# Identification and Estimation of Continuous Time Dynamic Discrete Choice Games*

JASON R. BLEVINS

*The Ohio State University*

May 8, 2018†
Preliminary and Incomplete Draft

**Abstract.** We consider the theoretical, computational, and econometric properties of a class of continuous time dynamic discrete choice games with stochastically sequential moves, recently introduced by Arcidiacono, Bayer, Blevins, and Ellickson (2016). We first establish existence of a Markov perfect equilibrium in a model with more general forms of heterogeneity across firms and states. Then we establish nonparametric identification when only discrete-time observations are available under slightly weaker conditions in the more general model with heterogeneity. Our conditions include cases where the overall decision rates may be unknown. We illustrate our results using three canonical models that provide the foundation for many applications in applied microeconomics: a single agent renewal model, a dynamic model of entry and exit, and a quality ladder model of oligopoly dynamics. Using these example models, we also examine the computational properties of the model and the statistical properties of estimators for the model through a series of small- and large-scale Monte Carlo experiments. Computing and estimating the model remains computationally tractable even in our largest experiment, which has over 58 million states, reflecting the large scale of empirical models with many heterogeneous firms.

**Keywords:** Continuous time, Markov decision processes, dynamic discrete choice, dynamic games, identification.

**JEL Classification:** C13, C35, C62, C73.

*I am grateful to seminar participants at Columbia University, Northwestern University, the University of British Columbia, the University of Chicago, and the University of Iowa and conference attendees at the 2013 Meeting of the Midwest Econometrics Group, the 2014 Meeting of the Midwest Economics Association, the 2014 University of Calgary Empirical Microeconomics Workshop, and the 2015 Econometric Society World Congress for useful comments and discussions. This work builds on previous joint research with Peter Arcidiacono, Patrick Bayer, and Paul Ellickson and has benefited tremendously from our many discussions together.

†This is a revision of a working paper whose first draft appeared in April 2014.

## 1. Introduction

This paper studies continuous-time econometric models of dynamic discrete choice games. Recent interest in continuous time dynamic games by Doraszelski and Judd (2012) and Arcidiacono, Bayer, Blevins, and Ellickson (2016) (henceforth ABBE) and others was motivated by their ability to allow researchers to compute and estimate more realistic, large-scale games and to carry out a wide array of complex counterfactual policy experiments which were previously infeasible due to either computational or theoretical limitations.

The goals of this paper are to extend several existing results for the ABBE model. Specifically, we generalize existing equilibrium characterizations, verify the identification conditions for specific models, examine the computational properties of the model, and study the small sample performance of the proposed two-step semiparametric estimator for the model parameters when only discrete time data are available. Our results build on the large body of work on the theoretical, computational, and econometric properties of discrete-time dynamic games as well as the recent theoretical and applied work using continuous time dynamic games. They also relate to the mathematical literature on continuous time stochastic processes and to work in linear algebra on matrix exponentials and matrix logarithms.

Modeling economic processes in continuous time dates back at least several decades and includes work in time series econometrics by Phillips (1972, 1973), Sims (1971), Geweke (1978), and Geweke, Marshall, and Zarkin (1986). For longitudinal models, Heckman and Singer (1986) promoted using continuous time models instead of discrete time models and put forward two main arguments for doing so. First, for most economic models there is typically not a natural, fixed time interval at which agents make decisions. Allowing agents to make decisions at (possibly unknown and stochastic) points in continuous time can be both more natural and easier computationally. Second, even if there is a compelling reason to use a model in which agents make decisions at fixed time intervals, there is no reason that this decision interval should coincide exactly with the interval at which the researcher observes or samples the process, which is typically annually or quarterly. While standard dynamic discrete choice models have a different functional forms when applied to different time intervals, continuous time models are invariant to the interval at which observations are recorded.

Despite this early work on continuous time models, discrete time dynamic discrete choice models are currently the de facto standard and now have a long, successful history in structural applied microeconometrics starting with the pioneering work of Gotz and McCall (1980), Miller (1984), Pakes (1986), Rust (1987), and Wolpin (1984). A recent series of papers (Aguirregabiria and Mira, 2007; Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky,

and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008) have shown how to extend two-step estimation techniques, originally developed by Hotz and Miller (1993) and Hotz, Miller, Sanders, and Smith (1994) in the context of single-agent dynamics, to more complex multi-agent settings. The computation of multi-agent models remains formidable, despite a growing number of methods for solving for equilibria (Pakes and McGuire, 1994, 2001; Doraszelski and Satterthwaite, 2010).

Dynamic decision problems are naturally high-dimensional and the computational challenges involved are even greater in the context of strategic games, where, traditionally, the simultaneous actions of players introduces a further dimensionality problem. In order to solve for optimal policies, one must calculate players' expectations over all combinations of actions of their rivals. The cost of computing these expectations grows exponentially in the number of players, making it difficult or impossible to compute the equilibrium in many economic environments. This unfortunate reality has severely limited the scale and the degree of heterogeneity in applied work using these methods.

Because of these limitations some authors have recently considered continuous time models which more closely reflect the nature and timing of actions by agents in the models while also reducing the computational burden. Doraszelski and Judd (2012) showed that continuous-time dynamic games have desirable computational properties, significantly decreasing the computational burden required to evaluate the Bellman operator, which can be used to compute equilibria. ABBE demonstrated the empirical tractability of continuous-time games, particularly for applications in industrial organization. They proposed an econometric model which still benefits from the computational aspects of continuous time models but more closely parallels the discrete choice features of discrete time models. They proposed a two-step conditional choice probability (CCP) estimator for their model, thus connecting continuous time games with a long line of work on estimation of discrete time dynamic games. They showed that it is feasible to estimate even extremely large-scale games, but that it is also now possible to carry out counterfactuals in those games, which would have been computationally prohibitive in a simultaneous-move discrete time model. ABBE demonstrated these advantages in the context of an empirical application which analyzed the entry, exit, expansion, contraction of grocery chain stores in urban markets throughout the United States from 1994–2006 with a particular focus on the effects of Walmart's entry into this sector.

The ABBE model was developed to make estimation of large-scale models in industrial organization feasible along with counterfactual simulations using those models. Similar continuous time models have subsequently been used in several applications including Schiraldi, Smith, and Takahashi (2012) to supermarkets in the U.K., Lee, Roberts, and Sweeting (2012) to baseball tickets, Nevskaya and Albuquerque (2012) to online games,

Jeziorski (2013) to the U.S. radio industry, and Cosman (2014) to bars in Chicago.

In addition to expanding on identification and estimation of the model, we use the ABBE model to carry out a series of Monte Carlo experiments. In solving for an equilibrium of a quality ladder model similar to that of Ericson and Pakes (1995), the computational time required is many orders of magnitude smaller than for a comparable discrete time game. A proposed two-step estimator is shown to provide additional computational gains even in the continuous time setting, where it is again orders of magnitude faster than the full solution estimator.

The remainder of this paper is organized as follows. In Section 2, we review a generalized version of the ABBE model that permits additional heterogeneity in the form of firm-specific discount rates and move arrival rates that may vary by firm and state. We establish a linear representation of the value function in terms of CCPs as well as the existence of a Markov perfect equilibrium in the more general model. We then develop new identification results for the model in Section 3 and show how to estimate the model using ideas from several widely-used approaches for estimating discrete time dynamic games in Section 4. We use two canonical examples throughout the paper to illustrate our results: a single agent renewal model based on Rust (1987) and a $2 \times 2$ entry model similar to example models used by Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008), and others. In Section 5 we introduce a third example: a quality ladder model of oligopoly dynamics based on the model of Ericson and Pakes (1995). Finally, in Section 6 we provide Monte Carlo evidence of the small sample performance of the estimator proposed by ABBE and demonstrate that the computational performance of the proposed methods is dramatically improved relative to the discrete time counterpart.

## 2. A Continuous Time Dynamic Discrete Choice Game with Stochastically Sequential Moves

We consider infinite horizon games in continuous time indexed by $t \in [0, \infty)$ with $N$ players indexed by $i = 1, \ldots, N$. The model we introduce is a heterogeneous generalization of the ABBE model where firms may have different discount rates and where the move arrival rates may differ by firm and across states. After formalizing the components of the structural model, we establish a linear representation of the value function in terms of conditional choice probabilities, as in Pesendorfer and Schmidt-Dengler (2008) and ABBE, as well as existence of a Markov perfect equilibrium in the more general model. We then show that the model yields a reduced form characterized by finite-state continuous-time Markov jump processes and discuss an alternative "clock process" that yields a convenient discrete time Markov process representation with desirable computational properties. We

conclude with a comparison of discrete- and continuous-time models.

### 2.1. The Structural Model and Basic Assumptions

*State Space* At any instant, the payoff-relevant market conditions that are common knowledge to all players can be summarized by a state vector $x$, which is a member of a finite state space $\mathcal{X}$ with $K = |\mathcal{X}| < \infty$. The states $x \in \mathcal{X}$ are typically represented as vectors of real numbers in a finite-dimensional Euclidean space. For example, $x = (x_0, x_1, \ldots, x_N)$ where the components $x_i$ are player-specific states for $i = 1, \ldots, N$, such as the number of stores operated by a retail chain, and $x_0$ is an exogenous market characteristic, such as population.

Because the state space is finite there is an equivalent, *encoded* state space representation $\mathcal{K} = \{1, \ldots, K\}$. We will use $\mathcal{K}$ throughout most of the paper, as it allows us to vectorize payoffs, value functions, and other quantities, but $\mathcal{X}$ is typically the most natural way to represent the state of the market.

*Endogenous State Changes* Player $i$ makes decisions in state $k$ at random times which occur according to a Poisson process with rate $\lambda_{ik}$. That is, at any point in time and in any state $k$ the holding time $\tau_i$ until the next move by player $i$ is exponentially distributed with rate parameter $\lambda_{ik}$. We assume these processes are independent across players, that the rates $\lambda_{ik}$ are known to the researcher, and that $\lambda_{ik} < \infty$ for all $i$ and $k$ reflecting the fact that monitoring the state is costly and so continuous monitoring ($\lambda_{ik} = \infty$) is infeasible.

When a move arrival for firm $i$ occurs, it chooses one of $J + 1$ actions in the choice set $\mathcal{J} = \{0, 1, 2, \ldots, J\}$. Endogenous state changes are induced by the actions of players. When player $i$ makes some action $j > 0$ in state $k$, the state jumps immediately and deterministically from $k$ the continuation state $l(i, j, k)$. By definition, the action $j = 0$ for each player represents inaction and so the continuation state is $l(i, 0, k) = k$ for all players $i$ and states $k$.

We now introduce two examples which we will return to throughout the paper to illustrate other assumptions and results.

**Example 1** (Single-Agent Renewal Model)**.** Consider a continuous-time version of the single-agent renewal model of Rust (1987). There is a single agent in this model, the manager of a bus company, so $N = 1$. The only state variable in the model is the accumulated mileage of a bus engine, so the state space can be represented as $\mathcal{K} = \{1, \ldots, K\}$. Suppose the manager decides whether or not to replace a bus engine ($J_1 = 1$) at a rate $\lambda$, which is constant across states. When a decision is made, the choice set is $\mathcal{J} = \{0, 1\}$. Importantly, this does not imply that the individual rates of replacement and

non-replacement are constant. There is a single agent in this model ($N = 1$)

**Example 2** ($2 \times 2$ Entry Model)**.** Consider a simple two-player entry game with a binary exogenous state variable. Each firm $i \in \{1, 2\}$ has two actions $j \in \{0, 1\}$. The exogenous state represents the level of demand, which can either be high or low. The state vector $x$ has three components: $x_1$ and $x_2$ are activity indicators for firms 1 and 2 and the level of demand is represented by $x_3 \in \{\text{L}, \text{H}\}$. Therefore, in vector form the state space is

$$\begin{aligned}\mathcal{X} = \{ \quad & (0,0,\text{L}), \quad (1,0,\text{L}), \quad (0,1,\text{L}), \quad (1,1,\text{L}), \\ & (0,0,\text{H}), \quad (1,0,\text{H}), \quad (0,1,\text{H}), \quad (1,1,\text{H}) \quad \}\end{aligned}$$

In encoded form, the state space is simply $\mathcal{K} = \{1, 2, 3, 4, 5, 6, 7, 8\}$.

*Payoffs* At each move arrival in state $k$, player $i$ also observes a vector $\varepsilon_{ik} = (\varepsilon_{i0k}, \ldots, \varepsilon_{iJk})^\top$ of choice-specific variables. These variables are private information of firm $i$. All firms and the researcher observe the state $k$ (and hence, $x_k$), but only firm $i$ observes $\varepsilon_{ik}$.

In discrete time models, because the actions and state changes effectively resolve at the same time, the period payoffs are defined to be functions of the actions of all firms, the state, and the unobservables. In our continuous-time model, the payoffs resulting from competition in the product market accrue as flows while the choice-specific payoffs (e.g., entry costs) accrue all at once. Therefore, we distinguish between the flow payoffs, denoted $u_{ik}$, and the instantaneous choice-specific payoffs, denoted $\psi_{ijk}$. Importantly, the choice-specific payoffs can depend on the market structure because they are indexed by $k$.

*Exogenous State Changes* In addition to move arrivals, a second type of event that can occur are exogenous state transitions attributed to nature ($i = 0$). When the model is in state $k$, let $q_{0kl}$ denote the rate at which transitions to state $l \neq k$ occur. The rate $q_{0kl}$ may be zero if the transition is not possible, or it may be some positive but finite value. Therefore, the overall rate at which the system leaves state $k$ is $\sum_{l \neq k}^{K} q_{0kl}$. Accordingly, we define $q_{0kk} = -\sum_{l=1}^{K} q_{0kl}$. The individual elements $q_{0kl}$ for $k \neq l$ are usually specified by the researcher to depend on a lower-dimensional vector of parameters $\theta$.

**Example 1, Continued.** Suppose the exogenous state transition process is characterized by two rate parameters $q_1$ and $q_2$ governing one- and two-state mileage jumps. The agent faces a cost minimization problem where the flow utility $u_{ik}$ is the cost of operating a bus with mileage $k$. Upon continuation, no cost is paid but a fixed amount is paid to replace the engine:

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

Upon continuation the state does not change but upon replacement the state jumps immediately to $k = 1$:

$$l(i, j, k) = \begin{cases} k & \text{if } j = 0, \\ 1 & \text{if } j = 1. \end{cases}$$

Following either choice, the agent receives an iid shock $\varepsilon_{ijk}$.

*Assumptions* Before turning to the equilibrium, we pause and collect our assumptions so far.

**Assumption 1** (Discrete States)**.** The state space is finite: $K = |\mathcal{X}| < \infty$.

**Assumption 2** (Discount Rates)**.** The discount rates $\rho_i > 0$ are known for all $i = 1, \ldots, N$.

**Assumption 3** (Move Arrival Rates)**.** The rates of move arrivals and exogenous state changes are strictly positive and bounded: for $i = 1, \ldots, N$ and $k = 1, \ldots, K$, $0 < \lambda_{ik} < \infty$ and $0 \leq q_{0kl} < \infty$ for all $l \neq k$.

**Assumption 4** (Bounded Payoffs)**.** The flow payoffs and choice-specific payoffs satisfy $|u_{ik}| < \infty$ and $|\psi_{ijk}| < \infty$ for all $i = 1, \ldots N$, $j = 0, \ldots, J$, and $k = 1, \ldots, K$.

**Assumption 5** (Additive Separability)**.** The instantaneous payoffs are additively separable as $\psi_{ijk} + \varepsilon_{ijk}$.

**Assumption 6** (Distinct Actions)**.** For all $i = 1, \ldots, N$ and $k = 1, \ldots, K$:

(a) $l(i, j, k) = k$ and $\psi_{ijk} = 0$ for $j = 0$,

(b) $l(i, j, k) \neq l(i, j', k)$ for all $j = 0, \ldots, J$ and $j' \neq j$.

**Assumption 7.** The choice-specific shocks $\varepsilon_{ik}$ are iid over time and across choices with a known joint distribution $F_{ik}$ which is absolutely continuous with respect to Lebesgue measure (with joint density $f_{ik}$), has finite first moments, and has support equal to $\mathbb{R}^{J+1}$.

Assumptions 1–7 are generalized counterparts of Assumptions 1–4 of Arcidiacono, Bayer, Blevins, and Ellickson (2016) that allow for firm heterogeneity and state dependent rates. The first part of Assumption 6 defines $j = 0$ to be the inaction choice, which does not change the state, and normalizes the corresponding instantaneous payoff to zero.[1] The second part of Assumption 6 requires actions $j > 0$ to be meaningfully distinct in the ways they change the state. This serves to rule out cases where two actions are indistinguishable.

