986\)](#), and [Rust \(1987\)](#). These methods have been used to study problems ranging from the optimal age of retirement to the choice of college major. However, due to the inherent complexity of nesting multi-agent DDC problems within iterative estimation routines, these methods were long considered intractable when it came to estimating multi-agent strategic games, despite a growing number of computational methods for solving for their equilibria ([Pakes and McGuire, 1994, 2001](#); [Doraszelski and Satterthwaite, 2007](#)). Recently, in a parallel series of papers, [Aguirregabiria and Mira \(2007\)](#), [Bajari, Benkard, and Levin \(2007\)](#), [Pesendorfer and Schmidt-Dengler \(2007\)](#), and [Pakes, Ostrovsky, and Berry \(2007\)](#), have showed how to extend the two-step estimation techniques, originally developed by [Hotz and Miller \(1993\)](#) and [Hotz, Miller, Sanders, and Smith \(1994\)](#) in the context of single-agent problems, to more complex multi-agent settings.¹ Ironically, the bottleneck is now on the computational side: while empirical researchers can estimate models with state spaces of ever-expanding complexity, post-estimation counterfactuals and simulations are limited by the curse of dimensionality inherent in simultaneous move games. In many cases, the model that researchers estimate is far richer than what they are able to use for simulations, leading some to suggest alternatives to the Markov Perfect Equilibrium concept in which firms condition on long run averages (regarding rivals) instead of current information ([Weintraub, Benkard, and Van Roy, 2008](#)). The goal of this paper is exploit the sequential structure of continuous time games to break the computational curse, create a tight link between estimation and counterfactuals, and open the door to more complex and realistic models of strategic interaction.

Ours is not the first paper to tackle these computational issues. Making full use of computing resources, [Pakes and McGuire \(2001\)](#) extend their seminal approach to solving dynamic games ([Pakes and McGuire, 1994](#)) by replacing explicit integration with simulation and utilizing an adaptive algorithm that targets the recurrent class of states. Their computational approach is able to alleviate the curse of dimensionality arising from the need to calculate expectations over successor states as well as the increasing size of the state space itself. In theoretical work that is closest in structure to ours, [Doraszelski and Judd \(2008\)](#) exploit the structure of continuous time to break the curse of dimensionality associated with the calculation of expectations over rival actions. In particular, because state changes occur only one agent at a time, the dimension of these expectations grows linearly in the number of players, rather than exponentially, resulting in computation times

¹Two-step estimation of dynamic discrete games was originally proposed by [Rust \(1994\)](#). Rust recommended substituting non-parametric estimates of rivals' reaction functions into each player's dynamic optimization problem, turning a complex equilibrium solution into a collection of simpler games against nature.

that are orders of magnitude faster than those of discrete time. Building on these insights, we seek to connect the computational advantages of continuous time with the empirical power of two-step estimation. To do so, we recast the dynamic discrete choice problem in discrete time as a sequential discrete choice problem in continuous time. In particular, rather than choosing simultaneous (and continuous) actions, as in [Doraszelski and Judd \(2008\)](#), players in our framework make sequential, discrete decisions in a random order that is determined by a stochastic jump process.

The use of Markov jump processes to model the state variables combined with the discrete choice nature of the problem results in a simple, but very flexible mathematical structure that is straightforward to estimate and yields sufficient computational gains to make full solution estimation feasible for very large problems. The model also shares many features with standard discrete choice models and so many of the insights and tools used in the discrete time framework, such as two-step CCP (conditional choice probability) estimation, are directly applicable to our continuous time model. Furthermore, the estimators we present are robust to many kinds of imperfectly sampled data, such as when observations are missing (e.g., passive actions such as the decision not to invest) or when data are sampled only at discrete intervals (e.g., quarterly or yearly observations). The model thus offers a seamless link between estimation and computation, allowing the same underlying model to be first estimated, using one of several methods with many different types of data, and later solved in order to simulate counterfactuals.

The structure of the paper is as follows. Section 2 introduces our modeling framework, discusses both the single- and multi-agent settings, and provides concrete (and canonical) examples of both. Section 3 develops our estimators, including both full-solution and two step approaches, and discusses issues associated with partial observation and time aggregation. Section 4 provides several Monte Carlo studies, including both full-solution and CCP estimators for both example models.

2 Model

The models we describe below are based on Markov jump processes. Therefore, before proceeding with the model itself we briefly review some relevant properties and results pertaining to Markov jump processes in Section 2.1. Next, in Section 2.2, we begin by describing a single-agent continuous-time dynamic discrete choice model in order to build intuition. We introduce a simple single-agent renewal model as an example in Section 2.3. We then show in Section 2.4 that extending the simple single-agent model to the case of dynamic discrete games with N players is simply a matter of modifying the intensity matrix governing the market-wide state vector in order to incorporate players' beliefs regarding the actions of their rivals. Following [Harsanyi \(1973\)](#), we treat the dynamic

discrete game as a single-agent game against nature, in which moves by rival agents are distributed in accordance with players' beliefs. As an example, we consider the quality ladder model of [Ericson and Pakes \(1995\)](#) in Section 2.5.

2.1 Markov Jump Processes

A Markov jump process is a stochastic process X_t indexed by $t \in [0, \infty)$ taking values in some discrete state space \mathcal{X} . If we begin observing this process at some arbitrary time t and state X_t , it will remain in this state for a duration of random length τ before transitioning to some other state $X_{t+\tau}$. The trajectory of such a process is a piecewise-constant, right-continuous function of time.

Jumps occur according to a Poisson process and the length of time between jumps is therefore Exponentially distributed. The pdf of the Exponential distribution with rate parameter $\lambda > 0$ is

$$f(x; \lambda) = \begin{cases} \lambda e^{-\lambda x}, & x \geq 0, \\ 0, & x < 0. \end{cases}$$

The cumulative distribution function is

$$F(x; \lambda) = \begin{cases} 1 - e^{-\lambda x}, & x \geq 0, \\ 0, & x < 0. \end{cases}$$

The mean is $\frac{1}{\lambda}$, the inverse of the rate parameter or frequency, and the variance is $\frac{1}{\lambda^2}$.

We consider stationary processes with finite state spaces $\mathcal{X} = \{1, \dots, K\}$. Before proceeding, we first review some fundamental properties of Markov Jump Processes, presented without proof. For details see [Karlin and Taylor \(1975, section 4.8\)](#).

A finite Markov jump process can be summarized by its intensity matrix

$$Q = \begin{bmatrix} -q_{11} & q_{12} & \dots & q_{1K} \\ q_{21} & -q_{22} & \dots & q_{2K} \\ \vdots & \vdots & \vdots & \vdots \\ q_{K1} & q_{K2} & \dots & -q_{KK} \end{bmatrix}$$

where

$$q_{ij} = \lim_{h \rightarrow 0} \frac{\Pr(X_{t+h} = j | X_t = i)}{h} \text{ for all } j \neq i$$

represents the probability per unit of time that the system transitions from i to j and

$$q_{ii} = \sum_{j \neq i} q_{ij}$$

denotes the rate at which the system transitions out of state i . Thus, transitions out of i follow an exponential distribution with rate parameter q_{ii} and, upon leaving state i , the system transitions to $j \neq i$ with probability

$$(1) \quad P_{ij} = \frac{q_{ij}}{\sum_{k \neq i} q_{ik}}.$$

Finally, let $P_{ij}(t) = \Pr(X_{t+s} = j | X_s = i)$ denote the probability that the system transitions to state j after a period of length t given that it was initially in state i and let $P(t) = (P_{ij}(t))$ denote the corresponding matrix of these probabilities. $P(t)$ can be found as the unique solution to the system of ordinary differential equations

$$\begin{aligned} P'(t) &= P(t)Q, \\ P(0) &= I. \end{aligned}$$

frequently referred to as the forward equations. It follows that

$$(2) \quad P(t) = e^{tQ} = \sum_{k=0}^{\infty} \frac{(tQ)^k}{k!}.$$

In some cases, we will need to calculate $P(t)$ for estimation, but in practice we cannot calculate the infinite sum (2) directly. If Q can be diagonalized as $Q = MDM^{-1}$ where $D = (d_{ij})$ is diagonal, then

$$(3) \quad P(t) = Me^{tD}M^{-1}$$

where

$$e^{tD} = \begin{bmatrix} e^{tD_{11}} & 0 & \dots & 0 \\ 0 & e^{tD_{22}} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e^{tD_{KK}} \end{bmatrix}.$$

In particular, a single element of $P(t)$ can be written as

$$(4) \quad P_{ij}(t) = \sum_{k=1}^K m_{ik} e^{td_{kk}} m^{kj}$$

where $M = (m_{ij})$ and m^{ki} denotes the (k, i) element of M^{-1} . However, for a general intensity matrix Q , we compute e^{tQ} using routines from Expokit, a Fortran package for calculating matrix exponentials (Sidje, 1998). We employ algorithms for dense intensity matrices in the case of single agent problems with small state spaces, and we use sparse matrix algorithms for more efficient computation in the case of dynamic games, which tend to have large state spaces and sparse intensity matrices.

Finally, we derive some properties of the Exponential distribution which will be required for constructing the value function later. In particular, we show that if there are n competing Poisson processes (or Exponential distributions) with rates λ_i for $i = 1, \dots, n$, then distribution of the minimum wait time is Exponential with rate $\sum_{i=1}^n \lambda_i$ and, furthermore, conditional on an arrival the probability that it is due to process i is $\lambda_i / \sum_{j=1}^n \lambda_j$.

Proposition 2.1. *If $\tau_i \sim \text{Exponential}(\lambda_i)$ for $i = 1, \dots, n$ and $\tau \equiv \min_i \tau_i$, then*

$$\tau \sim \text{Exponential}(\lambda_1 + \dots + \lambda_n).$$

Proof.

$$\begin{aligned} \Pr(\tau \leq t) &= \Pr\left(\min_i \tau_i \leq t\right) = 1 - \Pr(\tau_1 > t, \dots, \tau_n > t) \\ &= 1 - \prod_{i=1}^n \Pr(\tau_i > t) = 1 - \prod_{i=1}^n e^{-\lambda_i t} = 1 - e^{-(\sum_{i=1}^n \lambda_i)t}. \end{aligned}$$

■

Proposition 2.2. *Let τ_1, \dots, τ_n be random variables and let $\iota = \arg \min_i \tau_i$ be the index of the minimum. If $\tau_i \sim \text{Exponential}(\lambda_i)$, then*

$$\Pr(\iota = i) = \frac{\lambda_i}{\sum_{j=1}^n \lambda_j}.$$

Proof.

$$\begin{aligned} \Pr(\tau_i \leq \tau_j \forall j) &= \mathbb{E}_{\tau_i} [\Pr(\tau_j \geq \tau_i \forall j \neq i) \mid \tau_i] \\ &= \int_0^\infty [e^{-\sum_{j \neq i} \lambda_j t}] \lambda_i e^{-\lambda_i t} dt \\ &= \int_0^\infty \lambda_i e^{-(\sum_{j=1}^n \lambda_j)t} dt \\ &= -\frac{\lambda_i}{\sum_{j=1}^n \lambda_j} [e^{-(\sum_{j=1}^n \lambda_j)t}]_{\tau_i=0}^\infty \\ &= \frac{\lambda_i}{\sum_{j=1}^n \lambda_j}. \end{aligned}$$

■

These two propositions allow us to treat the n competing Poisson processes (τ_1, \dots, τ_n) as a single joint process (τ, ι) where the joint distribution is given above.

2.2 Single-Agent Dynamic Discrete Choice Models

Consider the problem of a single agent when time is a continuous variable $t \in [0, \infty)$. The state of the model at any time t can be summarized by a member x of some finite state space $\mathcal{X} = \{1, \dots, K\}$. Two competing Poisson processes drive the dynamics of the model: a continuous-time Markov jump process on \mathcal{X} with intensity matrix Q_0 which generates state changes and a Poisson arrival process with rate λ which governs when the agent can move (i.e., take an action). The Markov jump process represents moves by nature—state changes that aren't a direct result of actions by the agent. At each time t , if a jump occurs next, the state jumps immediately to the new value. The agent may not influence this process. If a move arrives next, the agent chooses an action a from the discrete choice set $\mathcal{A} = \{1, \dots, J\}$, conditional on the current state $k \in \mathcal{X}$. The set \mathcal{A} contains all possible actions the agent can take when given the opportunity to move.

The agent is forward looking and discounts future payoffs at a rate ρ . Thus, from the perspective of time t_1 , the present discounted value of some payoff π received at time t_2 is $e^{-\rho(t_2-t_1)}\pi$. While the model is in state k , the agent receives flow utility u_k . Thus, if the model remains in state k over the interval $[0, \tau)$, the present discounted value of the payoff obtained over this period from the perspective of time 0 is $\int_0^\tau e^{-\rho t} u_k dt$.

Upon receiving a move arrival at time τ when the current state is $k \in \mathcal{X}$ the agent chooses an action $j \in \mathcal{A}$. The agent then receives an instantaneous payoff $\psi_{jk} + \varepsilon_j$ associated with making choice j in state k , where ε_j is a choice-specific payoff shock that is iid over time and across choices. Let σ_{jk} denote the probability that the agent optimally chooses choice j in state k . Let w_{jk} denote the continuation value received by the agent after making choice j in state k . In most cases, w_{jk} will consist of a particular element of the value function, for example, if the state is unchanged after the action then we might have $w_{jk} = v_k$. On the other hand, if there is a terminal action after which the agent is no longer active, then we might have $w_{jk} = 0$. Finally, there might be uncertainty about the resulting state. In such cases we let ϕ_{jkl} denote the probability with which the model transitions to state l after the agent takes action j in state k , where for each j and k we have $\sum_{l=1}^K \phi_{jkl} = 1$. In many cases, such as an exit decision, these probabilities will be degenerate. In this notation, for example, one might express the future value term as $w_{jk} = \sum_{l=1}^K \phi_{jkl} v_l$. Since there are many possible scenarios, we use the notation w_{jk} for generality.

The instantaneous payoffs ψ_{jk} represent one-time changes to the agent's utility incurred as a direct result of action j . For example, in an entry-exit model ψ_{jk} might represent the

cost of entry or a scrap value earned upon exit. We typically assume, as is standard in the discrete choice literature, that for each j , ε_j follows the standard type 1 extreme value distribution, or Gumbel distribution, with cdf $F(x) = e^{-e^{-x}}$ and that $\varepsilon_j \perp\!\!\!\perp \varepsilon_k$ for all $k \neq j$.

We can now construct the Bellman equation, a recursive expression for the value function v which gives the present discounted value of all future payoffs obtained from starting in some state k and behaving optimally in future periods. Without loss of generality, we use time 0 as the initial time. In state k we have

$$(5) \quad v_k = \mathbb{E} \left[\int_0^\tau e^{-\rho t} u_k dt + e^{-\rho \tau} \frac{1}{\lambda + q_{kk}} \left(\sum_{l \neq k} q_{kl} v_l + \lambda \max_j \{ \psi_{jk} + \varepsilon_j + w_{jk} \} \right) \right].$$

Here we have used Propositions 2.1 and 2.2 to express the expectation over the joint distribution over $(\min_i \tau_i, \arg \min_i \tau_i)$ using first the distribution of $\tau \equiv \min_i \tau_i$ and then, conditional on an event at time τ , the distribution of $l \equiv \arg \min_i \tau_i$.

The first term in (5) represents the flow utility obtained in state k from the initial time until the next event (a move or jump), at time τ . The second term represents the discounted expected future value obtained from the time of the event onward, where $\lambda/(\lambda + q_{kk})$ is the probability that the event is a move opportunity and $q_{kl}/(\lambda + q_{kk})$ is the probability that the event is a jump to state $l \neq k$. The expectation is taken with respect to both τ and ε .

A policy rule is a function $\delta : \mathcal{X} \times \mathcal{E} \rightarrow \mathcal{A}$ which assigns to each state k and vector $\varepsilon = (\varepsilon_1, \dots, \varepsilon_J)$ an action from \mathcal{A} . The optimal policy rule satisfies the following inequality condition:

$$\delta(k, \varepsilon) = j \iff \psi_{jk} + \varepsilon_j + w_{jk} \geq \psi_{lk} + \varepsilon_l + w_{lk} \quad \forall l \in \mathcal{A}.$$

That is, when given the opportunity to choose an action, δ assigns the action that maximizes the agent's expected future discounted payoff. Thus, under the optimal policy rule, the conditional choice probabilities σ_{jk} satisfy

$$\sigma_{jk} = \Pr[\delta(k, \varepsilon) = j | k].$$

2.3 Example: A Single Agent Renewal Model

Our first example is a simple single-agent renewal model, based on the canonical bus engine replacement model analyzed by Rust (1987). The state space is $\mathcal{X} = \{1, \dots, K\}$, and the agent has a binary choice set $\mathcal{A} = \{0, 1\}$. The agent faces a cost minimization problem where the flow cost incurred in some state $k \in \mathcal{X}$ is $u_k = -\beta k$ where $\beta > 0$. The action $j = 0$ represents continuation, where the state remains unchanged, and the choice $j = 1$ causes the state to reset to $k = 1$.

The $K \times K$ intensity matrix for the jump process on \mathcal{X} is

$$Q_0 = \begin{bmatrix} -q_1 - q_2 & q_1 & q_2 & 0 & \dots & 0 \\ 0 & -q_1 - q_2 & q_1 & q_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & -q_1 - q_2 & q_1 & q_2 \\ 0 & 0 & \dots & 0 & -q_1 - q_2 & q_1 + q_2 \\ 0 & 0 & \dots & 0 & 0 & 0 \end{bmatrix}.$$

Thus, the state can only move forward until it reaches the final state K at which point it remains there until changed due to an action taken by the agent. For any state $1 \leq k < K - 1$ the state may jump forward either one state or two (and only one at state $K - 1$). Conditional on jumping, the probabilities of moving forward one state or two are $q_1/(q_1 + q_2)$ and $q_2/(q_1 + q_2)$ respectively.