[1] The role of the choice $j = 0$ is similar to the role of the "outside good" in models of demand. Because not all agents in the market are observed to purchase one of the goods in the model, their purchase is defined to be the outside good.

*Strategies, Best Responses, and Dynamic Payoffs* A stationary Markov policy for player $i$ is a function $\delta_i : \mathcal{K} \times \mathbb{R}^{J+1} \to \mathcal{J} : (k, \varepsilon_{ik}) \mapsto \delta_i(k, \varepsilon_{ik})$ which assigns to each state $k$ and vector $\varepsilon_{ik}$ an action from $\mathcal{J}$. For a given policy function $\delta_i$, we can define the conditional choice probabilities

$$\sigma_{ijk} = \Pr[\delta_i(k, \varepsilon_{ik}) = j \mid k] \tag{1}$$

for all choices $j$ and states $k$. Let $\varsigma_i$ denote player $i$'s beliefs regarding the actions of rival players, given by a collection of $(N-1) \times (J+1) \times K$ probabilities $\varsigma_{imjk}$ for each rival player $m \neq i$, state $k$, and choice $j$. Finally, let $V_{ik}(\varsigma_i)$ denote the expected present value for player $i$ being in state $k$ and behaving optimally at all points in the future given beliefs $\varsigma_i$. For given beliefs $\varsigma_i$, the optimal policy rule for player $i$ satisfies the following inequality condition:

$$\delta_i(k, \varepsilon_{ik}) = j \iff \psi_{ijk} + \varepsilon_{ijk} + V_{l(i,j,k)}(\varsigma_i) \geq \psi_{ij'k} + \varepsilon_{ij'k} + V_{l(i,j',k)}(\varsigma_i) \quad \forall j' \in \mathcal{J}. \tag{2}$$

That is, at each decision time $\delta_i$ assigns the action that maximizes the agent's expected future discounted payoff.

Given beliefs $\varsigma_i$ held by player $i$, we can define the value function (here, a $K$-vector) $V_i(\varsigma_i) = (V_{i1}(\varsigma_i), \ldots, V_{iK}(\varsigma_i))^\top$ where the $k$-th element $V_{ik}(\varsigma_i)$ is the present discounted value of all future payoffs obtained when starting in some state $k$ and behaving optimally in future periods given beliefs $\varsigma_i$. For small time increments $h$, under the Poisson assumption, the probability of an event with rate $\lambda_{ik}$ occurring is $\lambda_{ik}h$. Given the discount rate $\rho_i$, the discount factor for such increments is $1/(1+\rho_i h)$. Thus, for small time increments $h$ the present discounted value of being in state $k$ is (omitting the dependence on $\varsigma_i$ for brevity):

$$\begin{aligned} V_{ik} = \frac{1}{1+\rho_i h} \Bigg[ & u_{ik}h + \sum_{l \neq k} q_{0kl} h V_{il} + \sum_{m \neq i} \lambda_{mk} h \sum_{j=0}^{J} \varsigma_{imjk} V_{i,l(m,j,k)} \\ & + \lambda_{ik} h \, \mathrm{E} \max_j \left\{ \psi_{jk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \right\} + \left( 1 - \sum_{i=1}^{N} \lambda_{ik} h - \sum_{l \neq k} q_{0kl} h \right) V_{ik} + o(h) \Bigg]. \end{aligned}$$

Rearranging and letting $h \to 0$, we obtain the following recursive expression for $V_{ik}$ for beliefs $\varsigma_i$:

$$V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{0kl} V_{il} + \sum_{m \neq i} \lambda_{mk} \sum_{j=0}^{J} \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_{ik} \, \mathrm{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \}}{\rho_i + \sum_{l \neq k} q_{0kl} + \sum_{i=1}^{N} \lambda_{ik}}. \tag{3}$$

The denominator contains the sum of the discount factor and the rates of all events that might possibly change the state. The numerator is composed of the flow payoff

for being in state $k$, the rate-weighted values associated with exogenous state changes, the rate-weighted values associated with states that occur after moves by rival players, and the expected current and future value obtained when a move arrival for player $i$ occurs in state $k$. The expectation is with respect to the joint distribution of $\varepsilon_{ik} = (\varepsilon_{i0k}, \ldots, \varepsilon_{iJk})^\top$. Alternatively, and perhaps more intuitively, rearranging once again shows that the instantaneous discounted increment to the value function $V_{ik}$ for given beliefs $\varsigma_i$ is

$$\rho V_{ik} = u_{ik} + \sum_{l \neq k} q_{0kl}(V_{il} - V_{ik}) + \sum_{m \neq i} \lambda_{mk} \sum_{j=0}^{J} \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_{ik} \, \mathrm{E} \max_j \{\psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} - V_{ik}\}. \tag{4}$$

**Example 1, Continued.** In the renewal model, the value function can be expressed very simply as follows (where the $i$ subscript has been omitted since $N = 1$):

$$V_k = \frac{1}{\rho + q_1 + q_2 + \lambda} \left( u_k + q_1 V_{k+1} + q_2 V_{k+2} + \lambda \, \mathrm{E} \max \{\varepsilon_{0k} + V_k, -c + \varepsilon_{1k} + V_1\} \right).$$

**Example 2, Continued.** In the $2 \times 2$ entry model, the value function for player 1 in state $k$, where $x_k = (x_{k1}, x_{k2}, x_{k3}) \in \{0,1\} \times \{0,1\} \times \{\mathrm{L}, \mathrm{H}\}$, can be expressed recursively as

$$V_{1k} = \frac{1}{\rho_1 + \gamma_{x_{k3}} + h_{1k} + h_{2k}} \left( u_{1k} + \gamma_{x_{k3}} V_{1,l(0,x_{k3},k)} + h_{1k} V_{i,l(1,1-x_{k1},k)} + h_{2k} V_{i,l(2,1-x_{k2},k)} \right).$$

*Linear Representation of the Value Function* It will be convenient to express the Bellman equation in (3) in matrix notation, representing the Bellman operator, which we denote by $\Gamma$. Let $\Sigma_i(\sigma_i)$ denote the transition matrix implied by the choice probabilities $\sigma_i$ for player $i$ and the continuation state function $l(i, \cdot, \cdot)$. That is, the $(k, l)$ element of the matrix $\Sigma_i(\sigma_i)$ is the probability of transitioning from state $k$ to state $l$ as a result of an action by player $i$ under the given choice probabilities. For player $i$'s rivals, $\Sigma_m(\varsigma_{im})$ is the transition matrix induced by the actions of rival player $m$ according to the beliefs of player $i$ given by $\varsigma_{im}$. Let $Q_0 = (q_{0kl})$ denote the matrix of rates of exogenous state transitions and let $\tilde{Q}_0 = Q_0 - \mathrm{diag}(q_{011}, \ldots, q_{0KK})$ be the matrix formed by taking $Q_0$ and replacing the diagonal elements with zeros. Finally, define $\varsigma_{ii}$ to be the best response probabilities under the beliefs in $\varsigma_i$ about rival players.

Then, we can write the Bellman operator $\Gamma_i$ (for given beliefs $\varsigma_i$) as

$$\Gamma_i(V_i) = D_i \left[ u_i + \tilde{Q}_0 V_i + \sum_{m \neq i} L_m \Sigma_m(\varsigma_{im}) V_i + L_i \left\{ \Sigma_i(\varsigma_{ii}) V_i + C_i(\varsigma_{ii}) \right\} \right], \tag{5}$$

where $D_i = (d_{kl})$ is the $K \times K$ diagonal matrix containing the denominator of (3) for each $k$, hence $d_{kk} = 1/(\rho_i + \sum_{m=1}^{N} \lambda_{mk} + \sum_{l \neq k} q_{0kl})$, $L_m = \mathrm{diag}(\lambda_{m1}, \ldots, \lambda_{mK})$ is a diagonal matrix

containing the rates of decision times for player $m$, $C_i(\varsigma_{ii})$ is the $K \times 1$ vector containing the ex-ante expected value of the instantaneous payoff $c_{ijk} = \psi_{ijk} + \varepsilon_{ijk}$ for player $i$ in each state $k$ given the best response probabilities $\varsigma_{ii}$. That is, $k$-th element of $C_i(\varsigma_{ii})$ is $\sum_j \varsigma_{iijk} \left[\psi_{ijk} + e_{ijk}(\varsigma_{ii})\right]$, where $e_{ijk}(\varsigma_{ii})$ is the expected value of $\varepsilon_{ijk}$ given that action $j$ is chosen:

$$e_{ijk}(\varsigma_{ii}) \equiv \frac{1}{\varsigma_{iijk}} \int \varepsilon_{ijk} \cdot 1 \left\{ \varepsilon_{ij'k} - \varepsilon_{ijk} \leq \psi_{ijk} - \psi_{ij'k} + V_{i,l(i,j,k)} - V_{i,l(i,j',k)} \ \forall j' \right\} f(\varepsilon_{ik}) \, d\varepsilon_{ik}.$$

Hence, the value function $V_i$ corresponding to beliefs $\varsigma_i$ is a fixed point of $\Gamma_i$, $V_i = \Gamma(V_i)$.

A central result of Arcidiacono, Bayer, Blevins, and Ellickson (2016) showed that the differences in choice-specific value functions which appear in the definition of $e_{ijk}(\varsigma_{ii})$ above are identified as functions of the conditional choice probabilities. This result generalizes to the present setting as well:

**Lemma 1** (ABBE, 2016, Proposition 2)**.** *Under the maintained assumptions, for each player $i = 1, \ldots, N$, each state $k = 1, \ldots, K$, and each choice $j \in \mathcal{J}$ the choice-specific value function is identified up to differences with respect to some baseline choice $j' \in \mathcal{J}$:*

$$v_{ijk} \equiv (\psi_{ijk} + V_{i,l(i,j,k)}) - (\psi_{ij'k} - V_{i,l(i,j',k)}). \tag{6}$$

This is a continuous time analog of a similar result of Hotz and Miller (1993, Proposition 1). We build on this result to establish the following linear representation of the value function in terms of conditional choice probabilities, rate parameters, and payoffs. This representation generalizes Proposition 6 of ABBE and forms the basis of the identification results below and the estimator implemented in the Monte Carlo experiments. It is analogous to a similar result for discrete time games by Pesendorfer and Schmidt-Dengler (2008, eq. 6).

**Theorem 1.** *If Assumptions 1–7 hold, then for a given collection of beliefs $\sigma = (\sigma_1, \ldots, \sigma_N)$, $V_i$ has the following linear representation for each $i$:*

$$V_i(\sigma) = \Xi_i^{-1}(\sigma) \left[u_i + L_i C_i(\sigma_i)\right] \tag{7}$$

*where*

$$\Xi_i(\sigma) = \rho_i I_K + \sum_{m=1}^{N} L_m [I_K - \Sigma_m(\sigma_m)] - Q_0 \tag{8}$$

*is a nonsingular $K \times K$ matrix and $I_K$ is the $K \times K$ identity matrix.*

*Proof.* See Appendix A.

*Equilibrium* Following the literature, we focus on Markov perfect equilibria which are defined as follows.

**Definition.** A *Markov perfect equilibrium* is a collection of stationary policy rules $\{\delta_i\}_{i=1}^N$ such that (2) holds for all $i$, $k$, and $\varepsilon_{ik}$ given beliefs $\varsigma_i = (\sigma_1, \ldots, \sigma_N)$ generated by (1).

ABBE proved that such an equilibrium exists when players share common move arrival and discount rates and when the move arrival rates do not not vary across states (i.e., $\lambda_{ik} = \lambda$ and $\rho_i = \rho$ for all $i$ and $k$). The following theorem extends this to the more general model with heterogeneity.

**Theorem 2.** *If Assumptions 1–7 hold, then a Markov perfect equilibrium exists.*

*Proof.* See Appendix A.

### 2.2. *Reduced Form Markov Jump Processes Representation*

We now turn to focus on certain features of the reduced form of the model to provide some intuition. The reduced form of the models we consider is a *Markov jump process* on a finite state space $\mathcal{X} \subset \mathbb{R}^L$ with $K = |\mathcal{X}| < \infty$. Each element $x \in \mathcal{X}$ represents a possible state of the market and contains information about the market structure (e.g., which firms are active, the quality of each firm) and market conditions (e.g., demographic and geographic characteristics, input prices). The components of $x$ can be player-specific states, such as the number of stores operated by a retail chain, or exogenous market characteristics, such as population.

As in discrete time games, the players in our model can make actions and these actions influence the evolution of the market-wide state vector. These actions typically only affect certain individual components of the overall state vector. For example, when a new firm enters it may change the firm-specific activity indicator for that firm but not the population of the market. We also designate one player, indexed by $i = 0$, as "nature". This player is responsible for state changes that cannot be attributed to the action of any other player $i > 0$ (e.g., changes in population or per capita income).

The state space dynamics implied by the model can be characterized by a finite state Markov jump processes, a stochastic process $X_t$ indexed by $t \in [0, \infty)$ taking values in some finite 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 $\tau$ before transitioning to some other state $X_{t+\tau}$. The length of time $\tau$ is referred to as the holding time. A trajectory or sample path of such a process is a piecewise-constant, right-continuous function of time. This is illustrated in Figure 1, where a sample path $x_t$ for $t \in [0, \overline{T}]$ is plotted along with corresponding jump times $t_n$ and holding times $\tau_n$, with $n$ denoting the $n$-th jump.

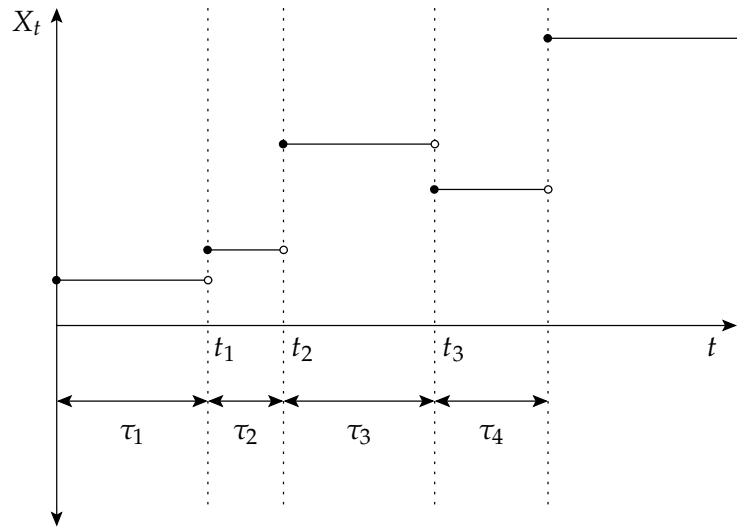


FIGURE 1. Markov jump process

A representative sample path $x_t$ for $t \in [0, \infty)$ with jump times $t_n$ and holding times $\tau_n$ shown for $n = 1, 2, 3, 4$.

Jumps occur according to a Poisson process and the holding times between jumps are therefore exponentially distributed. The mean of this distribution is $1/\lambda$, the inverse of the rate parameter, $\lambda$.

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) or Chung (1967, part II).

A finite Markov jump process can be summarized by it's *intensity matrix*, also known as the *infinitesimal generator matrix*,

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

where for $k \neq l$

$$q_{kl} = \lim_{h \to 0} \frac{\Pr\left(X_{t+h} = l \mid X_t = k\right)}{h}$$

is the probability per unit of time that the system transitions from state $k$ to state $l$ and the diagonal elements

$$q_{kk} = -\sum_{l \neq k} q_{kl}$$

are the negations of the rates at which the system leaves each state $k$. Thus, the holding times before transitions out of state $k$ follow an exponential distribution with rate parameter $-q_{kk}$. Then, conditional on leaving state $k$, the system transitions to state $l \neq k$ with probability $q_{kl} / \sum_{l \neq k} q_{kl} = -q_{kl}/q_{kk}$.

In the leading case we consider, the times at which actions are taken are not observed by the econometrician. In fact, the actions themselves are not observed either and must be inferred from observed state changes (e.g., we might say the "entry" action occurs when the number of stores a firm has changes from zero to one). With discrete time data, the state changes themselves are not directly observable either. With regularly-sampled data (e.g., annual or quarterly data) only the states at the beginning and end of each period of length $\Delta$ are observed. We cannot determine the exact sequence of actions from the observed state change over the period without additional assumptions (e.g., a firm might enter and exit in a single period, which would be undetected). However, we can make probabilistic statements about the likelihood of any particular transition occurring over an interval of given length using the *transition matrix*, which we will denote as $P(\Delta)$. This matrix will be used to relate the model to the data for estimation.