In the notation of the general model above, the continuation values are

$$w_{jk} = \begin{cases} v_k, & \text{if } j = 0, \\ v_1, & \text{if } j = 1. \end{cases}$$

That is, when in state k and the agent chooses to continue, $j = 0$, the state is unchanged and the continuation value is simply v_k . On the other hand when the agent chooses to reset the state, $j = 1$, the continuation value is v_1 , the present discounted value of being in state 1. Although no cost is incurred from continuation, the agent incurs a one-time cost of c when choosing to reset the state to the initial value:

$$\psi_{jk} = \begin{cases} 0, & \text{if } j = 0, \\ -c, & \text{if } j = 1. \end{cases}$$

The value function for this model is thus

$$v_k = \mathbb{E} \left[\int_0^\tau e^{-\rho t} u_k dt + e^{-\rho \tau} \left(\frac{q_1}{\lambda + q_1 + q_2} v_{k+1} + \frac{q_2}{\lambda + q_1 + q_2} v_{k+2} + \frac{\lambda}{\lambda + q_1 + q_2} \max\{\varepsilon_0 + v_k, -c + \varepsilon_1 + v_1\} \right) \right]$$

for $k \leq k - 2$. It is similar for $k - 1 \leq k \leq K$, with the appropriate adjustments being made at the state space boundary.

If we assume that the ε_j are iid with $\varepsilon_j \sim \text{TIEV}(0, 1)$ then we can simplify this expression further using the closed form representation of the expected future value (cf. [McFadden](#),

1984) and the law of iterated expectations (replacing $E_{\tau,\varepsilon}$ with $E_\tau E_{\varepsilon|\tau}$) to obtain:

$$\begin{aligned} E[\sigma_{0k}(\varepsilon_0 + v_l) + (1 - \sigma_{0k})(-c + \varepsilon_1 + v_1)] &= E[\max\{v_l + \varepsilon_0, v_1 - c + \varepsilon_1\}] \\ &= \ln[\exp(v_l) + \exp(v_1 - c)], \end{aligned}$$

and thus,

$$(6) \quad v_k = E \left[\int_0^\tau e^{-\rho t} u_k dt + e^{-\rho\tau} \left(\frac{q_1}{\lambda + q_1 + q_2} v_{k+1} + \frac{q_2}{\lambda + q_1 + q_2} v_{k+2} + \frac{\lambda}{\lambda + q_1 + q_2} \ln(\exp(v_k) + \exp(v_1 - c)) \right) \right].$$

The value function summarizes the present discounted value of all future cost flows from the perspective of some arbitrary point in time, without loss of generality taken to be time 0, and at an arbitrary state $k \in \mathcal{X}$. Here, τ represents the length of time until the arrival of the next event. At each point in time, the agent makes a decision based on an expected future utility comparison, with the expectation taken with respect to the next event time τ , and ε . Inside the expectation, the first term provides the expected utility accumulated over the time interval $[0, \tau)$. Since the agent does not move during this time, the state evolves undeterred according to the Markov jump process defined by the intensity matrix Q_0 , resulting in a cost flow u_k at each instant. The second term is the present discounted value of future utility from time τ onward, after the next event occurs. At the arrival time τ , the state jumps to $k + l$, $l \in \{1, 2\}$ with probability $q_l / (\lambda + q_1 + q_2)$ and with probability $\lambda / (\lambda + q_1 + q_2)$, the agent gets to move and makes an expected future utility maximizing choice of $j \in \{0, 1\}$. The agent may choose $j = 0$ and simply continue accumulating the flow cost until the next arrival, or choose $j = 1$ and reset the state to 1 by paying a cost c . The Type I Extreme Value assumption also yields closed forms for the associated CCPs:

$$(7) \quad \sigma_{jk} = \begin{cases} \frac{\exp(v_k - v_1 + c)}{\exp(v_k - v_1 + c) + 1}, & \text{if } j = 0, \\ \frac{1}{\exp(v_k - v_1 + c) + 1}, & \text{if } j = 1. \end{cases}$$

2.4 Multi-Agent Dynamic Discrete Games

Now, suppose there are N players indexed by $i = 1, \dots, N$. The state space \mathcal{X} is now a set of vectors of length N , where each component corresponds to the state of player i . Player i 's discount rate is ρ_i . We shall simplify the notation later by imposing symmetry and anonymity, but for generality we index all other quantities by i , including the flow utility in state k , u_{ik} , the choice probabilities, σ_{ijk} , instantaneous payoffs, ψ_{ijk} , and post-move state

transition probabilities, ϕ_{ijkl} .

Although it is still sufficient to have only a single state jump process on \mathcal{X} (with some intensity matrix Q_0) to capture moves by nature, there are now N competing Poisson processes with rates λ_i generating move arrivals for each of N players. The next event in the model is determined by the earliest arrival of one of these $N + 1$ processes.

By assuming that the iid shocks to the instantaneous payoffs are private information of the individual players, we can re-interpret the multi-agent model as a game against nature, and incorporate into the intensity matrix the uncertainty about the moves of rival firms, allowing us to construct the value function for the multi-agent model in much the same way as in the single-agent case. Let τ denote the time of the next event, a state jump or a move opportunity for any player, which is the minimum of a collection of competing Poisson processes with rates given by the intensity matrix Q_0 and the move arrival rates λ_i for $i = 1, \dots, N$.

In the interval between the previous event time and τ , no other events may take place since by definition τ is the time of the next event. At the time of the event, the probability that player i gets to move is proportional to λ_i and the probability that the state jumps from k to $l \neq k$ is proportional to q_{kl} . The denominator of these probabilities is the sum of all of the rates involved, so that the probability that the next event in state k is a move opportunity for player i is

$$\frac{\lambda_i}{\sum_{l=1}^N \lambda_l + q_{kk}},$$

where $q_{kk} = \sum_{l \neq k} q_{kl}$, and the probability that the state jumps from k to m is

$$\frac{q_{km}}{\sum_{l=1}^N \lambda_l + q_{kk}}.$$

As before, let σ_{ijk} denote the probability that action j is chosen optimally by player i in state k . These choice probabilities are determined endogenously in the model. The continuation values are denoted w_{ijk} , and ϕ_{ijkl} denotes the probability that immediately after player i takes action j , the state jumps to another state l .

Given the above notation, the value function for player i in state k is

$$(8) \quad v_{ik} = \mathbb{E} \left[\int_0^\tau e^{-\rho_i t} u_{ik} dt + e^{-\rho_i \tau} \frac{1}{\sum_{i=1}^N \lambda_i + q_{kk}} \left(\sum_{l \neq k} q_{kl} v_{il} + \sum_{l \neq i} \lambda_l \sum_{j=1}^J \sigma_{ljk} \sum_{m=1}^K \phi_{ljk m} v_{im} + \lambda_i \max_j \left\{ \psi_{ijk} + \varepsilon_{ij} + w_{ijk} \right\} \right) \right].$$

This expression is complicated only for the sake of generality. In many applications, it

will be the case that the $\phi_{ljk m}$ terms are degenerate, with deterministic state transitions following moves. Further simplifications are also possible when players are symmetric.

A policy rule in this model is a function $\delta_i : \mathcal{X} \times \mathcal{E}_i \rightarrow \mathcal{A}_i$ which assigns to each state k and vector $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{ij})$ an action from \mathcal{A}_i . Given a set of beliefs σ_{ljk} for each rival $l \neq i$ about the probability that player l chooses j in state k (which enter Q_{-i}), the optimal policy rule satisfies the following condition:

$$(9) \quad \delta_i(k, \varepsilon_i) = j \iff \psi_{ijk} + \varepsilon_{ij} + w_{ijk} \geq \psi_{ilk} + \varepsilon_{il} + w_{ilk} \quad \forall l \in \mathcal{A}_i.$$

That is, when given the opportunity to choose an action, δ_i assigns the action that maximizes the agent's expected future discounted payoff given the specified beliefs. Then, under a given policy rule, the conditional choice probabilities of player i , σ_{ijk} , satisfy

$$(10) \quad \sigma_{ijk} = \Pr[\delta_i(k, \varepsilon_i) = j | k].$$

A *Markov Perfect Equilibrium* is a collection of policy rules $(\delta_1, \dots, \delta_N)$ and a set of beliefs $\{\sigma_{ijk} : i = 1, \dots, N; j = 1, \dots, J; k = 1, \dots, K\}$ such that both (9) and (10) hold for all i .

2.5 Example: A Quality Ladder Model

To illustrate the application to dynamic games we consider the quality ladder model of [Ericson and Pakes \(1995\)](#). This model is widely used in industrial organization and has been examined extensively by [Pakes and McGuire \(1994, 2001\)](#), [Doraszelski and Satterthwaite \(2007\)](#), [Doraszelski and Pakes \(2007\)](#), and several others. The model consists of N players who compete in a single product market. The products are differentiated in that the product of firm i has some quality level $\omega_i \in \Omega$ where $\Omega = \{1, 2, \dots, \bar{\omega}, \bar{\omega} + 1\}$ is the set of possible quality levels, with $\bar{\omega} + 1$ denoting inactive firms. Firms with $\omega_i < \bar{\omega} + 1$ are incumbents.

We consider the particular case of price competition with a single differentiated product where firms make entry, exit, and investment decisions, however, the quality ladder framework is quite general and can be adapted quite easily to model other situations. For example, [Doraszelski and Markovich \(2007\)](#) use this framework in a model of advertising where, as above, firms compete in a differentiated product market by setting prices, but where the state ω_i is the share of consumers who are aware of firm i 's product. [Gowrisankaran \(1999a\)](#) develops a model of endogenous horizontal mergers where ω_i is a capacity level and the product market stage game is Cournot with a given demand curve and cost functions that enforce capacity constraints depending on each firm's ω_j .