Let $P_{kl}(\Delta) = \Pr(X_{t+\Delta} = l \mid X_t = k)$ denote the probability that the system is in state $l$ after a period of length $\Delta$ given that it was initially in state $k$. The transition matrix $P(\Delta) = (P_{kl}(\Delta))$ is the corresponding $K \times K$ matrix of these probabilities. For a finite-state continuous time Markov jump processes, the Kolmogorov forward equations form a system of matrix differential equations characterizing the transition matrix $P(\Delta)$ of a process with intensity matrix $Q$ (Karlin and Taylor, 1975, 4.8):

$$P'(\Delta) = QP(\Delta), \quad P(0) = I. \tag{9}$$

It follows that the unique solution to this system is

$$P(\Delta) = \exp(\Delta Q) = \sum_{j=0}^{\infty} \frac{(\Delta Q)^j}{j!}. \tag{10}$$

The transition matrix is the matrix exponential of the intensity matrix $Q$ scaled by $\Delta$. This is the matrix analog of the scalar exponential $\exp(x)$ for $x \in \mathbb{R}$.[2]

Finally, we review some properties of the exponential distribution which will be required for constructing the value functions in the dynamic games considered below. In particular, we note that if there are $N$ competing Poisson processes (or exponential distributions) with rates $\lambda_i$ for $i = 1, \ldots, N$, then distribution of the minimum wait time

[2] Although we cannot calculate the infinite sum (10) exactly, we can compute $\exp(\Delta Q)$ numerically using known algorithms implemented, for example, in the Fortran package Expokit (Sidje, 1998) or the `expm` command in Matlab.

is exponential with rate parameter $\sum_{i=1}^{N} \lambda_i$. Furthermore, conditional on an arrival the probability that it is due to process $i$ is $\lambda_i / \sum_{j=1}^{N} \lambda_j$. These properties are well known, but we present the following lemma for completeness.

**Lemma 2.** *Suppose $\tau_i \sim \text{Expo}(\lambda_i)$, for $i = 1, \ldots, N$, are independent and define $\tau \equiv \min_i \tau_i$ and $\iota \equiv \arg\min_i \tau_i$. Then*

$$\tau \sim \text{Expo}(\lambda_1 + \cdots + \lambda_N) \quad \textit{and} \quad \Pr(\iota = i) = \frac{\lambda_i}{\sum_{j=1}^{N} \lambda_j}.$$

*Proof.* See Appendix A.

This lemma allows us to treat the $N$ competing Poisson processes $(\tau_1, \ldots, \tau_N)$ as a single joint process $(\tau, \iota)$ where the joint distribution is given above. In the context of the dynamic games we consider, we can apply this proposition to decompose the state transition process at two levels. Consider a game currently in state $k$ with $N$ players. The first branching occurs at the player level, to determine which player moves next. If $N$ players move at rates $\lambda_{1k}, \ldots, \lambda_{Nk}$ in state $k$, then the overall rate of some player moving is $\sum_{i=1}^{N} \lambda_{ik}$ and the probability that player $i$ moves next is $\lambda_{ik} / \sum_{i=1}^{N} \lambda_{ik}$. The second level of branching occurs at the action level: conditional on a particular player moving, which action is chosen? Suppose player $i$ plays each action $j \in \{0, 1, \ldots, J\}$ in state $k$ at rate $h_{ijk}$. The sum of the individual action rates for player $i$ in state $k$, $\sum_{j=1}^{J} h_{ijk}$, must equal the overall rate of actions in state $k$, $\lambda_{ik}$. Conditional on player $i$ making some action, the probability that action $j$ is made is $h_{ijk} / \lambda_{ik}$.

In the notation of the lemma, we can let $\iota$ denote a player-action pair $(i, j)$ and $\tau$ denote the holding time until the next action by any player occurs in state $k$. The distribution of $\tau$ is exponential with rate parameter $\sum_{i=1}^{N} \lambda_{ik} = \sum_{i=1}^{N} \sum_{j=0}^{J} h_{ijk}$ and the probability that $\iota = (i, j)$ is $h_{ijk} / \left( \sum_{i=1}^{N} \sum_{j=0}^{J} h_{ijk} \right)$.

Now, in the context of the dynamic games we consider, the state space dynamics can be fully characterized by a collection of $N + 1$ competing Markov jump processes with intensity matrices $Q_0, Q_1, \ldots, Q_N$. Each process corresponds to some player $i$ and $i = 0$ denotes nature, to which we will attribute exogenous state changes. The aggregate intensity matrix, $Q$, is simply the sum of the player-specific intensity matrices: $Q = Q_0 + Q_1 + \ldots, Q_N$.

**Example 1, Continued.** Consider the $Q$ matrix implied by the continuous-time single-agent renewal model. The state variable in the model is the total accumulated mileage of a bus engine, $\mathcal{K} = \{1, \ldots, K\}$. The exogenous state transition process is characterized by a

$K \times K$ intensity matrix $Q_0$ on $\mathcal{K}$ with two parameters $\gamma_1$ and $\gamma_2$:

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

Let $\sigma_{1k}$ denote the probability of replacement in state $k$. We will discuss how these probabilities are determined in detail below, but for now suppose they are given. The intensity matrix for state changes induced by the the agent is

$$Q_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & \cdots & 0 \\ \lambda\sigma_{12} & -\lambda\sigma_{12} & 0 & 0 & \cdots & 0 \\ \lambda\sigma_{13} & 0 & -\lambda\sigma_{13} & 0 & \cdots & 0 \\ \lambda\sigma_{14} & 0 & 0 & -\lambda\sigma_{14} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \lambda\sigma_{1K} & 0 & 0 & 0 & \cdots & -\lambda\sigma_{1K} \end{bmatrix}.$$

The aggregate intensity matrix is $Q = Q_0 + Q_1$, which fully characterizes the reduced form state space dynamics.

A representative sample path generated by this model is shown in Figure 2. Inter-arrival times are indicated by $\tau_{in}$, where $i$ denotes the identity of the player (with $i = 0$ denoting nature) and $n$ denotes the event number. The agent's decisions ($a_{t_n}$) are indicated at each decision time. For example, at time $t_1$, the agent chooses to continue without replacement ($a_{t_1} = 0$), while at time $t_4$, the agent chooses to replace ($a_{t_4} = 1$), resetting the mileage.

**Example 2, Continued.** Let $h_{ik}$ be the hazard of firm $i$ switching from active to inactive or vice versa in state $k$. Let $\gamma_{\mathrm{L}}$ and $\gamma_{\mathrm{H}}$ be the rates at which nature switches between demand states (i.e., demand moves from low to high at rate $\gamma_{\mathrm{L}}$). The aggregate state space dynamics are illustrated in Figure 3.

The state transition hazards can be characterized by an $8 \times 8$ intensity matrix $Q$. Note that firms cannot change the demand state, firms cannot change each other's states, and nature cannot change the firms' states. Therefore, the overall intensity matrix has the form

$$Q = \begin{bmatrix} Q^{\mathrm{LL}} & Q^{\mathrm{LH}} \\ Q^{\mathrm{HL}} & Q^{\mathrm{HH}} \end{bmatrix} = \begin{bmatrix} Q_1^L + Q_2^L & Q_0^L \\ Q_0^H & Q_1^H + Q_2^H \end{bmatrix}$$

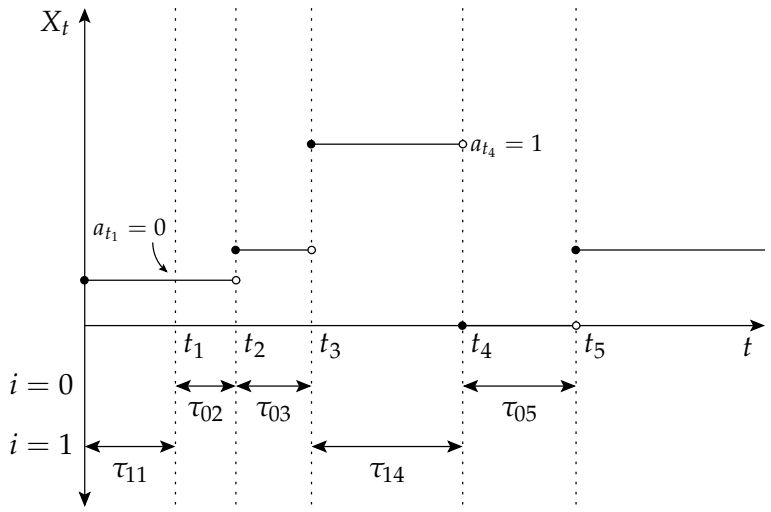


FIGURE 2. Single agent model: a representative sample path where $t_n$, $\tau_{in}$, and $a_{in}$ denote, respectively, the time, inter-arrival time, and action corresponding to $n$-th event. Moves by the agent are denoted by $i = 1$ while $i = 0$ denotes a state change (a move by nature).

The low demand state L corresponds to encoded sates $k = 1, \ldots, 4$. In this portion of the state space, firms change the state as follows (diagonal elements are omitted for simplicity):

$$Q_1^L = \begin{bmatrix} \cdot & h_{11} & 0 & 0 \\ h_{12} & \cdot & 0 & 0 \\ 0 & 0 & \cdot & h_{13} \\ 0 & 0 & h_{14} & \cdot \end{bmatrix}, \quad Q_2^L = \begin{bmatrix} \cdot & 0 & h_{21} & 0 \\ 0 & \cdot & 0 & h_{22} \\ h_{23} & 0 & \cdot & 0 \\ 0 & h_{24} & 0 & \cdot \end{bmatrix}$$

Importantly, the locations of the nonzero off-diagonal elements are distinct because the state-to-state communication patterns differ. A similar structure arises for the high demand state H, for $k = 5, 6, 7, 8$. Therefore, given $Q$ we can immediately determine $Q_0$, $Q_1$, and $Q_2$.

### *2.3. Normalizations with Respect to Alternative Clock Processes*

Consider a given finite Markov jump process with intensity matrix $Q = (q_{kl})$. The off-diagonal components of the matrix $Q$ are rates of state changes between distinct states. For a given state $k$, this is a model strictly for transitions that involve *leaving state k*. Such transitions occur at rate $\sum_{l \neq k} q_{kl} = -q_{kk}$. This is essentially a competing hazards model where each hazard is a transition to some other state $l$ with an associated rate $q_{kl}$. Note that when a hitting time occurs for one of these processes, by definition the process must

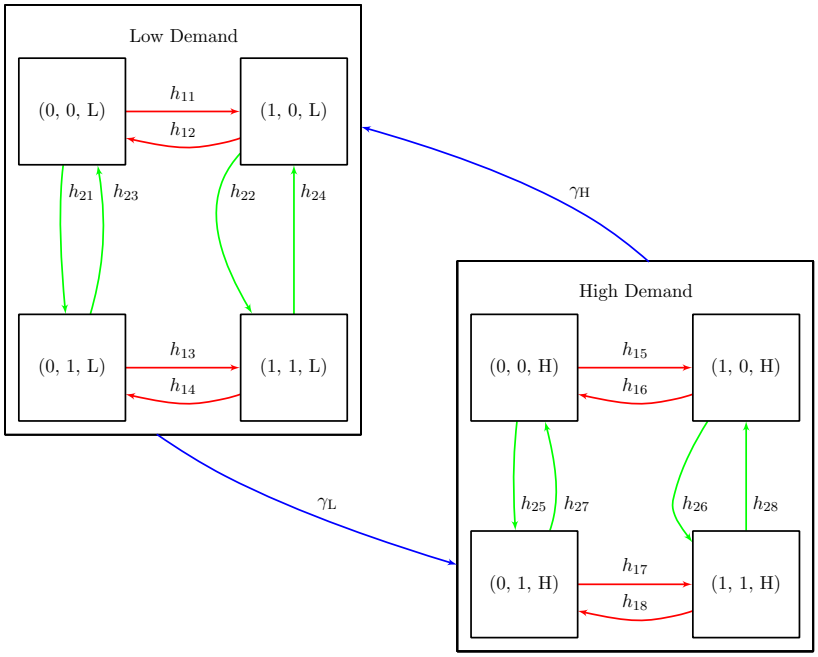


FIGURE 3. Two Player Entry Game with Exogenous Demand State

leave state $k$. The result is a branching process where at rate $-q_{kk}$ an event occurs and the probability that the event is a transition to state $l$ is $q_{kl}/(-q_{kk})$, which is proportional to $q_{kl}$.

We can characterize the same process with an alternative parametrization that will prove more useful in the context of dynamic discrete choice models. Let the rates $(\lambda_1, \ldots, \lambda_K)$ be given, where $\lambda_k$ is the rate for a Poisson "clock process" that governs *potential* transitions out of state $k$. For each $k$, we must have $\lambda_k \geq -q_{kk} = -\sum_{l \neq k} q_{kl}$ since the rate of potential transitions must be at least as large as the rate of *actual* transitions. At each hitting time under this process, the state may change to a state $l \neq k$ or it may remain in state $k$. In other words, this parametrization allows the process to remain at state $k$ with some probability when an event occurs. The overall rate $\lambda_k$ can be normalized to any value larger than the rate of actual state changes, $-q_{kk}$.

Specifically, let $\sigma_{kl} = q_{kl}/\lambda_k$ be the probability that the event is a transition from state $k$ to state $l \neq k$ and let $\sigma_{kk} = (\lambda_k - \sum_{l \neq k} q_{kl})/\lambda_k$ be the probability that the event is not a transition (i.e., the process remains at state $k$). To translate these probabilities into hazards, we simply multiply by the overall rate: $\tilde{q}_{kl} = \lambda_k \sigma_{kl}$ is the rate of transitions from state $k$ to $l \neq k$. Therefore, the diagonal elements—equal to the negatives of the rates of leaving

each state—are $\tilde{q}_{kk} = -\lambda_k(1 - \sigma_{kk})$. But the resulting intensity matrix $\tilde{Q} = (\tilde{q}_{kl})$ equals the original intensity matrix $Q = (q_{kl})$. To see this, note that $\tilde{q}_{kl} = \lambda_k \sigma_{kl} = \lambda_k(q_{kl}/\lambda_k) = q_{kl}$ for all $l \neq k$ and $\tilde{q}_{kk} = -\lambda_k + \lambda_k \sigma_{kk} = -\sum_{l \neq k} q_{kl} = q_{kk}$.

*2.4. Comparison with Discrete Time Models*

We conclude this section with a brief comparison of continuous time and discrete time models in terms of their reduced form implications. There are also related identification issues which we will return to in the next section. To be concrete, suppose that the unit of time is one year and that data are available on an annual basis (i.e., $\Delta = 1$).

First, consider a typical discrete time in which agents make decisions once per period of length $\delta = 1$ (i.e., annually). In an entry/exit setting where the choice set is $\mathcal{J} = \{0, 1\}$, this implies that there can be *at most* one entry or exit per year. In a chain store setting the choice variable is the net number of stores to open during the year, the choice set is $\mathcal{J} = \{-J, \ldots, J\}$. This implies that there an be at most $J$ openings or closings per year. Hence, $J$ must be chosen accordingly by the researcher to be the maximum number of possible stores opened or closed by any chain firm in any single year.

Now, consider a continuous time model with a common move arrival rate $\overline{\lambda}$ for all players and all states. In the entry/exit setting, the choice set is still $\mathcal{J} = \{0, 1\}$ which implies that there are *on average* $1/\overline{\lambda}$ entries or exits per year. In the chain store setting, there are two possible modeling choices in continuous time. The first, and perhaps most natural model would specify the choice set to be $\mathcal{J} = \{-1, 0, 1\}$ which implies that *on average* there are at most $1/\overline{\lambda}$ openings or closings per year. The choice set represents the set of possible *instantaneous* state changes, so the implicit assumption is that no more than one store is ever opened or closed simultaneously. The researcher should choose $\overline{\lambda}$ so that this average is not binding. This is analogous to choosing $J$ in the discrete time model so that the choice set is not binding. However, there is an additional degree of freedom in the continuous time setting because the choice set can be extended. The second approach would be to specify $\mathcal{J} = \{-J, \ldots, J\}$ for some $J$ which implies on average at most $J/\overline{\lambda}$ openings or closings per year.

In both cases, there is a trade-off between the time period chosen (which is usually fixed at one in discrete time) and the richness of the choice set. However, the continuous time model has an additional degree of freedom relative to the discrete time model. To allow for more frequent openings and closings, for example, one can either adjust the choice set (i.e., increase $J$), or increase $\overline{\lambda}$ to increase the potential frequency. The continuous time is more forgiving of misspecifications of the choice set because the chosen rate $\overline{\lambda}$ determines the *average* number of events rather than placing a hard limit on the realized number of events, which is the case in discrete time.