2.5.1 State Space Representation

We make the usual assumption that firms are symmetric and anonymous. That is, the primitives of the model are the same for each firm and only the distribution of firms across states, not the identities of those firms, is payoff-relevant. We also assume players share the same discount rate, $\rho_i = \rho$ for all i , and arrival rate, $\lambda_i = \lambda$, for all i . By imposing symmetry and anonymity, the size of the state space can be reduced from the total number of distinct market structures, $(\bar{\omega} + 1)^N$, to the number of possible distributions of N firms across $\bar{\omega} + 1$ states. The set of payoff-relevant states is thus the set of ordered tuples of length $\bar{\omega} + 1$ whose elements sum to N :²

$$\mathcal{S} = \{(s_1, \dots, s_{\bar{\omega}+1}) : \sum_j s_j = N\}.$$

In this notation, each vector $\omega = (\omega_1, \dots, \omega_N) \in \Omega^N$ maps to an element $s = (s_1, \dots, s_{\bar{\omega}+1}) \in \mathcal{S}$ with $s_j = \sum_{i=1}^N \mathbb{1}\{\omega_i = j\}$ for each j .

In practice we map the multidimensional space \mathcal{S} to an equivalent one-dimensional state space $\mathcal{X} \subset \mathbb{N}$ consisting of the integers $\{1, \dots, |\mathcal{S}|\}$. We use the same state-space encoding algorithm as Pakes and McGuire (1994) and Doraszelski and Judd (2008) to convert market structure tuples $s \in \mathcal{S}$ to integers $x \in \mathcal{X}$. The state of the market from the perspective of firm i is uniquely described by two integers (x, ω_i) where $x \in \mathcal{X}$ denotes the market structure and ω_i is firm i 's own quality level. This algorithm was studied in detail by Gowrisankaran (1999b).

2.5.2 Product Market Competition

Again, we make the same product market assumptions as Pakes and McGuire (1994), namely that there is a continuum of consumers with measure $M > 0$ and that consumer j 's utility from choosing the good produced by firm i is given by $g(\omega_i) - p_i + \varepsilon_{ij}$, where ε_i is iid across firms and consumers and has a Type I Extreme Value distribution. Pakes and McGuire (1994) introduce the g function in order to enforce an upper bound on profits. As in Pakes et al. (1993), for some constant ω^* we use the function

$$g(\omega_i) = \begin{cases} \omega_i & \text{if } \omega_i \leq \omega^*, \\ \omega_i - \ln(2 - \exp(\omega^* - \omega_i)) & \text{if } \omega_i > \omega^*. \end{cases}$$

²This representation is space-efficient if N is large relative to $\bar{\omega} + 1$. Otherwise, the algorithm used by Pakes and McGuire (1994), as described in Pakes, Gowrisankaran, and McGuire (1993), will be more parsimonious. See Gowrisankaran (1999b) for details.

Let $\sigma_i(\omega, p)$ denote firm i 's market share given the state ω and prices p . By [McFadden \(1974\)](#), we know that the share of consumers purchasing good i is

$$\sigma_i(\omega, p) = \frac{\exp(g(\omega_i) - p_i)}{1 + \sum_{j=1}^N \exp(g(\omega_j) - p_j)}.$$

In a market of size M , firm i 's demand is $q_i(\omega, p) = M\sigma_i$.

All firms have the same marginal cost $c \geq 0$. Taking as given the prices of other firms p_{-i} , the profit maximization problem of firm i is

$$\max_{p_i \geq 0} q_i(p, \omega)(p_i - c).$$

[Caplin and Nalebuff \(1991\)](#) show that there is a unique Bertrand-Nash equilibrium given by the solution to the first order conditions of the firm's problem:

$$\frac{\partial q_i}{\partial p_i}(p, \omega)(p_i - c) + q_i(p, \omega) = 0.$$

Given the functional forms above, the first order conditions become

$$-(p_j - c)(1 - \sigma_j) + 1 = 0.$$

We solve this nonlinear system of equations numerically using the Newton-Raphson method to obtain the equilibrium prices and the implied profits $\pi(\omega_i, \omega_{-i}) = q_i(p, \omega)(p_i - c)$ earned by each firm i in each state (ω_i, ω_{-i}) .

2.5.3 Incumbent Firms

We consider a simple model in which incumbent firms have three choices upon receiving a move arrival. Firms may continue without investing at no cost, they may invest an amount κ in order to increase the quality of their product from ω_i to $\omega'_i = \max\{\omega_i + 1, \bar{\omega}\}$, or they may exit the market and receive a scrap payment η . We denote these choices, respectively, by the choice set $\mathcal{A}_i = \{0, 1, 2\}$. When an incumbent firm exits the market, ω_i jumps deterministically to $\bar{\omega} + 1$. Associated with each choice j is a private shock ε_{ijt} . These shocks are iid over firms, choices, and time and have a Type I Extreme Value distribution. Given the future value associated with each choice, the resulting choice probabilities are defined by a logit system.

Due to the complexity of the state space, we now introduce some simplifying notation. For any market-wide state $k \in \mathcal{X}$, let $\omega_k = (\omega_{1k}, \dots, \omega_{Nk})$ denote the "decoded" counterpart in Ω^N . In the general notation introduced above, the instantaneous payoff ψ_{ijk} to firm

i from choosing choice j in state k is

$$\psi_{ijk} = \begin{cases} 0 & \text{if } j = 0, \\ -\kappa & \text{if } j = 1, \\ \eta & \text{if } j = 2. \end{cases}$$

Similarly, the continuation values are

$$w_{ijk} = \begin{cases} v_{ijk} & \text{if } j = 0, \\ v_{ijk'} & \text{if } j = 1, \\ 0 & \text{if } j = 2, \end{cases}$$

where state k' is the element of \mathcal{X} such that $\omega_{k'i} = \max\{\omega_{ki} + 1, \bar{\omega}\}$ and $\omega_{k'j} = \omega_{kj}$ for all $j \neq i$. Note that we are considering only incumbent firms with $\omega_{ki} < \bar{\omega} + 1$.

The value function for an incumbent firm in state k is thus

$$v_{ik} = \mathbb{E} \left[\int_0^\tau e^{-\rho t} \pi_{ik} dt + e^{-\rho \tau} \frac{1}{N\lambda + q_{kk}} \left(\sum_{l \neq k} q_{kl} v_{il} + \sum_{l \neq i} \lambda \sum_{j=1}^J \sigma_{ljk} \sum_{m=1}^K \phi_{ljk m} v_{im} \right. \right. \\ \left. \left. + \lambda (\sigma_{i0l}(v_{ik} + \varepsilon_{i0}) + \sigma_{i1l}(v_{ik'} - \kappa + \varepsilon_{i1}) + \sigma_{i2l}(\eta + \varepsilon_{i2})) \right) \right]$$

where π represents the flow profit accruing from product market competition and the expectation is with respect to τ and ε_{ij} for all i and j . Conditional upon moving while in state k , incumbent firms face the following maximization problem

$$\max \{v_{ik} + \varepsilon_{i0}, -\kappa + v_{ik'} + \varepsilon_{i1}, \eta + \varepsilon_{i2}\}$$

resulting in the corresponding choice probabilities

$$\sigma_{i0k} = \frac{\exp(v_{ik})}{\exp(v_{ik}) + \exp(-\kappa + v_{ik'}) + \exp(\eta)}, \\ \sigma_{i1k} = \frac{\exp(-\kappa + v_{ik'})}{\exp(v_{ik}) + \exp(-\kappa + v_{ik'}) + \exp(\eta)}, \\ \sigma_{i2k} = 1 - \sigma_{i0k} - \sigma_{i1k},$$

where as before, k' denotes the resulting state after investment.

2.5.4 Potential Entrants

Whenever the number of active firms n is smaller than N , potential entrants receive the opportunity to enter at a rate λ . Thus, if there are n active firms the rate at which

incumbents receive the opportunity to move is $n\lambda$ but the rate at which *any* type of move opportunity occurs is $(n + 1)\lambda$, the additional λ being for potential entrants. If firm i is a potential entrant with the opportunity to move it has two choices: it can choose to enter ($a_i = 1$), paying a setup cost η^e and entering the market immediately in a predetermined entry state $\omega^e \in \Omega$ (we choose $\omega^e = \lfloor \frac{\bar{\omega}}{2} \rfloor$) or it can choose not to enter ($a_i = 0$) at no cost. Associated with each choice j is a stochastic private payoff shock ε_{ijt}^e . These shocks are iid across firms, choices, and time and are distributed according to the Type I Extreme Value distribution.

In the general notation of Section 2.4, for entrants ($j = 1$) in state k , the instantaneous payoff is $\psi_{i1k} = -\eta^e$ and the continuation value is $w_{i1k} = v_{ik'}$ where k' is the element of \mathcal{X} with $\omega_{k'i} = \omega^e$ and $\omega_{k'j} = \omega_{kj}$ for all $j \neq i$. For firms that choose not to enter ($j = 0$) in state k , we have $\psi_{i0k} = v_{i0k} = 0$. Thus, conditional upon moving in state k , a potential entrant faces the problem

$$\max\{\varepsilon_{i0}^e, -\eta^e + v_{ik'} + \varepsilon_{i1}^e\}$$

yielding the conditional entry-choice probabilities

$$\sigma_{i1k} = \frac{\exp(v_{ik'} - \eta^e)}{1 + \exp(v_{ik'} - \eta^e)}.$$

2.5.5 State Transitions

In addition to state transitions that result directly from entry, exit, or investment decisions, the overall state of the market follows a jump process where at some rate γ , the quality of each firm i jumps from ω_i to $\omega'_i = \min\{\omega_i - 1, 1\}$. This process represents an industry-wide (negative) demand shock, interpreted as an improvement in the outside alternative.

Being a discrete-time model, Pakes and McGuire (1994) assume that each period this industry-wide quality depreciation happens with some probability δ , implying that the quality of all firms falls on average every $1/\delta$ periods. Our assumption of a rate γ is also a statement about this frequency in that $1/\gamma$ is the average length of time until the outside good improves.

Although the integer state-space encoding makes the Q_0 matrix intractable, we can construct it numerically where we map each market structure s to an integer x and the resulting state s' to an integer x' . The (x, x) element of Q_0 for each eligible state x is $-\gamma$ while the corresponding (x, x') element is γ . Note that player i 's state can never enter or leave $\omega_i + 1$ as a result of a move by nature. This is only possible when a firm enters or exits.

3 Estimation

Methods that solve for the value function v directly and use it to obtain the implied choice probabilities for estimation are referred to as full-solution methods. The nested-fixed point (NFXP) algorithm of Rust (1987), which uses value function iteration to compute v inside of an optimization routine which maximizes the likelihood, is the classic example of a full-solution method. Su and Judd (2008) provide an alternative MPEC (mathematical program with equilibrium constraints) approach which solves the constrained optimization problem directly, bypassing the repeated solution of the dynamic programming problem.

CCP-based estimation methods, on the other hand, are two-step methods pioneered by Hotz and Miller (1993) and Hotz et al. (1994) and later extended by Aguirregabiria and Mira (2007), Bajari et al. (2007), Pesendorfer and Schmidt-Dengler (2007), Pakes et al. (2007), and Arcidiacono and Miller (2008). The CCPs are estimated in a first step and used to approximate the value function in a closed-form inversion or simulation step. The approximate value function is then used in the likelihood function to estimate the model using a maximum pseudo-likelihood procedure.

Full-solution methods have the advantage that the exact CCPs are known once the value function is found—they do not have to be estimated—and thus the model can be estimated using full-information maximum likelihood. Although these methods are efficient in the statistical sense, they can become quite costly computationally for complex models with many players or a large state space. Many candidate parameter vectors must be evaluated during estimation and, if the value function is costly to compute, even if solving the model once might be feasible, doing so many times may not be. They are also not robust to multiplicity of equilibria, which can lead to discontinuities in the likelihood function, requiring more robust (and time-consuming) optimization routines. The Su and Judd (2008) MPEC approach provides one solution to both the computational bottleneck and the issue of multiplicity. CCP methods provide another attractive alternative, allowing the value function to be computed very quickly and the pseudo-likelihood function to condition upon the equilibrium that is played in the data.

Our model has the advantage of being estimable via either approach. It breaks one primary curse of dimensionality in that with probability one only a single player moves at any particular instant. Thus, a full-solution methods are more likely to be computationally feasible with our model than with a standard discrete-time framework such as that of Ericson and Pakes (1995), as solved by Pakes and McGuire (1994, 2001). A further benefit of our model is that standard CCP methods may also be used for estimation. Finally, because it is feasible to both estimate *and* solve our model, it preserves a tight link between the estimated model and the model used for post-estimation exercises such as simulating counterfactuals.

This section is organized as follows. We begin by discussing estimation via full-solution

methods with continuous time data in Section 3.1 before turning to cases where true continuous time data are not available. We consider the case when some moves may be unobserved in Section 3.2, and in Section 3.3 we consider the case where fully continuous-time observations are not available and the model is only observed at discrete intervals. Finally, we consider CCP-based estimation in Section 3.4.

3.1 Full-Solution Estimation

Consider a sample of size T consisting of a sequence of tuples $(\tau_t, i_t, a_t, x_t, x'_t)$ each describing a jump or move event where, for each observation t , τ_t is the time interval since the previous event, i_t is the player index associated with this event ($i_t = 0$ indicates a move by nature), a_t is the action taken by player i_t (undefined for moves by nature), x_t denotes the state at the time of the event, and x'_t denotes the state immediately after the event.

Let $L_t(\theta)$ denote the likelihood of observation t given the parameters θ . Before stating the likelihood function explicitly, we must first introduce some additional notation. For some quantities, we use a slightly different notation in this section in order to make the dependence on θ explicit. Let $g(\tau; \lambda) = \lambda e^{-\lambda\tau}$ and $G(\tau; \lambda) = 1 - e^{-\lambda\tau}$ denote, respectively, the pdf and cdf of the exponential distribution with rate parameter λ . For two states x and x' , let $q(x, x'; \theta)$ denote the corresponding element of the intensity matrix $Q_0(\theta)$ and let $p(x, x'; \theta)$ denote the corresponding transition probability, conditional on jumping, as defined in (1). Finally, let $\sigma(i_t, a_t, x_t; \theta)$ denote the conditional choice probability of player i_t taking action a_t in state x_t .

As we will see, $L_t(\theta)$ differs for moves and jumps. When a move is observed it provides information about the rate of move arrivals, through the density of the move arrival process $g(\tau_t; \lambda)$,³ and the payoff parameters, through the CCPs $\sigma(i_t, a_t, x_t; \theta)$. In addition, it provides information about the rate at which jumps occur in state x_t since we observe that a jump *did not* occur over the interval of length τ_t , which happens with probability $1 - G(\tau_t; q(x_t, x_t; \theta))$. Thus, in the case of a move, indicated by $i_t > 0$, the likelihood is

$$(11) \quad L_t(\theta) = g(\tau_t; \lambda) \cdot \sigma(i_t, a_t, x_t; \theta) \cdot [1 - G(\tau_t; q(x_t, x_t; \theta))].$$

When a jump is observed, it provides information first about the rate of jump arrivals in state x_t , through the density $g(\tau; q(x_t, x_t; \theta))$, and the state transitions themselves, through the conditional transition probability $p(x_t, x'_t; \theta)$. Similar to the case of moves considered above, a jump also provides information about the move arrival process with parameter λ since we know that over an interval of length τ_t we *did not* observe a move. Using the cdf,

³Or, for example, in a dynamic game with $n(x_t)$ active players in state x_t , the density would be $g(\tau_t; n(x_t)\lambda)$.

this happens with probability $1 - G(\tau; \lambda)$. Thus, the likelihood of a jump observation is

$$(12) \quad L_t(\theta) = g(\tau_t; q(x_t, x_t; \theta)) \cdot p(x_t, x'_t; \theta) \cdot [1 - G(\tau_t; \lambda)].$$

Combining (11) and (12), we can write the log likelihood for the complete sample of size T as

$$\begin{aligned} \ln L_T(\theta) &= \sum_{t=1}^T \mathbb{1}\{i_t > 0\} \left[\ln g(\tau_t; \lambda) + \ln \sigma(i_t, a_t, x_t; \theta) + \ln [1 - G(\tau_t; q(x_t, x_t; \theta))] \right] \\ &\quad + \sum_{t=1}^T \mathbb{1}\{i_t = 0\} \left[\ln g(\tau_t; q(x_t, x_t; \theta)) + \ln p(x_t, x'_t; \theta) + \ln [1 - G(\tau_t; \lambda)] \right]. \end{aligned}$$

3.2 Partially Observed Moves

We continue using the same notation as in the previous sections but now we suppose that the choice set is $\mathcal{A} = \{0, \dots, J-1\}$ and that only actions a_t for which $a_t > 0$ are observed by the econometrician. This complicates the estimation as now we only observe the truncated joint distribution of move arrival times and actions. Estimating λ using only the observed move times for observations with $a_t > 0$ would introduce a downward bias, corresponding to a longer average waiting time, because there could have been many unobserved moves in any interval between observed moves. Thus, in this setting τ_t is now the interval since the last *observed* event. For simplicity, we will consider only estimation of the single agent model of Section 2.2.

Over an interval where the state variable is constant at x_t , the choice probabilities for each action, $\sigma(a_t, x_t; \theta)$, are also constant. On this interval, conditional on receiving a move arrival, it will be observed with probability $1 - \sigma(a_t = 0, x_t; \theta)$.

For a given state x_t we can derive the likelihood of the waiting times between observed moves by starting with the underlying Poisson process generating the move arrivals. Let $N(t)$ denote the total cumulative number of move arrivals at time t and let $N_a(t)$ denote the number of move arrivals for which the agent chose action a . We will write $N_+(t)$ to denote $\sum_{a>0} N_a(t)$. We also define the waiting time before receiving a move arrival with corresponding action a , $W_a(t)$, defined as the smallest value of $\tau \geq 0$ such that $N_a(t + \tau) - N_a(t) \geq 1$. Let $W_+(t)$ and $W(t)$ be defined similarly.