## 3. Identification Analysis

Our identification analysis proceeds in two steps corresponding to the reduced form and the structural model. If we view deriving the implications of the model as a forward or "bottom-up" problem, then identification is an inverse or "top-down" problem. Our arguments follow the same natural progression.

Deriving the implications of the structural model can be viewed as a bottom-up exercise: the structural primitives $u$ and $\psi$ imply value functions $V$ which imply choice probabilities $\sigma$. These probabilities along with the rates of moves, $\lambda$, and state transitions by nature, $Q_0$, in turn imply an intensity matrix $Q$. Finally, given the $Q$ matrix and a process for sampling data, this implies a data generating process. For example, for a fixed sampling interval $\Delta$ the distribution of observable data is $P(\Delta) = \exp(\Delta Q)$.

On the other hand, the identification problem requires us to consider the inverse problem, working from the top down. First, given knowledge of the data generating process, the transition matrix $P(\Delta)$ for a fixed interval $\Delta$ for example, we derive conditions under which we can uniquely determine the reduced form intensity matrix $Q$. If the complete continuous time record is potentially observable, then $Q$ is trivially identified. Otherwise, if the state of the game is potentially observable only at times that do not necessarily correspond to state changes, we show that $Q$ can still be identified. In the leading case of discrete time data sampled at a regular intervals, we show that this is possible under very mild conditions on the structure of the game. We also briefly consider other possibilities such as shrinking and random sampling intervals.

Second, with $Q$ in hand we turn to identification of the structural primitives of the model, namely the flow payoffs $u$ and instantaneous payoffs $\psi$. We show that knowledge of $Q$ allows us to recover these structural primitives with fewer identifying restrictions than required in discrete time models. This is due to the absence of simultaneous moves at any given instant, which is also the source of the computational efficiency of the model.

### *3.1. Identification of Q*

With continuous-time data, identification and estimation of the intensity matrix for finite-state Markov jump processes is straightforward and well-established (Billingsley, 1961). However, when a continuous-time process is only sampled at discrete points in time, the parameters of the underlying continuous-time model may not be point identified.[3] In the

[3]This is known as the *aliasing problem* and it has been studied extensively in the context of continuous-time systems of stochastic differential equations (Sims, 1971; Phillips, 1973; Hansen and Sargent, 1983; Geweke, 1978; Kessler and Rahbek, 2004; McCrorie, 2003; Blevins, 2017). See Figure 1 of Blevins (2017) for an illustration in the frequency domain, where the problem is perhaps most obvious.

present model, the concern is that there may be multiple $Q$ matrices which give rise to the same data generating process, which is the potentially observable transition probability matrix $P(\Delta)$ in the leading case of fixed sampling intervals.[4]

Before proceeding, we note that there is a similar question of identification in discrete time models that is hidden by the usual assumption that the unknown frequency of moves is equal to the (known) sampling frequency (Hong, Li, and Wang, 2015). If instead agents move at a discrete interval of length $\delta$ that differs from the sampling interval $\Delta$, then there are in general multiple transition matrices $P_0$ such that $P_0^{\Delta/\delta} = P(\Delta)$ (Gantmacher, 1959; Singer and Spilerman, 1976).

To illustrate this issue in the continuous time setting, Figure 4 displays two distinct paths which coincide both before and after an interval of length $\Delta$, but which take different intermediate steps. Consider the possible paths of the process between times $t_2 - \Delta$ and $t_2$. The dashed path first moves to a higher state before arriving at the resulting state $k_{t_2}$, while the dashed and dotted path first moves to a lower state and arrives in $k_{t_2}$ at a later time (but before $t_2$). There are an infinite number of such paths since time is continuous, but the dynamics of the process over the interval are summarized by the transition matrix $P(\Delta)$.

Much of the previous work on this identification problem seeks conditions on the observable discrete-time transition matrix $P(\Delta)$. We briefly review some of these results in the next subsection, but our approach is to show that one can instead identify $Q$ via identifying restrictions on the primitives of the underlying structural model and that such restrictions easily arise from the statement of the model itself. These can be viewed as exclusion restrictions.

For example, in applications there are typically player-specific components of the state vector where player $i$ is not permitted to change the players-specific state of player $j$ and vice-versa. In an entry-exit model, such a state is incumbency status: players can enter and exit by their own action, but no player can enter or exit on behalf of another player. Similarly, if the overall state vector has components that are exogenous state variables, such as population, then we know that any state changes involving those variables must

[4] A related issue is the embeddability problem: could the transition matrix $P(\Delta)$ have been generated by a continuous-time Markov jump process for some intensity matrix $Q$ or some discrete-time chain over fixed time periods of length $\delta$? This is a model specification issue, also arising in both discrete time and continuous time: could the data have been generated by the structural model (here, a continuous-time Markov jump process)? We assume throughout that the model is well-specified and therefore, such an intensity matrix $Q$ exists. This problem was first proposed by Elfving (1937). Kingman (1962) derived the set of embeddable processes with $K = 2$ and Johansen (1974) gave an explicit description of the set for $K = 3$. Singer and Spilerman (1976) summarize several known necessary conditions for embeddability involving testable conditions on the determinant and eigenvalues of $P(\Delta)$. It might be possible to use such conditions to construct model specification tests.

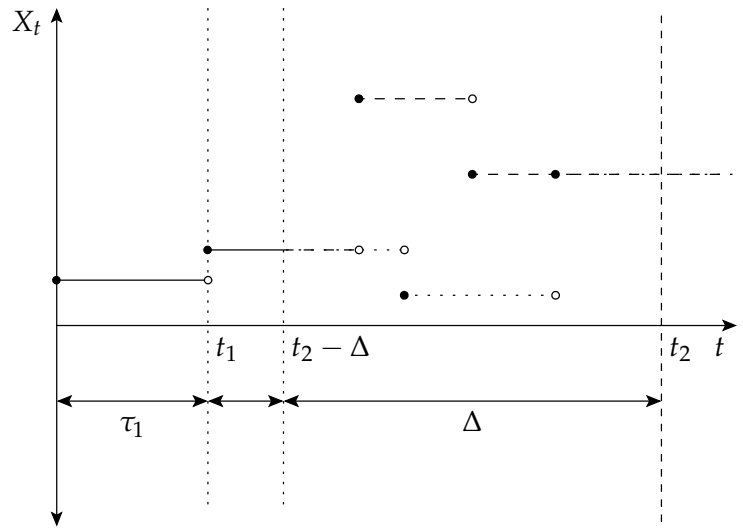


FIGURE 4. Time aggregation: Two distinct paths which end in the same state at $t_2$ and begin in the same state at $t_2 - \Delta$ and but differ over intermediate interval of length $\Delta$.

be due to an "action" by nature and not by an action of any other player. This natural structure implies many linear restrictions on the $Q$ matrix. We show that restrictions of this form limit the domain of the mapping $Q \mapsto \exp(\Delta Q) = P(\Delta)$ in such a way as to guarantee an almost surely unique intensity matrix $Q$ for any given discrete time transition matrix $P(\Delta)$.

### *3.1.1. Identification of Unrestricted Q Matrices*

Returning to the general problem of identification of $Q$, recall that the question is whether there exists a unique matrix $Q$ that leads to the observed transition matrix $P(\Delta) = \exp(\Delta Q)$ when the process is sampled at uniform intervals of length $\Delta$. The question amounts to determining when the matrix logarithm $\ln P(\Delta)$ is unique, in which case $Q = \Delta^{-1} \ln P(\Delta)$. In general, the matrix logarithm is not unique (see Gantmacher, 1959; Singer and Spilerman, 1976).

Previous mathematical treatments have tended to view the relationship $\exp(\Delta Q) = P(\Delta)$ from the perspective of the transition matrix $P(\Delta)$. In such cases there is not an underlying model that generates $Q$, so $Q$ is the model primitive of interest and is unrestricted (aside from requirement that it must be a valid intensity matrix). As a result, most previous work on the aliasing problem focused on finding sufficient conditions on the matrix $P(\Delta)$ (rather than $Q$) to guarantee that $\ln P(\Delta)$ is unique. For example, if the eigenvalues of $P(\Delta)$ are distinct, real, and positive, then $Q$ is identified (Culver,

1966). More generally, Culver (1966) proved that $Q$ is identified if the eigenvalues of $P(\Delta)$ are positive and no elementary divisor (Jordan block) of $P(\Delta)$ belonging to any eigenvalue appears more than once. Other sufficient conditions for identification of $Q$ include $\min_k\{P_{kk}(\Delta)\} > 1/2$ (Cuthbert, 1972) and $\det P(\Delta) > e^{-\pi}$ (Cuthbert, 1973). See Singer and Spilerman (1976) for a summary of these results and others.

Other sufficient conditions for identification of $Q$ involve alternative sampling schemes. For example, $Q$ can always be identified for some sufficiently small sampling interval $\Delta$ (Cuthbert, 1973; Singer and Spilerman, 1976; Hansen and Sargent, 1983). A useful result for experimental studies is that $Q$ is identified if the process is sampled at two distinct intervals $\Delta_1$ and $\Delta_2$ where $\Delta_2 \neq k\Delta_1$ for any integer $k$ (Singer and Spilerman, 1976, 5.1).

The first type of conditions above—restrictions on $P(\Delta)$—are based on a "top down" approach and are less desirable in cases where there is an underlying economic model that generates the $Q$ matrix. The second type of restrictions are based on changing how the continuous time process is sampled, which is not possible in observational studies for which the data have already been collected at regular intervals. Instead, we take a "bottom up" approach which allows economic theory to inform our identification conditions via restrictions on $Q$ that guarantee uniqueness of $\ln P(\Delta)$. For applied economists, more compelling conditions are likely to involve cross-row and cross-column restrictions on the $Q$ matrix and the locations of known zeros of the $Q$ matrix. As we discuss below, such restrictions arise naturally once the collection of firms, actions, and the resulting state transitions are defined.

### *3.1.2. Structural Restrictions for Identification of Q*

The problem of identifying continuous time models with only discrete time data has also appeared previously in the econometrics literature, in work by Phillips (1973) on continuous time regression models. He considered multivariate, continuous-time, time-homogeneous regression models of the form $y'(t) = Ay(t) + \xi(t)$, where $y(t)$ is an $n \times 1$ vector and $A$ is an $n \times n$ structural matrix. He discusses the role of prior information on the matrix $A$ and how it can lead to identification. He showed that $A$ is identified given only discrete time observations on $y$ if $A$ satisfies certain rank conditions.

Our proposed identification strategy is inspired by this work on multivariate regression models, but our model is different because the $Q$ matrix is known to be an intensity matrix (rather than an arbitrary matrix of regression coefficients) and has a rather sparse structure which is dictated by an underlying structural model. Yet, there are a number of similarities: the present model can also be characterized by a system of differential equations as in (9), where the intensity matrix $Q$ plays a role similar to the matrix $A$ above. If $Q$ is an valid intensity matrix, then the functions $P(\Delta)$ which solve this system are the transition

matrices of continuous-time stationary Markov chains (Chung, 1967, p. 251–257).

The structural model restricts $Q$ to a lower-dimensional subspace since it is sparse and must satisfy both within-row and across-row restrictions, and given the results above it seems likely that these restrictions could lead to identification of $Q$. That is, even if there are multiple matrix solutions to the equation $P(\Delta) = \exp(\Delta Q)$, it is unlikely that two of them simultaneously satisfy the restrictions of the structural model. We return to the two examples introduced previously to illustrate this idea.

**Example 1, Continued.** In the single-agent renewal model the aggregate intensity matrix is

$$Q = \begin{bmatrix} -q_1 - q_2 & q_1 & q_2 & 0 & 0 \\ h_{12} & -q_1 - q_2 - h_{12} & q_1 & q_2 & 0 \\ h_{13} & 0 & -q_1 - q_2 - h_{13} & q_1 & q_2 \\ h_{14} & 0 & 0 & -q_1 - q_2 - h_{14} & q_1 + q_2 \\ h_{15} & 0 & 0 & 0 & -h_{15} \end{bmatrix}, \tag{11}$$

where we have dropped the $i$ subscript on the hazards for simplicity. The number of parameters to be estimated in this matrix is substantially less than if each of the intensities were allowed to vary freely. Of the 20 non-trivial state-to-state transitions, only 11 are permitted: seven due to nature and four by action of the player. The remaining nine transitions are not possible in a single step. Nature cannot decrease mileage and can only increase it by one or two states at once. The agent can only reset mileage to the initial state. This results in nine known zeros of the aggregate $Q$ matrix. As we show below, these restrictions are sufficient to identify $Q$. Note that given $Q$, we can separately determine both $Q_0$ and $Q_1$. Additionally, the hazards $h_{1k}$ are the products of the overall move arrival rates and the conditional choice probabilities, which introduces shape restrictions on $h_{1k}$ across states $k$.

**Example 2, Continued.** In the $2 \times 2 \times 2$ entry example, the aggregate intensity matrix is $Q = Q_0 + Q_1 + Q_2$:

$$Q = \left[\begin{array}{cccc|cccc} \cdot & h_{11} & h_{21} & 0 & \gamma_L & 0 & 0 & 0 \\ h_{12} & \cdot & 0 & h_{22} & 0 & \gamma_L & 0 & 0 \\ h_{23} & 0 & \cdot & h_{13} & 0 & 0 & \gamma_L & 0 \\ 0 & h_{24} & h_{14} & \cdot & 0 & 0 & 0 & \gamma_L \\ \hline \gamma_H & 0 & 0 & 0 & \cdot & h_{15} & h_{25} & 0 \\ 0 & \gamma_H & 0 & 0 & h_{16} & \cdot & 0 & h_{26} \\ 0 & 0 & \gamma_H & 0 & h_{27} & 0 & \cdot & h_{17} \\ 0 & 0 & 0 & \gamma_H & 0 & h_{28} & h_{18} & \cdot \end{array}\right], \tag{12}$$

where the diagonal elements have been omitted for simplicity. Note that some transitions cannot happen at all, such as $(0, 1, \mathrm{L})$ to $(1, 0, \mathrm{L})$. The remaining transitions can happen only due to the direct action of one of the firms, but not the other. For example, moving from $(0, 0, \mathrm{H})$ to $(1, 0, \mathrm{H})$ is only possible if firm 1 chooses to become active. From any state, the set of other states to which either firm can move the state as a result of an action is limited naturally by the model and the definition of the state space. This structure yields intensity matrices that are rather sparse, which in turn makes identification of $Q$ more likely even with time aggregation since any observationally equivalent $Q$ matrix must have the same structure. Finally, given $Q$ we can again separately recover $Q_0$, $Q_1$, and $Q_2$.

Similar sparse structures arise in even models with large numbers of firms and millions of states, as in the application of Arcidiacono, Bayer, Blevins, and Ellickson (2016). In light of this lower-dimensional structure, we build on the results of Blevins (2017) who gave sufficient conditions for identification of the intensity matrix $Q$ of a general finite state Markov jump processes. These conditions were based on structural restrictions on the matrix $Q$ of the general linear form $R\,\mathrm{vec}(Q) = r$. For the $K \times K$ matrix $Q = (q_{kl})$, $\mathrm{vec}(Q)$ is the vector obtained by stacking the columns of $Q$: $\mathrm{vec}(Q) = (q_{11}, q_{21}, \ldots, q_{K1}, \ldots, q_{1K}, \ldots, q_{KK})^\top$.

These restrictions will serve to rule out alternative $Q$ matrices. Gantmacher (1959) showed that all solutions $\tilde{Q}$ to $\exp(\Delta\tilde{Q}) = P(\Delta)$ have the form

$$\tilde{Q} = Q + UDU^{-1}$$

where $U$ is a matrix whose columns are the eigenvectors of $Q$ and $D$ is a diagonal matrix containing differences in the complex eigenvalues of $Q$ and $\tilde{Q}$. This means that both the eigenvectors $U$ and the real eigenvalues of $Q$ are identified. Any other such matrices $\tilde{Q}$ must also satisfy the prior restrictions, so $R\,\mathrm{vec}(\tilde{Q}) = r$. By the relationship between $Q$ and $\tilde{Q}$ above, we also have $R\,\mathrm{vec}(Q + UDU^{-1}) = r$. But $R\,\mathrm{vec}(Q) = r$ and by linearity of the vectorization operator, $R\,\mathrm{vec}(UDU^{-1}) = 0$. An equivalent representation is

$$R(U^{-\top} \otimes U)\,\mathrm{vec}(D) = 0.$$

Here, adapting Theorem 1 of Blevins (2017) to the special case of finite-state Markov jump processes, when there are at least $\left\lfloor \frac{K-1}{2} \right\rfloor$ linear restrictions and $R$ has full rank, then $D$ must be generically zero and therefore the eigenvalues of $\tilde{Q}$ and $Q$ are equal. If the eigenvectors and all eigenvalues of $\tilde{Q}$ are the same as those of $Q$, the matrices must be equal and therefore $Q$ is identified.