By the properties of Poisson processes we know that $W(t)$, the waiting time until the next move arrival (both observed and unobserved) is independent of t and has an Exponential distribution with parameter λ . We will show a similar result for $W_+(t)$. Because the probability of truncation (the probability of choosing $a = 0$) depends on x , so will the distribution of $W_+(t)$. We will derive the distribution for intervals where the state

is constant which will be sufficient for the purposes of the likelihood function. We have

$$\begin{aligned}
\Pr(W_+(t) \geq \tau) &= \Pr[N_+(t + \tau) - N_+(t) = 0] \\
&= \sum_{k=0}^{\infty} \Pr[N(t + \tau) - N(t) = k, N_0(t + \tau) - N_0(t) = k] \\
&= \sum_{k=0}^{\infty} \Pr[N(t + \tau) - N(t) = k] \sigma(0, x)^k \\
&= \sum_{k=0}^{\infty} \frac{e^{-\lambda\tau} (\lambda\tau)^k}{k!} \sigma(0, x)^k \\
&= e^{-\lambda\tau} \sum_{k=0}^{\infty} \frac{(\sigma(0, x)\lambda\tau)^k}{k!} \\
&= e^{-\lambda\tau} e^{\sigma(0, x)\lambda\tau} \\
&= e^{-(1-\sigma(0, x))\lambda\tau},
\end{aligned}$$

and therefore the cdf of $W_+(t)$ is

$$\Pr(W_+(t) \leq \tau) = 1 - e^{-(1-\sigma(0, x))\lambda\tau}.$$

For a given x , this is precisely the cdf of the exponential distribution with parameter $(1 - \sigma(0, x))\lambda$.

As in (11) and (12) we have two cases for the likelihood of an observation $(\tau_t, i_t, a_t, x_t, x'_t)$. For a move, when $i_t > 0$, we have

$$L_t(\theta) = g(\tau_t; (1 - \sigma(0, x_t))\lambda) \cdot \frac{\sigma(a_t, x_t; \theta)}{1 - \sigma(0, x_t; \theta)} \cdot [1 - G(\tau_t; q(x_t, x_t; \theta))].$$

The only differences here are the first and second terms in which the parameter λ is now scaled to reflect the rate of *observed* move arrivals and the choice probabilities are scaled to condition on the fact that the move was actually observed. For a jump the likelihood becomes

$$L_t(\theta) = g(\tau_t; q(x_t, x_t; \theta)) \cdot p(x_t, x'_t; \theta) \cdot [1 - G(\tau_t; (1 - \sigma(0, x_t))\lambda)].$$

Here, only one term is different: the probability of *not* having observed a move over the interval τ_t . Estimation can now proceed as usual by constructing and maximizing the log-likelihood function of the full sample.

3.3 Time Aggregation

Suppose we only observe the stochastic process X_t at n discrete points in time $\{t_1, t_2, \dots, t_n\}$. Let $\{x_1, x_2, \dots, x_n\}$ denote the corresponding states. Through the aggregate intensity matrix Q , where $Q = Q_0 + \sum_i Q_i$, these discrete-time observations provide information about the underlying state jump process as well as the rate of move arrivals and the conditional choice probabilities. We use these observations to estimate the parameters of Q_0 and the structural parameters implicit in the conditional choice probabilities σ_{ijk} which appear in Q_i for each $i = 1, \dots, N$. We thus use a full-solution method which solves for the value function for each value of θ in order to obtain the implied CCPs to construct each Q_i .

Let $P(t)$ denote the transition probability function from (2) corresponding to the aggregate intensity matrix Q . These probabilities summarize the relevant information about a pair observations (t_{j-1}, x_{j-1}) and (t_j, x_j) . That is, $P_{x_{j-1}, x_j}(t_j - t_{j-1})$ is the probability of the process moving from x_{j-1} to x_j after an interval of length $t_j - t_{j-1}$. This includes cases where $x_j = x_{j-1}$ as the transition probabilities account for there having been no jump or any of an infinite number of combinations of jumps to intermediate states before coming back to the initial state. The likelihood for a sample $\{(t_j, x_j)\}_{j=1}^n$ is thus

$$\ln L_n(\theta) = \sum_{j=1}^n \ln P_{x_{j-1}, x_j}(t_j - t_{j-1}).$$

To be more concrete, consider the single agent renewal model of Section 2.3 with $K = 5$ states. The intensity matrix Q_0 gives the rates at which the state changes due to nature. Suppose that the state increases one state at rate γ_1 and two states at rate γ_2 . Then, Q_0 for this model is

$$Q_0 = \begin{bmatrix} -\gamma_1 - \gamma_2 & \gamma_1 & \gamma_2 & 0 & 0 \\ 0 & -\gamma_1 - \gamma_2 & \gamma_1 & \gamma_2 & 0 \\ 0 & 0 & -\gamma_1 - \gamma_2 & \gamma_1 & \gamma_2 \\ 0 & 0 & 0 & -\gamma_1 - \gamma_2 & \gamma_1 + \gamma_2 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Let σ_k denote the conditional choice probability of choosing to renew—moving the state back to 1 deterministically—in state k . Note that σ_k is determined endogenously and depends on the parameters θ through the value function as in (7). If λ is the rate at which

moves arrive then Q_1 is

$$Q_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ \lambda\sigma_2 & -\lambda\sigma_2 & 0 & 0 & 0 \\ \lambda\sigma_3 & 0 & -\lambda\sigma_3 & 0 & 0 \\ \lambda\sigma_4 & 0 & 0 & -\lambda\sigma_4 & 0 \\ \lambda\sigma_5 & 0 & 0 & 0 & -\lambda\sigma_5 \end{bmatrix}.$$

The first row contains only zeros because the model remains at state 1 regardless of which action is taken. The remaining diagonal elements are $-\lambda + \lambda(1 - \sigma_k)$ where λ is the rate at which the model potentially *leaves* state k and $\lambda(1 - \sigma_k)$ is the rate at which the state potentially remains unchanged yielding a net exit rate of $-\lambda\sigma_k$. The aggregate intensity matrix in this case is $Q = Q_0 + Q_1$, where the corresponding probability function $P(t)$ is used for estimation.

3.4 CCP-Based Estimation

We consider CCP estimation in terms of the single-agent model, but application to the multi-agent model follows directly and is discussed briefly in Section 3.4.2. CCP estimation relies on finding a mapping from CCPs σ_{jk} to the value function v_k . When separated at the time of the next event, the value function as expressed in (5) contains both terms involving v_k directly, as well as the familiar “social surplus” term which is typically used to obtain the inverse mapping. These extra terms preclude the use of the usual inverse CCP mapping, however, when the value function is separated at the time of the player’s next move, the inverse mapping is straightforward.

The derivation is very similar to the next-event representation of Section 2.2, but we now need to consider that between any two moves, any number of other state jumps could have occurred. For example, if the model is initially in state k and no move arrival occurs on the interval $[0, \tau)$ while the state follows the dynamics of the underlying Markov jump process, we know that the probability of being in any state l at time $t \in [0, \tau)$ is $P_{kl}(t)$, where $P(t)$ are the jump probabilities associated with the intensity matrix Q_0 . The total payoff obtained over $[0, \tau)$, discounted to the beginning of the interval, is therefore $\int_0^\tau e^{-\rho t} \sum_{l=1}^K P_{kl}(t) u_l dt$.

The next-move representation of the value function in state k , is

$$(13) \quad v_k = E \left[\int_0^\tau e^{-\rho t} \sum_{l=1}^K P_{kl}(t) u_l dt + e^{-\rho\tau} \sum_{l=1}^K P_{kl}(\tau) \sum_{j=1}^J \sigma_{jl} (\psi_{jl} + \varepsilon_j + w_{jl}) \right].$$

Note that this simply an alternate representation of the value function in (5), expressed in terms of the next *move* time instead of the next *event* time. Both representations are

equivalent.

The first term above represents the flow utility obtained from the initial time until the first move arrival at time τ . The second term represents the expected instantaneous and future utility obtained from making a choice at time τ . The resulting state l at time τ is stochastic, as is the optimal choice j and, possibly even the continuation value w_{jl} . The expectation operator is needed because τ itself is random and is not known a priori.

If $\varepsilon_j \sim \text{TIEV}(0, 1)$, then the CCPs admit the following closed form:

$$(14) \quad \sigma_{jk} = \frac{\exp(\psi_{jk} + w_{jk})}{\sum_{m=1}^J \exp(\psi_{mk} + w_{mk})}.$$

Suppose we wish to express this probability with respect to another state, say state 1, then we can write

$$(15) \quad \sigma_{jk} = \frac{\exp(\psi_{jk} + w_{jk} - \psi_{j1} - w_{j1})}{\sum_{m=1}^J \exp(\psi_{mk} + w_{mk} - \psi_{m1} - w_{m1})}.$$

The ψ_{jk} 's typically have closed forms in terms of the parameters. Thus, if we know differences in the continuation values $w_{jk} - w_{j1}$, we effectively know the CCPs and can estimate the model. In what follows, we show how to obtain these differences using first stage estimates of the CCPs and a closed form inverse relationship with the value function.

First, note that from (14) we can write

$$(16) \quad \ln \left[\sum_{m=1}^J \exp(\psi_{mk} + w_{mk}) \right] = -\ln \sigma_{ijk} + \psi_{jk} + w_{jk}.$$

The left side of this expression is precisely the closed form for the ex-ante future value term in the value function.

In many model specifications we can then obtain an expression for the differences in (15) by choosing an appropriate normalizing state. We use the example model of Section 2.3 to illustrate this point. In terms of this model, we can write (16), for $j = 1$ as

$$(17) \quad \ln [\exp(v_k) + \exp(v_1 - c)] = -\ln \sigma_{1k} + v_1 - c.$$

Note that the left-hand side of the above equation is exactly the expression in the value function as expressed in (6). Substituting (17) into (6) gives the following expression for

the value function for each state k :

$$\begin{aligned} v_k &= \mathbb{E} \left[\int_0^\tau e^{-\rho t} \sum_{l=1}^K P_{kl}(t) u_l dt + e^{-\rho \tau} \sum_{l=1}^K P_{kl}(\tau) (-\ln \sigma_{1l} + v_1 - c) \right] \\ &= \mathbb{E} \left[\int_0^\tau e^{-\rho t} \sum_{l=1}^K P_{kl}(t) u_l dt - e^{-\rho \tau} \sum_{l=1}^K P_{kl}(\tau) \ln \sigma_{1l} + e^{-\rho \tau} (v_1 - c) \right] \end{aligned}$$

where in the second equality we have used the fact that $v_1 - c$ does not depend on l and that the probabilities $P_{kl}(t)$ must sum to one over $l = 1, \dots, K$. Evaluating the above expression at $k = 1$ and differencing gives

$$v_k - v_1 = \mathbb{E} \left[\int_0^\tau e^{-\rho t} \sum_{l=1}^K [P_{kl}(t) - P_{1l}(t)] u_l dt - e^{-\rho \tau} \sum_{l=1}^K [P_{kl}(\tau) - P_{1l}(\tau)] \ln \sigma_{1l} \right].$$

This expression gives differences in the value function in terms of the conditional choice probability σ_{1l} . With first-stage estimates of σ_{1l} for each l we can use this expression to “invert” the estimated CCPs to obtain an approximation of $v_k - v_1$ which can then be used, along with (7), to approximate $\sigma(a_t, x_t; \theta)$ in the likelihood. The result is a pseudo-likelihood function which can be maximized to obtain an estimate of θ .

3.4.1 Computational Issues

There are three issues we must address in order to calculate v numerically. The first problem is the expectation with respect to τ . Here, we simply approximate the integral using Monte Carlo integration by drawing R values of τ , $\{\tau^s\}_{s=1}^R$, and forming the approximation

$$(18) \quad v_k \approx \frac{1}{R} \sum_{s=1}^R \left[\int_0^{\tau^s} e^{-\rho t} \sum_{l=1}^K P_{kl}(t) u_l dt + e^{-\rho \tau^s} \sum_{l=1}^K P_{kl}(\tau^s) \sum_{j=1}^J \sigma_{jk} (\psi_{jk} + \varepsilon_j + w_{jk}) \right].$$

The second problem is finding a closed form for the flow utility term in (20). Let $b_i(\tau) = \int_0^\tau e^{-\rho s} \sum_j P_{ij}(s) u(x_j) ds$, $B(\tau) = (b_1(\tau), b_2(\tau), \dots, b_K(\tau))^\top$ and $U = (u(x_1), \dots, u(x_n))^\top$. Define $C \equiv -(\rho I - Q)$ for simplicity. Then we can write the first term inside the expectation

in matrix notation as

$$\begin{aligned}
B(\tau) &= \int_0^\tau e^{-\rho s I} P(s) U ds \\
&= \int_0^\tau e^{-\rho s I} e^{s Q} U ds \\
&= \left[\int_0^\tau e^{-s(\rho I - Q)} ds \right] U \\
&= \left[\int_0^\tau C^{-1} C e^{s C} ds \right] U \\
&= C^{-1} \left[\int_0^\tau C e^{s C} ds \right] U \\
&= C^{-1} \left[e^{s C} - e^{0 C} \right]_0^\tau U \\
&= C^{-1} \left[e^{\tau C} - I \right] U.
\end{aligned}$$

Finally, substituting for C we have

$$(19) \quad B(\tau) = -(\rho I - Q)^{-1} \left[e^{-\tau(\rho I - Q)} - I \right] U.$$

The final issue is the expectation over ε in the future value term in (20). We can isolate this term using the law of iterated expectations, replacing $E_{\tau, \varepsilon}$ with $E_\tau E_{\varepsilon|\tau}$. If we then make the standard assumption that the ε_j are iid and distributed according to the Type I Extreme Value distribution then we can simplify this expression using the known closed form for the maximum of J values $\{\delta_1 + \varepsilon_1, \dots, \delta_J + \varepsilon_J\}$:

$$E \left[\max\{\delta_1 + \varepsilon_1, \dots, \delta_J + \varepsilon_J\} \right] = \ln \left[\exp(\delta_1) + \dots + \exp(\delta_J) \right].$$

See, for example, [McFadden \(1984\)](#) for details.

3.4.2 Multi-Agent Models

In dynamic games, in the interval between an arbitrary time $t < \tau_i$ and τ_i , any combination of state jumps and moves by other players may take place. Q_0 describes the dynamics of state jumps, and we can construct a similar intensity matrices Q_i that describe the dynamics of events caused by the actions of rival players. In any state $k \in \mathcal{X}$, player i moves at a rate λ_i which is constant across k .

Thus, the rate at which the model leave state k due to player i is λ_i . The rate at which the model enters another state $l \neq k$, the (k, l) element of Q_i , is given by the sum

$$\lambda_i \sum_{j=1}^J \sigma_{ijk} \phi_{ijkl},$$

which accounts for uncertainty both over the choice and the resulting state. Intuitively, this is the probability of moving to state l expressed as a proportion of λ_i , the rate at which the model leaves state k . Note that we must also allow for the state to remain at k , in which case the diagonal (k, k) element of Q_i is

$$-\lambda_i + \lambda_i \sum_{j=1}^J \sigma_{ijk} \phi_{ijkk}.$$

From the perspective of player i , the dynamics of the model follow an intensity matrix $Q_{-i} \equiv Q_0 + \sum_{j \neq i} Q_j$ which captures all events caused by nature and player i 's rivals. With this intensity matrix in hand, the flow utility portion of the value function can be expressed exactly as before with $P^{-i}(t)$ being constructed using the intensity matrix Q_{-i} : $\int_0^\tau e^{-\rho t} \sum_{l=1}^K P_{kl}^{-i}(t) u_{il} dt$. The value function for player i is then

$$(20) \quad v_{ik} = \mathbb{E} \left[\int_0^{\tau_i} e^{-\rho t} \sum_{l=1}^K P_{kl}^{-i}(t) u_{il} dt + e^{-\rho \tau} \sum_{l=1}^K P_{kl}^{-i}(\tau_i) \sum_{j=1}^J \sigma_{ijl} (\psi_{ijl} + \varepsilon_{ij} + w_{ijl}) \right].$$

CCP estimation can now proceed as in the single agent case by applying the inverse mapping and expressing the value function in differences.

4 Monte Carlo Experiments

4.1 Single Agent Dynamic Discrete Choice

Here, we generate data according to the simple single player binary choice model of Section 2.3. The primitives of the model are the payoff parameter β , the intensity matrix parameters q_1 and q_2 , the reset cost c , the discount rate ρ , and the move arrival rate λ . We fix $\rho = 0.05$ and focus on estimating $\theta = (\lambda, q_1, q_2, \beta, c)$.

To generate data according to this model we first choose values for θ and then use numerical fixed point methods to determine the value function over the state space \mathcal{X} to within a tolerance of $\epsilon = 10^{-6}$ in the relative sup norm. To evaluate the expectation over τ in (6), we use Monte Carlo integration as described in Section 3.4.1, drawing R arrival intervals according to the appropriate exponential distribution and approximating the integral using the sample average. We then use the resulting value function to generate data for various values of T .

In the first set of experiments, described in Table 1, we estimate the model using a full solution method, obtaining the value function through value function iteration for each value of θ while maximizing the likelihood function using the L-BFGS-B algorithm (Byrd, Lu, and Nocedal, 1995; Zhu, Byrd, Lu, and Nocedal, 1997). We generate 100 data

sets over the interval $[0, T]$ with $T = 25,000$ for an average of 10,000 events and then estimate the model under several sampling regimes: true continuous time data, continuous time data when passive actions ($a = 0$, the choice not to renew) are unobserved, and discrete time data with observed at intervals $\Delta \in \{0.625, 1.25, 2.5, 5.0, 10.0\}$. The mean estimates and standard errors are reported. The continuous time samples result in very small standard errors relative to the discrete time data, and as the resolution of the discrete time observations decreases, the standard errors increase, reflecting a loss of information.

We also carry out experiments using CCP estimation in the single agent model with continuous time data. The results are displayed in Table 2 for several values of T between 6,500 and 50,000, resulting in between 2,500 and 20,000 continuous-time events.

4.2 A Dynamic Discrete Game

Our second set of Monte Carlo experiments corresponds to the quality ladder model described in Section 2.5. We obtain estimates of $\theta = (\lambda, \gamma, \kappa, \eta, \eta^e)$ for each of 25 simulated datasets and report the means and standard deviations (in parenthesis).

Table 3 summarizes the results for full-solution estimation, where we obtain the value function using value function iteration for each trial value of θ . We vary the number of players N and the number of quality levels $\bar{\omega}$, reporting K , the number of states from the perspective of player i —the number of distinct (x, ω_i) combinations—as well as the mean estimates and standard errors for 25 replications. In these experiments, we used samples containing $T = 200$ continuous time events in each of $M = 200$ markets.

Table 4 presents similar results obtained using CCP estimation. In this table, we simulated 25 datasets, each consisting of $T = 100$ events in $M = 50$ markets for a total of 5,000 continuous time observations.