The following theorem establishes that there are sufficiently many restrictions of full rank to identify $Q$ in a broad class of continuous time games. This theorem includes exogenous market-specific state variables and shows that such states increase the number

of zero restrictions and make identification of $Q$ more likely, as do player-specific state variables.

**Theorem 3.** *Suppose the state vector is $x = (x_0, x_1, \ldots, x_N) \in \mathcal{X}_0 \times \mathcal{X}_1 \times \cdots \times \mathcal{X}_N$ where the component $x_0 \in \mathcal{X}_0$ is an exogenous market characteristic taking $|\mathcal{X}_0| = K_0$ values and for each $i = 1, \ldots, N$ the component $x_i$ is a player-specific state affected only by the action of each player with $|\mathcal{X}_i| = K_i$ possible distinct values. If $Q$ has distinct eigenvalues that do not differ by an integer multiple of $2\pi i/\Delta$, then $Q$ is generically identified when*

$$K_0 \prod_{i=1}^{N} K_i - K_0 - \sum_{i=1}^{N} J_i + \frac{1}{2} \geq 0. \tag{13}$$

*The quantity on the left hand side is the number of overidentifying restrictions and these are strictly increasing in $K_i$ for $i = 1, \ldots, N$, strictly increasing in $K_0$ when $K_i > 1$ for any $i = 1, \ldots, N$, and strictly decreasing in $J_i$ for $i = 1, \ldots, N$.*

*Proof.* See Appendix A.

Generic identification here means that $Q$ is identified with the exception of a measure zero set of population $Q$ matrices (see Phillips, 1973; Blevins, 2017). Sparsity of $Q$ aids in identification and is increasing in both the number of exogenous states $K_0$ and player-specific states $K_i$ but decreasing in the number of choices $J_i$. Therefore, for identification we need either a sufficiently large number of exogenous or player-specific states or a sufficiently small number of choices. In typical applications $J_i$ is small relative to $K_i$ and $K_0$.

One result of the theorem is that any binary choice game with meaningful player-specific states ($N > 1$ with $J_i = 1$ and $K_i > 1$ for all $i$) is identified, regardless of the number of players or exogenous market states. The sufficient condition in this case simplifies to $K_0(\prod_i K_i - 1) \geq N - \frac{1}{2}$. When $K_0 \geq 1$ and $K_i \geq 2$ we have $K_0(\prod_i K_i - 1) \geq 2^N - 1$ which exceeds $N - \frac{1}{2}$ for integers $N > 1$.

### *3.1.3. Identification of $Q_i$*

For identification purposes, we make the following assumption which requires that given the aggregate intensity matrix $Q$, we can determine the player-specific intensity matrices $Q_i$ for $i = 0, \ldots, N$.

**Assumption 8.** The mapping $Q \to \{Q_0, Q_1, \ldots, Q_N\}$ is known.

This assumption is satisfied in most applications, where firms cannot change each other's state variables and where actions by nature can be distinguished from the actions of firms. Note also that the diagonal elements are unimportant: if the off-diagonal elements

of each $Q_i$ can be identified from $Q$, then diagonal elements are equal to the negative of the sum of the off-diagonal elements. This assumption can be verified by inspection of $Q$ Examples 1 and 2. For example, in the single-agent renewal example $Q$ is given in (11) and for a two-player entry model, $Q$ is given in (12). A sufficient condition in general is that the continuation states resulting from actions of different players are distinct: for all players $i$ and $m \neq i$ and all states $k$,

$$\{l(i,j,k) : j = 1,\ldots,J\} \cap \{l(m,j,k) : j = 1,\ldots,J\} = \varnothing.$$

*3.2. Identification of the Value Functions and Payoffs*

We now establish that the value functions, instantaneous payoffs, and utility functions are identified. Let $V_i = (V_{i1},\ldots,V_{iK})^\top$ denote the $K$-vector of valuations for player $i$ in each state. Let $\psi_{ij} = (\psi_{ij1},\ldots,\psi_{ijK})^\top$ denote the $K$-vector of instantaneous payoffs for player $i$ making choice $j$ in each state and let $\psi_i = (\psi_{i1}^\top,\ldots,\psi_{iJ}^\top)^\top$. Given an appropriate collection of linear restrictions on these quantities, we show below that they are identified.

First, we consider a more general setting where there are variables that may shift the overall rate of moves but not the payoffs. Suppose we can partition the state variable into two components $(x,z)$, where $x$ are the payoff-relevant variables and $z$ are variables that affect the rate of moves. Let $h_{ij}(x,z)$ be the hazard of player $i$ choosing action $j$ in state $(x,z)$.

Importantly, we note that when $j = 0$ is a latent or unobserved continuation action, it is not possible to identify the rates $h_{i0}(x,z)$ even with continuous time data. Hence, the overall rates of decisions $\lambda_i(x,z) = \sum_j h_{ij}(x,z)$ are also unidentified. We will consider here whether the structural model or functional form restrictions provide enough information to identify these rates.

For ease of exposition, we consider the case where where the choice-specific errors have a type 1 extreme value distribution. Noting that $h_{ij}(x,z) = \lambda_i(x,z)\sigma_{ij}(x)$ and invoking Lemma 1, in this case differences in log hazards can be written as

$$\ln h_{ij}(x,z) - \ln h_{i0}(x,z) = \ln \sigma_{ij}(x) - \ln \sigma_{i0}(x) = \psi_{ij}(x) - \psi_{i0}(x) + V_i(l(i,j,k)) - V_i(x).$$

Rearranging, we have

$$\ln h_{ij}(x,z) = [\ln h_{i0}(x,z) - \psi_{i0}(x)] + \psi_{ij}(x) + V_i(l(i,j,k)) - V_i(x).$$

The hazards on the left hand size for $j > 0$ are identified from $Q$, while the quantities on the right hand size are unknown.

Suppose the state space is $\mathcal{X} \times \mathcal{Z}$ with $|\mathcal{X}| = K$, $|\mathcal{Z}| = L$, and $J_i = J$ for all $i$. Then we have a linear system of equations with $NJKL$ identified hazards, $NKL$ unknown

hazards, $N(J+1)K$ unknown instantaneous payoff, and $NK$ unknown valuations. The total number of unknowns is $NKL + N(J+1)K + NK = NK(J+L+2)$. So, the order condition necessary for full identification is $NJKL \geq NK(J+L+2)$ or $JL \geq J+L+2$.

Before proceeding, we define $S_{ij}$ to be the state transition matrix induced by the continuation state function $l(i,j,\cdot)$. In other words, $S_{ij}$ is a permutation matrix where the $(k,l)$ element is 1 if playing action $j$ in state $k$ results in a transition to state $l$ and 0 otherwise. Let $I_K$ denote the $K \times K$ identity matrix.

$$\left[\begin{array}{c} \ln h_{i1}^1 \\ \vdots \\ \ln h_{iJ}^1 \\ \hline \vdots \\ \hline \ln h_{i1}^L \\ \vdots \\ \ln h_{iJ}^L \end{array}\right] = \left[\begin{array}{cccc|cccc|c} I_K & \dots & 0 & & -I_K & I_K & 0 & \dots & 0 & S_{i1} - I_K \\ I_K & \dots & 0 & & -I_K & 0 & I_K & \dots & 0 & S_{i2} - I_K \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ I_K & \dots & 0 & & -I_K & 0 & 0 & \dots & I_K & S_{iJ} - I_K \\ \hline \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \hline 0 & \dots & I_K & & -I_K & I_K & 0 & \dots & 0 & S_{i1} - I_K \\ 0 & \dots & I_K & & -I_K & 0 & I_K & \dots & 0 & S_{i2} - I_K \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & I_K & & -I_K & 0 & 0 & \dots & I_K & S_{iJ} - I_K \end{array}\right] \left[\begin{array}{c} \ln h_{i0}^1 \\ \vdots \\ \ln h_{i0}^L \\ \hline \psi_{i0} \\ \psi_{i1} \\ \vdots \\ \psi_{iJ} \\ \hline V_i \end{array}\right]$$

If we define $X_i$ to be the large partitioned matrix, then we have a system in the following form:

$$\begin{bmatrix} X_i \\ R_i \end{bmatrix} \begin{bmatrix} \ln h_i^0 \\ \psi_i \\ V_i \end{bmatrix} = \begin{bmatrix} \ln h_i^+ \\ r_i \end{bmatrix}.$$

Note that under Assumption 6, for any action $j > 0$ in any state $k$, the resulting state is always different from $k$. Therefore, the diagonal elements of $S_{ij}$ are all zero and $S_{ij} - I_K$ has full rank for each $j > 0$. It follows that $X_i$ is a $JKL \times (J+L+2)K$ matrix with rank $(J+L+1)K$. Hence, we will need $JKL - (J+L+1)K = K(JL - J - L - 1)$ full-rank restrictions for identification.

**Theorem 4.** *If for player $i$ there exists a collection of linear restrictions represented by a matrix $R_i$ and vector $r_i$ such that*

$$R_i \begin{bmatrix} \ln h_i^0 \\ \psi_i \\ V_i \end{bmatrix} = r_i$$

*and the matrix $\begin{bmatrix} X_i \\ R_i \end{bmatrix}$ has rank $(J+L+2)K$, then $h_i^0$, $\psi_i$, and $V_i$ are identified.*

First, we note that the number of restrictions per player is independent of the total number of players in the game. Therefore, the total number of required identifying restrictions is only linear in $N$. On the other hand, for discrete time models the number of restrictions needed is exponential in $N$ (Pesendorfer and Schmidt-Dengler, 2008).

It is helpful now to consider some more examples. Suppose the move arrival rate exclusion is a binary variable (e.g., fast and slow states), meaning $L = 2$, and suppose we are considering a binary choice model, $J = 1$. There are $2K$ observable hazards and $5K$ unknowns, so we need $3K$ full-rank restrictions. Since the instantaneous payoffs $\psi_{ij}(x)$ are only identified relative to the cost of inaction $\psi_{i0}(x)$, lets also assume that inaction is costless. This yields $K$ restrictions of the form $\psi_{i0k} = 0$ for $k = 1, \ldots, K$. If we further assume that the move arrival rates are independent of $x$ (but not $z$), this yield $(K-1)L$ restrictions of the form $\sum_{j=0}^{J} h_{ij}(x, z) = \sum_{j=0}^{J} h_{ij}(x', z)$ for all $x, x'$ and all $z$. With only these relatively straightforward identifying assumptions we have $3K - 2$ restrictions. With only two more restrictions, the model is identified.

Now consider the case without a move arrival rate exclusion (i.e., $L = 1$). If inaction is again costless ($\psi_{i0k} = 0$ for all $k$), then there are only $(J + 2)K$ remaining unknowns and $JK$ identified hazards, leaving only $2K$ additional restrictions needed. If we assume that the move arrival rate is constant across states ($\sum_{j=0}^{J} h_{ij}(x) = \sum_{j=0}^{J} h_{ij}(x')$ for all $x, x'$) then we have $K - 1$ restrictions. Only $K + 1$ additional restrictions are needed to identify the model.

Finding additional full-rank restrictions is not difficult in most applications. Examples include states where the value function is known, for example, if $V_{ik} = 0$ when a firm has permanently exited. Exclusion restrictions of the form $V_{ik} = V_{ik'}$ are also common, where $k$ and $k'$ are two states that differ only by a rival-specific state and are payoff equivalent to firm $i$. Finally, states where the instantaneous payoffs are the same provide restrictions, for example we may assume that entry costs or scrap values are constant across states which would imply $\psi_{ijk} - \psi_{ijk'} = 0$ for all $i$, some choice $j$, all states $k$ and $k'$. In these cases, the rank condition can be verified almost by inspection in applications.

When there is no move arrival rate exclusion ($L = 1$), another possible restriction that mirrors the assumed fixed decision times in discrete time models is $\sum_{j=0}^{J} h_{ij}(x) = 1$ for all $i$ and $x$. This $K$ restrictions per player and reduces the number of unknowns.

It remains to identify the $K$-vector of payoffs $u_i$ for each player $i$. In light of the linear representation in (7),

$$u_i = \Xi_i(Q)V_i - L_i C_i(\sigma_i)$$

where $\Xi_i$ is the matrix function defined in (8). Under the maintained assumptions and restrictions, $V_i$ and $\psi_i$ are identified for each player. The choice probabilities $\sigma$ are also

identified since $Q$ is identified. Therefore, $u_i$ can be obtained from the equation above.

**Theorem 5.** *Under the maintained assumptions, if for any player $i$ the quantities $V_i$, $\psi_i$, and $Q$ are identified, then the flow payoffs $u_i$ are identified.*

**Example 1, Continued.** In the single-agent renewal model, since the replacement cost does not depend on the mileage state we have $\psi_{1k} = c$ for all $k$. This alone yields $K-1$ restrictions of full rank of the form $\psi_{1k} - \psi_{11} = 0$ for all $k$. We also assumed the rate of move arrivals is constant across states, yielding $K-1$ additional restrictions. The linearity of the utility function imposes restrictions on $V$, and although this does not fit in the linear restriction framework of Theorem 4 it also contributes to identification of $\psi$ and $V$.

**Example 2, Continued.** In the simple two-player entry-exit model, we may suppose that the entry costs and scrap values are independent of the market state (high or low demand) and whether a rival is present. In other words, $\psi_{i1k} - \psi_{i11}$ for all states $k$, yielding $K-1$ restrictions per player or $2K-2$ total restrictions.

## 4. Estimation

In this section, we review the two-step CCP or pseudo maximum likelihood (PML) estimator which was proposed by Arcidiacono, Bayer, Blevins, and Ellickson (2016) and was based on the ideas of Hotz and Miller (1993). We will examine this estimator in the Monte Carlo experiments of Section 6.

Since it is unknown whether there is a unique equilibrium of the model, we maintain a widely-used assumption from the discrete time literature on the data generating process: we require that the data be generated by a single Markov perfect equilibrium at the true values and that all players expect the same equilibrium to be played at all instants both in and out of sample. This assumption allows us to estimate the model even in the presence of multiple equilibria.[5]

**Assumption 9.** (a) In each market $m = 1, \ldots, M$, players expect the same equilibrium, with intensity matrix $Q^0$, to be played for all times $t \in [0, \infty)$. (b) The distribution of state transitions in each market $m$ and each time period $t$ is consistent with row $k_{mt}$ of the transition matrix $P^0(\Delta) = \exp(\Delta Q^0)$.

Now, let $h$ denote the vector of all distinct hazards appearing in the matrices $Q_0, Q_1, \ldots, Q_N$. Namely, let $h_{ijk}$ denote the hazard of player $i$ choosing action $j$ in state $k$,

$$h = (h_{012}, \ldots, h_{0,K-1,K}, h_{111}, \ldots, h_{1JK}, \ldots, h_{ijk}, \ldots, h_{N11}, \ldots, h_{NJK}), \tag{14}$$

[5]Note that partial-identification-based approaches such as Tamer (2003) and Ciliberto and Tamer (2009) do not require this assumption.

and let $\mathcal{H} \subset \mathbb{R}^{K(K-1)+NJK}$ denote the space of all possible vectors $h$. This vector includes hazards for nature for $i = 0$, where $h_{0lk}$ denotes the hazard of jumps from state $k$ to state $l$. However, note that the continuation action hazards for $j = 0$ will not appear in $Q$ unless the continuation choice is separately observed. Furthermore, if the move arrival rates $\lambda_{ik}$ are known we can impose this structure and write $h$ as

$$h = (q_{012}, \ldots, q_{0,K-1,K}, \lambda_{11}\sigma_{111}, \ldots, \lambda_{ik}\sigma_{ijk}, \ldots, \lambda_{N1}\sigma_{N11}, \ldots, \lambda_{NK}\sigma_{NJK}).$$

Given a vector $h$ the corresponding matrix $Q$ can be reconstructed.

Let $\theta$ denote the finite-dimensional vector of parameters which determine the move arrival rates (if estimated), the flow payoffs $u$, and the instantaneous payoffs $\psi$. Then, let $\Lambda : \Theta \times \mathcal{H} \to \mathcal{H} : (\theta, h) \mapsto \Lambda(\theta, h)$ denote the mapping which, given a value of $\theta$ and a vector of hazards $h$, a new vector of hazards $h' = \Lambda(\theta, h)$ is determined using the inverse CCP mapping of Theorem 1. This proceeds in two steps: given $\theta$ and $h$, we reconstruct $Q$ and $\sigma$ and obtain new value functions using (7). Given the new value functions, we determine the new choice probabilities $\sigma'$ and then form the new vector of hazards $h'$. In equilibrium, the conditional choice probabilities are such that $h = \Lambda(\theta, h)$.