References

- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53. [2, 17]
- Arcidiacono, P. and R. A. Miller (2008). CCP estimation of dynamic discrete choice models with unobserved heterogeneity. Working paper, Duke University. [17]
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370. [2, 17]
- Byrd, R. H., P. Lu, and J. Nocedal (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific and Statistical Computing* 16, 1190–1208. [26]

- Caplin, A. and B. Nalebuff (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica* 59, 25–59. [14]
- Doraszelski, U. and K. L. Judd (2008). Avoiding the curse of dimensionality in dynamic stochastic games. Working paper, Harvard University. [1, 2, 3, 13]
- Doraszelski, U. and S. Markovich (2007). Advertising dynamics and competitive advantage. *The RAND Journal of Economics* 38, 557–592. [12]
- Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in IO. In M. Armstrong and R. H. Porter (Eds.), *Handbook of Industrial Organization*, Volume 3, Chapter 30, pp. 1887–1966. North Holland. [12]
- Doraszelski, U. and M. Satterthwaite (2007). Computable markov-perfect industry dynamics: Existence, purification, and multiplicity. Working paper, Harvard University. [2, 12]
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economics and Statistics* 62, 53–82. [4, 12, 17]
- Gowrisankaran, G. (1999a). A dynamic model of endogenous horizontal mergers. *The RAND Journal of Economics* 30, 56–83. [12]
- Gowrisankaran, G. (1999b). Efficient representation of state spaces for some dynamic models. *Journal of Economic Dynamics and Control* 23, 1077–1098. [13]
- Harsanyi, J. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory* 2, 1–23. [3]
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–529. [2, 17]
- Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61, 265–289. [2, 17]
- Karlin, S. and H. M. Taylor (1975). *A First Course in Stochastic Processes* (Second ed.). Academic Press. [4]
- McFadden, D. L. (1974). Conditional logit analysis of qualitative choice analysis. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press. [14]
- McFadden, D. L. (1984). Econometric analysis of qualitative response models. In Z. Griliches and M. Intriligator (Eds.), *Handbook of Econometrics*, Volume 2, Chapter 24, pp. 1396–1456. Amsterdam: Elsevier. [9, 25]

- Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political Economy* 92(6), 1086–1120. [2]
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica* 54, 755–784. [2]
- Pakes, A., G. Gowrisankaran, and P. McGuire (1993). Implementing the pakes-mcguire algorithm for computing markov perfect equilibria in gauss. Unpublished manuscript, Harvard University. [13]
- Pakes, A. and P. McGuire (1994). Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *The RAND Journal of Economics* 25, 555–589. [2, 12, 13, 16, 17]
- Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the ‘curse’ of dimensionality. *Econometrica* 69, 1261–1281. [2, 12, 17]
- Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics* 38, 373–399. [2, 17]
- Pesendorfer, M. and P. Schmidt-Dengler (2007). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75, 901–928. [2, 17]
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55, 999–1013. [2, 8, 17]
- Rust, J. (1994). Estimation of dynamic structural models, problems and prospects: Discrete decision processes. In C. Sims (Ed.), *Advances in Econometrics: Sixth World Congress*, Volume 2, Chapter 4, pp. 119–170. Cambridge University Press. [2]
- Sidje, R. B. (1998). Expokit: A software package for computing matrix exponentials. *ACM Trans. Mathematical Software* 24, 130–156. [6]
- Su, C.-L. and K. L. Judd (2008). Constrained optimization approaches to estimation of structural models. Working paper, University of Chicago. [17]
- Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica* 76, 1375–1411. [2]
- Zhu, C., R. H. Byrd, P. Lu, and J. Nocedal (1997). Algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization. *ACM Transactions on Mathematical Software* 23, 550–560. [26]

Sampling	n	q_1	q_2	λ	β	c
Population	∞	0.050	0.150	0.200	1.500	1.000
Continuous Time	10,000	0.050 (0.002)	0.150 (0.002)	0.200 (0.003)	1.498 (0.051)	0.991 (0.050)
Passive Moves	7,676	0.050 (0.002)	0.150 (0.002)	0.200 (0.008)	1.492 (0.127)	0.985 (0.076)
$\Delta = 0.625$	40,000	0.050 (0.002)	0.150 (0.003)	0.203 (0.014)	1.540 (0.447)	1.080 (0.459)
$\Delta = 1.25$	20,000	0.050 (0.002)	0.150 (0.003)	0.202 (0.010)	1.516 (0.400)	1.049 (0.413)
$\Delta = 2.5$	10,000	0.050 (0.002)	0.150 (0.003)	0.202 (0.011)	1.490 (0.465)	1.027 (0.518)
$\Delta = 5.0$	5,000	0.051 (0.003)	0.150 (0.004)	0.216 (0.094)	1.464 (0.700)	1.037 (0.728)
$\Delta = 10.0$	2,500	0.051 (0.004)	0.151 (0.008)	0.223 (0.095)	1.574 (0.873)	1.156 (0.986)

The mean and standard deviation (in parenthesis) of the parameter estimates for 100 different samples are given for various sampling regimes. Δ denotes the observation interval for discrete time data. n denotes the average number of observations (continuous-time events or discrete-time intervals) when observing the model on the interval $[0, T]$. We fixed the discount rate, $\rho = 0.05$, the number of states, $K = 10$, and the number of draws used for Monte Carlo integration, $R = 25$.

Table 1: Single Player Monte Carlo Results ($T = 25,000$)

T	n	q_1	q_2	λ	β	c
∞	∞	0.050	0.150	0.200	1.500	1.000
50,000	20,000	0.050 (0.001)	0.149 (0.002)	0.200 (0.002)	1.553 (0.033)	0.993 (0.033)
25,000	10,000	0.050 (0.002)	0.150 (0.002)	0.200 (0.003)	1.554 (0.054)	0.988 (0.050)
12,500	5,000	0.050 (0.002)	0.149 (0.004)	0.200 (0.004)	1.568 (0.081)	0.991 (0.067)
6,500	2,500	0.050 (0.003)	0.150 (0.005)	0.201 (0.006)	1.554 (0.106)	0.971 (0.086)

The mean and standard deviation (in parenthesis) of the parameter estimates for 100 different samples are given for various choices of T , the length of the observation window. n denotes the average number of observations (continuous-time events) when observing the model on the interval $[0, T]$. We fixed the discount rate, $\rho = 0.05$, the number of states, $K = 10$, and the number of draws used for Monte Carlo integration, $R = 25$.

Table 2: Single Player Monte Carlo Results: CCP Estimation

N	$\bar{\omega}$	K	λ	γ	κ	η	η^e
	Population		1.000	0.500	2.000	6.000	8.000
10	9	437 580	0.998 (0.016)	0.498 (0.004)	2.017 (0.103)	5.970 (0.131)	7.976 (0.100)
11	9	831 402	0.998 (0.015)	0.498 (0.004)	1.999 (0.108)	5.979 (0.147)	7.983 (0.095)
12	9	1 511 640	0.999 (0.015)	0.498 (0.004)	2.029 (0.084)	5.940 (0.097)	7.956 (0.090)
13	9	2 645 370	0.998 (0.014)	0.497 (0.005)	2.024 (0.088)	5.971 (0.151)	7.990 (0.113)
14	9	4 476 780	1.000 (0.010)	0.498 (0.050)	2.004 (0.602)	6.005 (0.133)	8.046 (0.184)
15	9	7 354 710	1.003 (0.015)	0.499 (0.005)	2.059 (0.135)	5.933 (0.094)	7.957 (0.107)
16	9	11 767 536	1.007 (0.009)	0.498 (0.006)	1.958 (0.131)	6.038 (0.125)	8.065 (0.144)

The mean and standard deviation (in parenthesis) of the parameter estimates for 25 different samples are given for various choices of N and $\bar{\omega}$. K denotes the number of distinct states or (x, ω_i) combinations. The number of markets and observed events were held fixed at $M = 200$ and $T = 100$ for a total of 20,000 continuous-time events. We fixed $\rho = 0.05$ and use $R = 25$ draws for Monte Carlo integration.

Table 3: Quality Ladder Monte Carlo Results

N	$\bar{\omega}$	K	$N\lambda$	γ	κ	η	η^e
	Population		1.0000	0.5000	2.0000	6.0000	8.0000
2	5	30	1.0065 (0.0191)	0.5005 (0.0116)	2.0156 (0.0669)	5.9271 (1.0626)	7.9357 (0.9637)
3	8	360	1.0049 (0.0174)	0.5001 (0.0117)	1.9390 (0.0710)	6.0659 (0.6442)	8.0777 (0.5798)
4	10	2860	1.0043 (0.0160)	0.4997 (0.0112)	1.9490 (0.0758)	6.0774 (0.4441)	8.0242 (0.4128)
5	9	6435	1.0009 (0.0245)	0.4978 (0.0072)	1.9896 (0.1110)	6.0212 (0.4004)	8.0300 (0.3433)
5	10	10010	1.0047 (0.0185)	0.4993 (0.0097)	1.9700 (0.0762)	6.0055 (0.3529)	7.9803 (0.2896)

The mean and standard deviation (in parenthesis) of the parameter estimates for 25 different samples are given for various choices of N and $\bar{\omega}$. K denotes the number of distinct states or (x, ω_i) combinations. The number of markets and observed events were held fixed at $M = 50$ and $T = 100$.

Table 4: Quality Ladder Monte Carlo Results: CCP Estimation