The remaining conditions are standard regularity conditions requiring markets to be independent, the parameter space to be compact, the population parameter vector to be identified, and the hazard mapping $\Lambda$ to be sufficiently smooth.

**Assumption 10.** (a) The observations $\{k_{mt} : m = 1, \ldots, M, t = 1, \ldots, T\}$, sampled on the lattice $\{t\Delta : t = 1, \ldots, T\}$, are independent across markets $m$ and $\Pr(k_{mt} = k) > 0$ for all $k \in \mathcal{K}$. (b) $\Theta$ is compact and $\theta^0 \in \text{int}(\Theta)$. (c) For any $\theta \in \Theta$ with $\theta \neq \theta^0$ and any $h$ such that $h = \Lambda(\theta, h)$, we have $h \neq h^0$. (d) $\Lambda : \Theta \times \mathcal{H} \to \mathcal{H} : (\theta, h) \mapsto \Lambda(\theta, h)$ is twice continuously differentiable.

### *4.1. Maximum Likelihood Estimation*

The model can be estimated using maximum likelihood if either the equilibria can be enumerated (which is may be feasible, given the computational tractability of the continuous-time model) or there is a unique equilibrium. It is likely that methods proposed for discrete time models, such as the homotopy method (Borkovsky, Doraszelski, and Kryukov, 2010; Besanko, Doraszelski, Kryukov, and Satterthwaite, 2010; Bajari, Hong, Krainer, and Nekipelov, 2010) or recursive lexicographical search (Iskhakov, Rust, and Schjerning, 2016), could be adapted to the present model as well, but investigating such possibilities is beyond the scope of this paper.

Define the pseudo log likelihood function for parameters $\theta$ and arbitrary hazards $h$ as

$$L_M(\theta, h) = \frac{1}{M} \sum_{m=1}^{M} \sum_{t=1}^{T} \ln P\left(k_{m,t-1}, k_{mt}; \Delta, \Lambda(\theta, h)\right),$$

where $P(k,l;\Delta,h)$ denotes the $(k,l)$ element of the transition matrix induced by $h$. As in Aguirregabiria and Mira (2007), we say this is a pesudo log likelihood function because it can be evaluated at hazards $h$ that are not necessarily equilibrium values (e.g., for two-step estimation we will use estimates $\hat{h}$ of equilibrium hazards).

If we were to estimate the model under conditions known to yield a unique equilibrium for each $\theta \in \Theta$, then we could define the maximum likelihood estimator simply as

$$\hat{\theta}_{\text{ML}} = \arg\max_{\theta\in\Theta} L_M(\theta,h)$$
$$\text{subject to } h = \Lambda(\theta,h).$$

Alternatively, when there are multiple equilibria we could define the estimator in terms of the equilibrium that maximizes the likelihood. Let $\mathcal{E}(\theta) = \{h : h = \Lambda(\theta,h)\}$ denote the set of equilibrium hazards for given parameters $\theta$. Then define

$$\hat{\theta}_{\text{ML}} = \arg\max_{\theta\in\Theta} \left\{ \max_{(h,V)\in\mathcal{E}(\theta)} L_M(\theta,h) \right\}$$

*4.2. Two-Step Pseudo Maximum Likelihood Estimation*

Suppose we have a $\sqrt{M}$-consistent first stage M-estimator $\hat{h}$ for $h^0$. ABBE defined a feasible two-step pseudo maximum likelihood estimator $\hat{\theta}$ of $\theta$ as

$$\hat{\theta}_{\text{PML}} = \arg\max_{\theta\in\Theta} L_M(\theta,\Lambda(\theta,\hat{h}))$$

and show that under Assumption 10 above $\hat{\theta}_{\text{PML}}$ is consistent and asymptotically normal. This is similar to both the CCP estimator of Hotz and Miller (1993) and the pseudo maximum likelihood PML estimator of Aguirregabiria and Mira (2007). We examine the finite sample and computational properties of this estimator and the computational properties of the model more generally in the Monte Carlo experiments presented in Section 6.

## 5. A Continuous-Time Quality Ladder Model of Oligopoly Dynamics

To illustrate the application to dynamic games used in empirical industrial organization we consider a discrete control version of the quality ladder model proposed by Ericson and Pakes (1995). This model has been examined extensively by Pakes and McGuire (1994, 2001), Doraszelski and Satterthwaite (2010), Doraszelski and Pakes (2007), and several others. The model consists of at most $N$ firms 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 $\Omega = \{1,2,\ldots,\bar{\omega},\bar{\omega}+1\}$ is the finite set of possible quality levels, with

$\bar{\omega}+1$ denoting the "quality" of inactive firms. Firms with $\omega_i < \bar{\omega}+1$ are incumbents. In contrast to Pakes and McGuire (1994), all controls here are discrete: given a move arrival, firms choose whether or not to move up the quality ladder, not how much to spend to increase their chances of doing so.

We consider the particular example 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 easily adapted to other settings. 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 $\omega_i$ is the share of consumers who are aware of firm $i$'s product. Gowrisankaran (1999a) develops a model of endogenous horizontal mergers where $\omega_i$ 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 $\omega_i$.

### *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. 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 $N$, given by $\mathcal{S} = \{(s_1,\ldots,s_{\bar{\omega}+1}) : \sum_j s_j = N, s_j \in \mathbb{Z}^*\}$, where $\mathbb{Z}^*$ is the set of nonnegative integers. In this notation, each vector $\omega = (\omega_1,\ldots,\omega_N) \in \Omega^N$ maps to an element $s = (s_1,\ldots,s_{\bar{\omega}+1}) \in \mathcal{S}$ with $s_j = \sum_{i=1}^N 1\{\omega_i = j\}$ for each $j$.

In practice we map the multidimensional space $\mathcal{S}$ to an equivalent one-dimensional state space $\mathcal{X} = \{1,\ldots,|\mathcal{S}|\}$.[6] Payoff relevant market configurations from the perspective of firm $i$ are then uniquely described by two integers $(x,\omega_i)$, where $x \in \mathcal{X}$ denotes the market structure and $\omega_i$ is firm $i$'s own quality level.

### *5.2. Product Market Competition*

Again, we follow Pakes and McGuire (1994) in assuming a continuum of consumers with measure $M > 0$ and that each consumer's utility from choosing the good produced by firm $i$ is $g(\omega_i) - p_i + \varepsilon_i$, where $\varepsilon_i$ is iid across firms and consumers and follows a type I extreme value distribution. The $g$ function is used to enforce an upper bound on profits.

[6] In particular, we use the "probability density space" encoding algorithm described in Gowrisankaran (1999b), to map market structure tuples $s \in \mathcal{S}$ to integers $k \in \mathcal{X}$.

As in Pakes, Gowrisankaran, and McGuire (1993), for some constant $\omega^*$ we specify the function

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

Let $\varsigma_i(\omega, p)$ denote firm $i$'s market share given the state $\omega$ and prices $p$. From McFadden (1974), we know that the share of consumers purchasing good $i$ is

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

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

All firms have the same constant marginal cost $c \geq 0$. Taking the prices of other firms, $p_{-i}$, as given, 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 (in this single-product firm setting) there is a unique Bertrand-Nash equilibrium, which is 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

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

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 $(\omega_i, \omega_{-i})$.

### *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 $\kappa$ in order to increase the quality of their product from $\omega_i$ to $\omega_i' = \min\{\omega_i + 1, \bar{\omega}\}$, or they may exit the market and receive some scrap value $\varphi$. We denote these choices, respectively, by the choice set $\mathcal{J}_i = \{0, 1, 2\}$. When an incumbent firm exits the market, $\omega_i$ jumps deterministically to $\bar{\omega} + 1$. Associated with each choice $j$ is a private shock $\varepsilon_{ijt}$. These shocks are iid over firms, choices, and time and follow a standard 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 $\omega_k = (\omega_{k1}, \ldots, \omega_{kN})$ denote its counterpart in $\Omega^N$. In the general notation introduced above, the instantaneous payoff $\psi_{ijk}$ to firm $i$ from choosing choice $j$ in state $k$ is

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

The state resulting from continuing ($j = 0$) is simply $l(i, 0, k) = k$. Similarly, for investment ($j = 1$), $l(i, 1, k) = k'$ where state $k'$ is the element of $\mathcal{X}$ such that $\omega_{k'i} = \min\{\omega_{ki} + 1, \bar{\omega}\}$ and $\omega_{k'm} = \omega_{km}$ for all firms $m \neq i$. Note that we are considering only incumbent firms with $\omega_{ki} < \bar{\omega} + 1$. Exiting is a terminal action with an instantaneous payoff, but no continuation value.

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

$$V_{ik} = \frac{1}{\rho + \sum_{l \neq k} q_{kl} + \sum_{m=1}^{N} \lambda_{mk}} \left( \pi_{ik} + \sum_{l \neq k} q_{kl} V_{il} + \sum_{m \neq i} \lambda_{mk} \sum_{j} \sigma_{mjk} V_{i,l(m,j,k)} \right. \\ \left. + \lambda_{ik} \, \mathrm{E} \max \left\{ V_{ik} + \varepsilon_{i0}, V_{i,l(i,1,k)} - \kappa + \varepsilon_{i1}, \varphi + \varepsilon_{i2} \right\} \right)$$

where $\pi_{ik} = \pi(\omega_{ki}, \omega_{k,-i})$ represents the flow profit accruing from product market competition and $\lambda_{ik} = \lambda$ for incumbents and potential entrants and $\lambda_{ik} = 0$ if firm $i$ is not active in state $k$. Conditional upon moving while in state $k$, incumbent firms face the maximization problem $\max \{V_{ik} + \varepsilon_{i0}, -\kappa + V_{ik'} + \varepsilon_{i1}, \varphi + \varepsilon_{i2}\}$. The resulting choice probabilities are

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

where, as before, $k' = l(i, 2, k)$ denotes the resulting state after investment by firm $i$.

### *5.4. Potential Entrants*

Whenever the number of incumbents is smaller than $N$, a single potential entrant receives the opportunity to enter at rate $\lambda$. Potential entrants are short-lived and do not consider the option value of delaying entry. If firm $i$ is a potential entrant with the opportunity to

move it has two choices: it can choose to enter ($j = 1$), paying a setup cost $\eta$ and entering the market immediately in a predetermined entry state $\omega^e \in \Omega$ or it can choose not to enter ($j = 0$) at no cost. Associated with each choice $j$ is a stochastic private payoff shock $\varepsilon^e_{ijt}$. These shocks are iid across firms, choices, and time and are distributed according to the type I extreme value distribution.

In our general notation, for actual entrants ($j = 1$) in state $k$ the instantaneous payoff is $\psi_{i1k} = -\eta$ and the continuation state is $l(i, 1, k) = k'$ where $k'$ is the element of $\mathcal{X}$ with $\omega_{k'i} = \omega^e$ and $\omega_{k'm} = \omega_{km}$ for all $m \neq i$. For firms that choose not to enter ($j = 0$) in state $k$, we have $\psi_{i0k} = 0$ and the firm leaves the market with no continuation value. Thus, upon moving in state $k$, a potential entrant faces the problem

$$\max\{\varepsilon^e_{i0}, -\eta + V_{ik'} + \varepsilon^e_{i1}\}$$

yielding the conditional entry-choice probabilities

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

### *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 $\gamma$, the quality of each firm $i$ jumps from $\omega_i$ to $\omega_i' = \max\{\omega_i - 1, 1\}$. This process represents an industry-wide (negative) demand shock, interpreted as an improvement in the outside alternative.

## 6. Monte Carlo Experiments

In this section we describe Monte Carlo experiments conducted using the single-agent renewal model of Example 1 and the quality ladder model described in Section 5.

### *6.1. Single-Agent Dynamic Discrete Choice*

Here, we generate data according to the simple single player binary choice model described in Example 1. The primitives of the model are the payoff (mileage cost) parameter $\beta$, the intensity matrix (mileage transition) parameters $q_1$ and $q_2$, the reset (engine replacement) cost $c$, the discount rate $\rho$, and the move arrival rate $\lambda$. We fix $\rho = 0.05$ and focus on estimating $\theta = (\lambda, q_1, q_2, \beta, c)$. We first choose values for $\theta$ and then use value function iteration to determine the value function over the state space $\mathcal{X}$ to within a tolerance of $\varepsilon = 10^{-3}$ in the relative sup norm. We then use the resulting value function to generate data for various values of $T$ and $\Delta$.

In the first set of experiments, we use the full-solution maximum likelihood approach to estimate the model. The value functions are obtained through value function iteration for each value of $\theta$ in an inner loop while maximizing the likelihood function in an outer loop using the L-BFGS-B algorithm (Byrd, Lu, and Nocedal, 1995; Zhu, Byrd, Lu, and Nocedal, 1997) with numerical derivatives with step size $h = 10^{-8}$. We estimate the model under several different scenarios including full continuous-time data, continuous-time data when the continuation decision is not observed, and discrete time data. In the experiments involving discrete time data, we compare the estimator across several choices of the observation interval $\Delta$ while keeping the overall observation window $[0, T]$ fixed. In each experiment we fixed the discount rate at $\rho = 0.05$ and the number of mileage states at $K = 10$. The population parameters are $(q_1, q_2, \lambda, \beta, c) = (0.15, 0.05, 0.2, -0.1, 1.25)$.

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 continuation choice) are unobserved, and discrete time data observed at intervals $\Delta \in \{0.625, 1.25, 2.5, 5.0, 10.0\}$.[7] For each specification, we report the mean bias and the root mean square error (RMSE) of the parameter estimates over 100 replications in Table 1. All are parameters are estimated quite precisely and with little finite-sample bias. The loss in precision from moving away from continuous time data is initially greatest for the move arrival rate, $\lambda$, yet all estimates of this parameter are still very precise. The replacement cost, $c$, also loses precision with more coarsely sampled data, but the increases are not large until we move to seeing only one in four events on average in the sampling period.

We also carry out the same experiments using CCP-based estimation in the single-agent model. The results are displayed in Table 2. Again, we estimate the model with full continuous-time data, a continuous-time dataset with missing passive actions, and several discrete-time datasets of varying granularity. For the full continuous-time dataset, we can nonparametrically estimate the CCPs using a simple frequency estimator. When accounting for passive moves, we approximate the CCPs by dividing the number of times each particular observed choice was made in each state by the implied expected number of move arrivals in that state. Finally, when estimating the model with discrete-time data, we first jointly estimate the first-stage parameters ($\lambda$, $q_1$, and $q_2$) and the parameters of a logistic regression model for the probability of renewal with parameters $\alpha$. The regressors in our logit model are a constant, the state $x$, and $\ln x$. Then, we invert the predicted CCPs

[7] One could view 25,000 as the number of months in the data with $\{0.625, 1.25, 2.5, 5.0, 10.0\}$ indicating the number of months (or fraction of months) between samples. While 25,000 implies having over 2,000 years of data, this is following *one* time series. An almost equivalent structure would follow 1000 decision-makers over two years.

obtained using the estimated parameters $\hat{\alpha}$ to obtain the value function which we use to estimate the remaining second stage parameters.

When passive moves are not observed we cannot estimate the CCPs directly. Since the full choice set is not observed, we first store $\sum_{t=1}^{T} 1\{a_t = j, x_t = k\} / \sum_{t=1}^{T} 1\{x_t = k\}\tau_t$ for each state $k$ and each choice $j$. The numerator is the number of observations for which $j$ is chosen in state $k$ and the denominator is the total time spent in state $k$. Dividing this quantity by $\lambda$ yields an estimate of the true CCP since $1/\lambda$ is the average time between move arrivals.

Using CCPs increases the mean square error slightly, reflecting noise from the first stage. However, the estimates are still very good, particularly when the average number of state changes per sampling interval is small.

Finally, we also estimated the model with continuous-time data while allowing for buses to be of two distinct types, where the type is not observed by the econometrician. In this specification, the type affected both the mileage transition probabilities and payoff parameters. In particular, with probability $\pi$, the bus is of the first type and with probability $1 - \pi$, the bus is of the second type. For buses of type $m = 1, 2$, the mileage jumps forward one unit at rate $q_1$ and two units at rate $q_{2m}$, the cost of mileage is $\beta$, and the cost of replacement is $c_m$. Again, estimation proceeded quickly with little difficulty in separating the unobserved heterogeneity from the other model parameters. The results are reported in Table 3 for varying numbers of agents, $M$, and observations per agent, $N$.

### *6.2. A Dynamic Discrete Game*

Our second set of Monte Carlo experiments corresponds to the quality ladder model described in Section 5. We estimate models ranging from 10 to 30 firms and obtain estimates of $\theta = (\lambda, \gamma, \kappa, \eta, \varphi)$. The population parameters used were $(\lambda, \gamma, \kappa, \eta, \varphi) = (1.9, 0.1, 2.0, 8.0, 2.0)$. In all experiments, as before, we fixed $\rho = 0.05$. In these experiments, we used samples containing $T = 200$ continuous time events in each of $M = 200$ markets. For each of 25 simulated datasets, we report the mean bias and root mean square error (RMSE) of the parameter estimates.

In all experiments, we hold the number of possible quality levels fixed at $\bar{\omega} = 7$, set $\omega^e = \lfloor \frac{\bar{\omega}}{2} \rfloor$, and vary the maximum number of players, $N$, and the market size, $M$. We increase the market size ($M$) so that the average number of active players ($n_{\text{avg}}$) grows with the total number of possible players ($N$). The average quality level of active firms is denoted $\omega_{\text{avg}}$. We also report $K$, the number of states from the perspective of player $i$—the number of distinct $(\omega, \omega_i)$ combinations. The size of the state space for the largest problem is over 58 million.

Table 4 summarizes the results for full-solution estimation, where we obtain the value

function to within a tolerance of $\varepsilon = 10^{-5}$ in the relative sup norm using value function iteration for each trial value of $\theta$. Table 5 presents the analogous results obtained using CCP estimation, where we assume the true CCPs are available. In all cases, both full-solution methods and CCP estimation perform extremely well in terms of having small finite sample bias and mean square error.

For CCP estimation, we substitute the true CCPs in the $\Gamma$ mapping to estimate the value function. In practice, the CCPs must be estimated in a preliminary step. However, because there are many possible methods for doing so, and because they tend to be application- and data-specific, we simply present the results for the second-stage parameters as if the true CCPs were known. We have estimated the CCPs nonparametrically using locally weighted averages with little change in the results.

We then compare the computational time required for both full-solution and CCP estimation in Table 6. We first report the number of players $N$, the market size $M$, and the total number of states $K$. In all cases we use L-BFGS-B to maximize the log-likelihood function or pseudo log-likelihood function with starting values $(0.5, 0.5, 0.5, 6.0, 1.0)$.[8] For each model, computational times are reported for only one replication. Since we consider many models, the overall trends are clear despite the fact that we do not report averages.[9]

The first timing column reports the time required to obtain the value function $V$ for each model specification. This step is necessary to either generate a dataset or to simulate the model (e.g., to perform counterfactuals). In particular, we use datasets consisting of $M = 200$ markets with $T = 200$ continuous time events observed in each. We estimate the parameters $\lambda$ and $\gamma$ in a first step, but we do not report times because they were less than two seconds in all cases. This step is independent of the method used to obtain the value function. Next, we report the total time required to estimate the second stage parameters $\kappa$, $\eta$, and $\varphi$ via full-solution estimation. For each new trial value of $\theta$, we use the value function at the previous $\theta$ as the new starting value for the value function iteration. Finally, we report the time required to estimate the second-stage parameters.[10]

Even in a game with 30 players and over 58 million states, full-solution estimation took under seven hours. Conditional on already having the CCPs from a first stage, two-step estimation times were incredibly fast, with the longest taking less than three minutes.

[8] We use central finite-difference gradients with a stepsize of $h = 10^{-6}$ for full-solution estimation and $h = 10^{-9}$ for CCP estimation.

[9] All reported times are seconds CPU time required for estimation of a single replication of each specification on a dual processor Intel Xeon X5670 system. Our programs are written in Fortran and take advantage of parallel processing, however, the reported times are total CPU time used across all processors.

[10] This table does not address the time required to estimate the first-stage CCPs, which can vary significantly depending on which method is used. Parametric methods can clearly be quite fast while fully nonparametric methods can be computationally intensive.

To put these numbers in perspective, Doraszelski and Judd (2012) note that it would take about a year to just solve for the equilibrium of a 14 player game using the basic Pakes-McGuire algorithm.[11] For the continuous time model, it takes less than one minute to solve the game and twenty minutes to estimate the parameters using a full-solution (NFXP) approach. CCP estimation requires only two seconds. These computational times suggest that very large classes of problems can be easily estimated in a continuous-time framework. Furthermore, the computational time required to calculate the fixed point once in continuous time is small even for very large problems. This implies that simulating counterfactuals from large-scale models will not be an issue.

## 7. Conclusion

In this paper we have developed new results on the theoretical and econometric properties of the empirical framework introduced by Arcidiacono, Bayer, Blevins, and Ellickson (2016) for continuous time dynamic discrete choice games. We established equilibrium existence in a more general class of models with heterogeneous firms and state-dependent move arrival rates, we developed more general conditions for nonparametric identification with discrete time data, we explored these results in the context of three canonical examples widely used in applied work, and we examined the computational properties of the model as well as the finite sample properties of the estimator through a series of large-scale Monte Carlo experiments that are representative of models used in applied work.

## References


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

Arcidiacono, P., P. Bayer, J. R. Blevins, and P. B. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *Review of Economic Studies 83*, 889–931.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica 75*, 1331–1370.

Bajari, P., H. Hong, J. Krainer, and D. Nekipelov (2010). Computing equilibria in static games of incomplete information using the all-solution homotopy.


[11] Similar computational times are also reported in Doraszelski and Pakes (2007). These times are for a model with $\bar{\omega} = 9$ without entry and exit, which for a fixed value of $N$ is roughly comparable to our model with $\bar{\omega} = 7$, which includes entry and exit.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica 78*, 453–508.

Billingsley, P. (1961). *Statistical Inference for Markov Processes*. Chicago: University of Chicago Press.

Blevins, J. R. (2017). Identifying restrictions for finite parameter continuous time models with discrete time data. *Econometric Theory 33*, 739–754.

Borkovsky, R. N., U. Doraszelski, and Y. Kryukov (2010). A user's guide to solving dynamic stochastic games using the homotopy method. *Operations Research 58*(4), 1116–1132.

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.

Caplin, A. and B. Nalebuff (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica 59*, 25–59.

Chung, K. L. (1967). *Markov Chains with Stationary Transition Probabilities*. Berlin: Springer.

Ciliberto, F. and E. Tamer (2009). Market structure and multiple equilibria in the airline industry. *Econometrica 77*, 1791–1828.

Cosman, J. (2014). Industry dynamics and the value of variety in nightlife: Evidence from Chicago. Working paper, University of British Columbia.

Culver, W. J. (1966). On the existence and uniqueness of the real logarithm of a matrix. *Proceedings of the American Mathematical Society 17*, 1146–1146.

Cuthbert, J. R. (1972). On uniqueness of the logarithm for Markov semi-groups. *Journal of the London Mathematical Society s2-4*, 623–630.

Cuthbert, J. R. (1973). The logarithm function for finite-state Markov semi-groups. *Journal of the London Mathematical Society s2-6*, 524–532.

Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics 3*, 53–93.

Doraszelski, U. and S. Markovich (2007). Advertising dynamics and competitive advantage. *The RAND Journal of Economics 38*, 557–592.

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.

Doraszelski, U. and M. Satterthwaite (2010). Computable Markov-perfect industry dynamics. *The RAND Journal of Economics 41*, 215–243.

Elfving, G. (1937). Zur Theorie der Markoffschen Ketten. *Acta Societatis Scientiarum Fennicæ Nova Series A 2*(8), 1–17.

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

Gantmacher, F. R. (1959). *The Theory of Matrices*, Volume 1. New York: Chelsea.

Geweke, J. (1978). Temporal aggregation in the multiple regression model. *Econometrica 46*, 643–661.

Geweke, J., R. C. Marshall, and G. A. Zarkin (1986). Exact inference for continuous time Markov chain models. *Review of Economic Studies 53*, 653–669.

Gotz, G. A. and J. J. McCall (1980). Estimation in sequential decision-making models: A methodological note. *Economics Letters 6*, 131–136.

Gowrisankaran, G. (1999a). A dynamic model of endogenous horizontal mergers. *The RAND Journal of Economics 30*, 56–83.

Gowrisankaran, G. (1999b). Efficient representation of state spaces for some dynamic models. *Journal of Economic Dynamics and Control 23*, 1077–1098.

Hansen, L. P. and T. J. Sargent (1983). The dimensionality of the aliasing problem in models with rational spectral densities. *Econometrica 51*, 377–387.

Heckman, J. J. and B. Singer (1986). Econometric analysis of longitudinal data. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics*, Volume 3, Chapter 29. Amsterdam: North Holland.

Hong, H., W. Li, and B. Wang (2015). Estimation of dynamic discrete models from time aggregated data. *Journal of Econometrics 188*(2), 435–446.

Horn, R. A. and C. R. Johnson (1985). *Matrix Analysis*. New York: Cambridge University Press.

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

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.

Iskhakov, F., J. Rust, and B. Schjerning (2016). Recursive lexicographical search: Finding all Markov perfect equilibria of finite state directional dynamic games. *Review of Economic Studies 83*, 658–703.

Jeziorski, P. (2013). Empirical model of dynamic merger enforcement—choosing ownership caps in U.S. radio. Working paper, UC Berkeley, Haas School of Business.

Johansen, S. (1974). Some results on the imbedding problem for finite Markov chains. *Journal of the London Mathematical Society s2-8*, 345–351.

Karlin, S. and H. M. Taylor (1975). *A First Course in Stochastic Processes* (Second ed.). San Diego, CA: Academic Press.

Kessler, M. and A. Rahbek (2004). Identification and inference for multivariate cointegrated and ergodic Gaussian diffusions. *Statistical Inference for Stochastic Processes 7*, 137–151.

Kingman, J. F. C. (1962). The imbedding problem for finite Markov chains. *Probability Theory and Related Fields 1*, 14–24.

Lee, C.-Y., J. W. Roberts, and A. Sweeting (2012). Competition and dynamic pricing in a perishable goods market. Working paper, Duke University.

McCrorie, J. R. (2003). The problem of aliasing in identifying finite parameter continuous time stochastic models. *Acta Applicandae Mathematicae 79*, 9–16.

McFadden, D. L. (1974). Conditional logit analysis of qualitative choice analysis. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press.

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

Nevskaya, Y. and P. Albuquerque (2012). A continuous time model of product usage: Measuring the effect of product design and rewards in online games. Working paper, University of Rochester, Simon School of Business.

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

Pakes, A., G. Gowrisankaran, and P. McGuire (1993). Implementing the Pakes-McGuire algorithm for computing Markov perfect equilibria in Gauss. Working paper, Harvard University.

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.

Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric Markov perfect equilibrium, and the 'curse' of dimensionality. *Econometrica 69*, 1261–1281.

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.

Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies 75*, 901–928.

Phillips, P. C. B. (1972). The structural estimation of a stochastic differential equation system. *Econometrica 40*, 1021–1041.

Phillips, P. C. B. (1973). The problem of identification in finite parameter continuous time models. *Journal of Econometrics 1*, 351–362.

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

Schiraldi, P., H. Smith, and Y. Takahashi (2012). Estimating a dynamic game of spatial competition: The case of the U.K. supermarket industry. Working paper, London School of Economics.

Sidje, R. B. (1998). Expokit: A software package for computing matrix exponentials. *ACM Transactions on Mathematical Software 24*, 130–156.

Sims, C. A. (1971). Discrete approximations to continuous time distributed lags in econometrics. *Econometrica 39*, 545–563.

Singer, B. and S. Spilerman (1976). The representation of social processes by Markov models. *The American Journal of Sociology 82*(1), 1–54.

Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies 70*, 147–165.

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

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.

## A. Proofs

*Proof of Lemma 2.* The result follows directly from the joint distribution function under independence:

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

Therefore, $\tau$ has an exponential distribution with rate parameter $\sum_{i=1}^{N} \lambda_i$.

Furthermore,

$$\begin{aligned}\Pr(\tau_i \leq \tau_j \; \forall j) &= \mathrm{E}\left[\Pr(\tau_j \geq \tau_i \; \forall j \neq i) \mid \tau_i\right] \\ &= \int_0^\infty \left[\mathrm{e}^{-\sum_{j\neq i} \lambda_j}\right] \lambda_i \, \mathrm{e}^{-\lambda_i \tau_i} \; d\tau_i \\ &= \int_0^\infty \lambda_i \, \mathrm{e}^{-(\sum_{j=1}^{N} \lambda_j)\tau_i} \; d\tau_i \\ &= -\frac{\lambda_i}{\sum_{j=1}^{N} \lambda_j} \left[\mathrm{e}^{-(\sum_{j=1}^{N} \lambda_j)\tau_i}\right]_{\tau_i=0}^{\infty} \\ &= \frac{\lambda_i}{\sum_{j=1}^{N} \lambda_j}.\end{aligned}$$

■

*Proof of Theorem 1.* Given a collection of equilibrium best response probabilities $\{\sigma_i\}_{i=1}^{N}$, we can obtain a matrix expression for the value function $V_i(\sigma_i)$. By Proposition 2 of Arcidiacono, Bayer, Blevins, and Ellickson (2016), the difference $V_{i,l(i,j,k)}(\sigma_i) - V_{i,l(i,j',k)}(\sigma_i)$ can be expressed as a function of payoffs and choice probabilities $\sigma_i$ and so we can write $C_i$ as a function of only conditional choice probabilities and payoffs (i.e., so that it no longer depends on the value function).

Note that we can write the value function in vector form as follows:

$$\begin{aligned}V_i(\sigma_i) &\left[\left(\rho_i I_K + \sum_{m=1}^{N} L_m\right) - (Q_0 - \tilde{Q}_0)\right] \\ &= u_i + \tilde{Q}_0 V_i(\sigma_i) + \sum_{m\neq i} L_m \Sigma_m(\sigma_m) V_i(\sigma_i) + L_i \left[\Sigma_i(\sigma_i) V_i(\sigma_i) + C_i(\sigma_i)\right].\end{aligned}$$

| Sampling | $n$ | | $q_1$ | $q_2$ | $\lambda$ | $\beta$ | $c$ |
|---|---|---|---|---|---|---|---|
| Continuous Time | 10,000 | Bias | 0.000 | 0.000 | -0.000 | 0.010 | 0.005 |
| | | RMSE | 0.002 | 0.001 | 0.003 | 0.068 | 0.053 |
| Passive Moves | 7,176 | Bias | 0.000 | 0.000 | -0.002 | 0.052 | -0.030 |
| | | RMSE | 0.002 | 0.001 | 0.025 | 0.180 | 0.200 |
| $\Delta = 0.625$ | 40,000 | Bias | -0.013 | 0.003 | -0.011 | 0.110 | 0.055 |
| | | RMSE | 0.013 | 0.004 | 0.022 | 0.238 | 0.246 |
| $\Delta = 1.25$ | 20,000 | Bias | -0.005 | 0.001 | -0.009 | 0.075 | -0.055 |
| | | RMSE | 0.006 | 0.002 | 0.026 | 0.224 | 0.301 |
| $\Delta = 2.5$ | 10,000 | Bias | -0.003 | 0.001 | -0.003 | 0.015 | -0.084 |
| | | RMSE | 0.005 | 0.002 | 0.027 | 0.335 | 0.419 |
| $\Delta = 5.0$ | 5,000 | Bias | 0.001 | 0.000 | -0.005 | 0.087 | -0.019 |
| | | RMSE | 0.007 | 0.003 | 0.020 | 0.265 | 0.403 |
| $\Delta = 10.0$ | 2,500 | Bias | 0.008 | -0.002 | 0.000 | 0.008 | -0.153 |
| | | RMSE | 0.020 | 0.007 | 0.022 | 0.398 | 0.634 |

The mean bias and root mean square error (RMSE) are reported for 100 simulated datasets under several sampling regimes. Passive moves refers to datasets for which the choice $a = 0$ is not observed while $\Delta$ 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 fixed interval $[0, T]$ with $T = 25,000$.

TABLE 1. Single player Monte Carlo results: NFXP estimation.

| Sampling | $n$ | | $q_1$ | $q_2$ | $\lambda$ | $\beta$ | $c$ |
|---|---|---|---|---|---|---|---|
| Continuous Time | 10,000 | Bias | -0.000 | 0.000 | 0.000 | 0.016 | 0.007 |
| | | RMSE | 0.002 | 0.001 | 0.003 | 0.066 | 0.053 |
| Passive Moves | 7,176 | Bias | -0.000 | 0.000 | -0.014 | -0.167 | -0.094 |
| | | RMSE | 0.002 | 0.001 | 0.018 | 0.220 | 0.133 |
| $\Delta = 0.625$ | 40,000 | Bias | -0.013 | 0.003 | -0.003 | 0.112 | 0.129 |
| | | RMSE | 0.013 | 0.004 | 0.040 | 0.298 | 0.296 |
| $\Delta = 1.25$ | 20,000 | Bias | -0.005 | 0.001 | 0.009 | 0.064 | 0.113 |
| | | RMSE | 0.006 | 0.002 | 0.048 | 0.303 | 0.334 |
| $\Delta = 2.5$ | 10,000 | Bias | -0.003 | 0.001 | 0.011 | 0.115 | 0.103 |
| | | RMSE | 0.005 | 0.002 | 0.071 | 0.351 | 0.398 |
| $\Delta = 5.0$ | 5,000 | Bias | 0.001 | 0.000 | 0.021 | 0.098 | 0.102 |
| | | RMSE | 0.007 | 0.003 | 0.088 | 0.390 | 0.514 |
| $\Delta = 10.0$ | 2,500 | Bias | 0.004 | -0.001 | 0.050 | 0.176 | 0.151 |
| | | RMSE | 0.018 | 0.008 | 0.219 | 0.562 | 0.746 |

The mean bias and root mean square error (RMSE) are reported for 100 simulated datasets under several sampling regimes. Passive moves refers to datasets for which the choice $a = 0$ is not observed while $\Delta$ denotes the observation interval for discrete-time data. The CCPs were estimated in a first step using a frequency estimator for continuous-time data and via logistic regression on $x$ and $\ln x$ for estimation with time aggregated data. $n$ denotes the average number of observations (continuous-time events or discrete-time intervals) when observing the model on the fixed interval $[0, T]$ with $T = 25,000$.

TABLE 2. Single player Monte Carlo results: CCP estimation.

Rearranging to collect terms involving $V_i(\sigma_i)$ yields

$$V_i(\sigma_i)\left[\rho_i I_k + \sum_{m=1}^{N} L_m[I_K - \Sigma_m(\sigma_m)] - Q_0\right] = u_i + L_i C_i(\sigma_i).$$

The matrix in square brackets side is strictly diagonally dominant: for each $m$ $\rho_m > 0$ by Assumption 2, $L_m$ is a diagonal matrix with strictly positive elements by Assumption 3, $\Sigma_m(\sigma_m)$ has elements in $[0,1]$ with row sums equal to one, and elements of $Q_0$ satisfy $|q_{0kk}| = \sum_{l \neq k} |q_{0kl}|$ in each row $k$. Therefore, by the Levy-Desplanques theorem (Horn and Johnson, 1985, Theorem 6.1.10) this matrix is nonsingular and the representation in Theorem 1 holds. ■

*Proof of Theorem 2.* Define the mapping $\Upsilon : [0,1]^{N \times J \times K} \to [0,1]^{N \times J \times K}$ by stacking the best response probabilities. This mapping defines a fixed point problem for the equilibrium choice probabilities $\sigma$ as follows:

$$\Upsilon_{ijk}(\sigma) = \int 1\left\{\varepsilon_{ij'k} - \varepsilon_{ijk} \leq \psi_{ijk} - \psi_{ij'k} + V_{i,l(i,j,k)}(\sigma_{-i}) - V_{i,l(i,j',k)}(\sigma_{-i}) \quad \forall j' \in \mathcal{J}_i\right\} f(\varepsilon_{ik})\, d\varepsilon_{ik}.$$

The mapping $\Upsilon$ is a continuous function from a compact space onto itself, recalling that $V_{ik}$ has the linear representation of Theorem 1. By Brouwer's theorem, it has a fixed point. The fixed point probabilities imply Markov strategies that constitute a Markov perfect equilibrium. ■

*Proof of Theorem 3.* To establish generic identification of $Q$ we can specialize the proof of Theorem 1 of Blevins (2017) to the present setting, where $Q$ is an intensity matrix with row sums equal to zero and therefore has one real eigenvalue equal to zero and therefore at most $K-1$ complex eigenvalues. In this setting, $P(\Delta)$ is observed and is the solution to the system of differential equations in (9) while $Q$ is a matrix of unknown parameters with $q_{kl}$ for $l \neq k$ being the hazard of jumps from state $k$ to state $l$. The unique solution to this system is the transition matrix $P(\Delta) = exp(\Delta Q)$, which has the same form as the matrix $B$ in equation (3) of Blevins (2017) and $Q$ in this model is analogous to $A$ in (ïIJś). Therefore, and identification of $Q$ depends on establishing a unique solution to an equation involving a matrix exponential of a parameter matrix. Furthermore, in this setting $Q$ is known to have row sums equal to zero, and therefore the vector of ones is a right eigenvector of $Q$ with zero as the corresponding eigenvalue. In this case, the number of required restrictions on $Q$ is reduced to $\lfloor (n-1)/2 \rfloor$ because we know $Q$ has at least one real eigenvalue.

Under the assumptions the number of distinct states in the model is $K = K_0 \prod_{i=1}^{N} K_i$. Therefore, we will require at least $\lfloor \frac{K-1}{2} \rfloor$ linear restrictions of the form $R\,\text{vec}(Q) = r$ where $R$ has full rank. We proceed by showing that the present model admits an intensity matrix

$Q$ with a known sparsity pattern and so we can use the locations of zeros as homogeneous restrictions, where $r$ will be a vector of zeros.

Recall that each player $i$ has $J_i + 1$ choices, but since one is a continuation choice this results in $J_i$ non-zero off-diagonal elements per row of $Q$. There are (at most) $K_0 - 1$ non-zero off-diagonal elements due to exogenous state changes by nature. The only other non-zero elements of each row are the diagonal elements and therefore there are $K - \sum_i J_i - (K_0 - 1) - 1 = K_0 \prod_i K_i - \sum_i J_i - K_0$ zeros per row of $Q$. The order condition we need to show is that the *total* number of zero restrictions is at least $\lfloor (K-1)/2 \rfloor$. For simplicity, it will suffice to show that there are $K/2 \geq \lfloor (K-1)/2 \rfloor$ restrictions. Summing across rows, this condition is satisfied when $(K_0 \prod_i K_i)(K_0 \prod_i K_i - K_0 - \sum_i J_i) \geq K_0 \prod_i K_i / 2$. Simplifying yields the sufficient condition in (13).

To see how this condition changes with $K_0$, $K_i$, and $J_i$, we take derivatives. The derivative with respect to $K_0$ is $\prod_i K_i - 1 \geq 0$. This value is always non-negative, since $K_i \geq 1$ for all $i$, and is strictly positive when $K_i > 1$ for any $i$. The derivative with respect to $K_i$ for $i = 1, \ldots, N$ is $K_0 \prod_{m \neq i} K_m > 0$. This value is always strictly positive since $K_0 \geq 1$ and $K_i \geq 1$ for all $i$. Finally, the derivative with respect to $J_i$ is $-1$. ■

| $M$ | $N$ | | $q_1$ | $q_{21}$ | $q_{22}$ | $\pi$ | $\lambda$ | $\beta$ | $c_1$ | $c_2$ |
|---|---|---|---|---|---|---|---|---|---|---|
| 25 | 100 | Bias | 0.000 | 0.001 | 0.001 | -0.023 | 0.001 | 0.040 | -0.014 | 0.003 |
| | | RMSE | 0.006 | 0.004 | 0.005 | 0.118 | 0.005 | 0.306 | 0.112 | 0.255 |
| 50 | 100 | Bias | 0.000 | 0.000 | 0.000 | -0.007 | 0.001 | 0.045 | -0.005 | 0.000 |
| | | RMSE | 0.004 | 0.003 | 0.004 | 0.070 | 0.004 | 0.194 | 0.067 | 0.141 |
| 100 | 100 | Bias | 0.000 | 0.000 | 0.000 | -0.115 | 0.001 | 0.023 | -0.006 | -0.006 |
| | | RMSE | 0.003 | 0.002 | 0.002 | 0.059 | 0.003 | 0.139 | 0.049 | 0.107 |
| 25 | 200 | Bias | 0.000 | 0.000 | 0.000 | -0.015 | 0.000 | 0.025 | 0.004 | 0.002 |
| | | RMSE | 0.003 | 0.003 | 0.003 | 0.093 | 0.004 | 0.177 | 0.061 | 0.118 |
| 50 | 200 | Bias | 0.000 | 0.000 | 0.000 | -0.006 | 0.000 | 0.033 | 0.009 | 0.008 |
| | | RMSE | 0.003 | 0.002 | 0.002 | 0.074 | 0.003 | 0.140 | 0.042 | 0.102 |
| 100 | 200 | Bias | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.014 | 0.002 | -0.005 |
| | | RMSE | 0.002 | 0.001 | 0.002 | 0.047 | 0.002 | 0.097 | 0.029 | 0.062 |

TABLE 3. Single player Monte Carlo results with unobserved heterogeneity.

| $N$ | $M$ | $K$ | | $\lambda$ | $\gamma$ | $\kappa$ | $\eta$ | $\varphi$ |
|---|---|---|---|---|---|---|---|---|
| 4 | 0.5 | 840 | Bias | 0.008 | 0.001 | 0.002 | 0.007 | 0.007 |
| | | | RMSE | 0.012 | 0.002 | 0.016 | 0.099 | 0.118 |
| 6 | 1.0 | 5,544 | Bias | 0.009 | 0.001 | -0.002 | 0.000 | 0.007 |
| | | | RMSE | 0.013 | 0.002 | 0.016 | 0.056 | 0.062 |
| 8 | 3.0 | 24,024 | Bias | 0.008 | 0.001 | -0.002 | 0.005 | 0.011 |
| | | | RMSE | 0.012 | 0.002 | 0.015 | 0.068 | 0.055 |
| 10 | 5.0 | 80,080 | Bias | 0.009 | 0.001 | -0.004 | 0.005 | 0.001 |
| | | | RMSE | 0.013 | 0.002 | 0.019 | 0.080 | 0.055 |
| 12 | 8.0 | 222,768 | Bias | 0.008 | 0.000 | -0.001 | 0.029 | -0.004 |
| | | | RMSE | 0.012 | 0.002 | 0.022 | 0.080 | 0.059 |
| 14 | 10.0 | 542,640 | Bias | 0.009 | 0.001 | 0.002 | -0.007 | 0.009 |
| | | | RMSE | 0.013 | 0.002 | 0.026 | 0.067 | 0.050 |
| 16 | 13.0 | 1,193,808 | Bias | 0.009 | 0.001 | 0.006 | 0.001 | -0.001 |
| | | | RMSE | 0.013 | 0.002 | 0.030 | 0.076 | 0.043 |
| 18 | 17.0 | 2,422,728 | Bias | 0.008 | 0.001 | -0.002 | 0.007 | 0.009 |
| | | | RMSE | 0.012 | 0.002 | 0.029 | 0.094 | 0.050 |
| 20 | 21.0 | 4,604,600 | Bias | 0.008 | 0.001 | 0.003 | 0.000 | 0.000 |
| | | | RMSE | 0.012 | 0.002 | 0.037 | 0.112 | 0.051 |
| 22 | 25.0 | 8,288,280 | Bias | 0.008 | 0.000 | 0.006 | 0.030 | 0.019 |
| | | | RMSE | 0.012 | 0.002 | 0.039 | 0.125 | 0.063 |
| 24 | 29.0 | 14,250,600 | Bias | 0.008 | 0.001 | 0.002 | 0.021 | 0.016 |
| | | | RMSE | 0.013 | 0.002 | 0.039 | 0.120 | 0.055 |
| 26 | 33.0 | 23,560,992 | Bias | 0.009 | 0.001 | 0.000 | 0.051 | 0.033 |
| | | | RMSE | 0.013 | 0.002 | 0.042 | 0.124 | 0.074 |
| 28 | 37.0 | 37,657,312 | Bias | 0.008 | 0.001 | -0.004 | 0.002 | 0.032 |
| | | | RMSE | 0.013 | 0.002 | 0.044 | 0.179 | 0.072 |
| 30 | 41.0 | 58,433,760 | Bias | 0.010 | 0.000 | 0.007 | 0.048 | 0.035 |
| | | | RMSE | 0.014 | 0.002 | 0.045 | 0.165 | 0.081 |

TABLE 4. Quality ladder Monte Carlo results: NFXP estimation.

| $N$ | $M$ | $K$ | | $\lambda$ | $\gamma$ | $\kappa$ | $\eta$ | $\varphi$ |
|---|---|---|---|---|---|---|---|---|
| 4 | 0.5 | 840 | Bias | 0.008 | 0.001 | 0.001 | 0.016 | 0.009 |
| | | | RMSE | 0.012 | 0.002 | 0.017 | 0.091 | 0.109 |
| 6 | 1.0 | 5,544 | Bias | 0.009 | 0.000 | -0.003 | 0.014 | 0.011 |
| | | | RMSE | 0.013 | 0.002 | 0.017 | 0.053 | 0.046 |
| 8 | 3.0 | 24,024 | Bias | 0.008 | 0.000 | -0.006 | 0.034 | 0.015 |
| | | | RMSE | 0.012 | 0.002 | 0.017 | 0.086 | 0.041 |
| 10 | 5.0 | 80,080 | Bias | 0.009 | 0.001 | -0.003 | 0.044 | 0.008 |
| | | | RMSE | 0.013 | 0.002 | 0.019 | 0.103 | 0.039 |
| 12 | 8.0 | 222,768 | Bias | 0.009 | 0.000 | 0.005 | 0.078 | 0.002 |
| | | | RMSE | 0.012 | 0.002 | 0.027 | 0.131 | 0.039 |
| 14 | 10.0 | 542,640 | Bias | 0.009 | 0.000 | 0.018 | 0.055 | 0.008 |
| | | | RMSE | 0.013 | 0.002 | 0.042 | 0.110 | 0.034 |
| 16 | 13.0 | 1,193,808 | Bias | 0.009 | 0.001 | 0.028 | 0.070 | 0.003 |
| | | | RMSE | 0.013 | 0.002 | 0.054 | 0.135 | 0.028 |
| 18 | 17.0 | 2,422,728 | Bias | 0.008 | 0.000 | 0.023 | 0.065 | 0.002 |
| | | | RMSE | 0.012 | 0.002 | 0.052 | 0.138 | 0.032 |
| 20 | 21.0 | 4,604,600 | Bias | 0.008 | 0.000 | 0.036 | 0.071 | -0.006 |
| | | | RMSE | 0.012 | 0.002 | 0.075 | 0.160 | 0.035 |
| 22 | 25.0 | 8,288,280 | Bias | 0.008 | 0.000 | 0.043 | 0.094 | 0.000 |
| | | | RMSE | 0.013 | 0.002 | 0.084 | 0.168 | 0.044 |
| 24 | 29.0 | 14,250,600 | Bias | 0.008 | 0.000 | 0.043 | 0.082 | -0.005 |
| | | | RMSE | 0.014 | 0.002 | 0.090 | 0.147 | 0.043 |
| 26 | 33.0 | 23,560,992 | Bias | 0.009 | 0.000 | 0.045 | 0.104 | 0.004 |
| | | | RMSE | 0.013 | 0.002 | 0.097 | 0.173 | 0.059 |
| 28 | 37.0 | 37,657,312 | Bias | 0.008 | 0.000 | 0.040 | 0.052 | 0.000 |
| | | | RMSE | 0.013 | 0.002 | 0.100 | 0.197 | 0.059 |
| 30 | 41.0 | 58,433,760 | Bias | 0.010 | 0.000 | 0.061 | 0.113 | 0.001 |
| | | | RMSE | 0.014 | 0.002 | 0.113 | 0.215 | 0.066 |

TABLE 5. Quality ladder Monte Carlo results: CCP estimation.

| $N$ | $M$ | $K$ | $n_{\text{avg}}$ | $\omega_{\text{avg}}$ | Obtain $V$ | NFXP | CCP |
|---|---|---|---|---|---|---|---|
| 4 | 0.5 | 840 | 3.94 | 4.07 | $< 1$ | 12 | $< 1$ |
| 6 | 1.0 | 5,544 | 4.88 | 4.17 | 1 | 47 | $< 1$ |
| 8 | 3.0 | 24,024 | 6.35 | 4.88 | 6 | 149 | $< 1$ |
| 10 | 5.0 | 80,080 | 7.86 | 4.92 | 13 | 262 | $< 1$ |
| 12 | 8.0 | 222,768 | 9.94 | 4.75 | 34 | 503 | 1 |
| 14 | 10.0 | 542,640 | 11.42 | 4.46 | 57 | 1,040 | 2 |
| 16 | 13.0 | 1,193,808 | 13.25 | 4.20 | 112 | 1,214 | 3 |
| 18 | 17.0 | 2,422,728 | 15.24 | 3.94 | 220 | 1,643 | 7 |
| 20 | 21.0 | 4,604,600 | 17.17 | 3.77 | 415 | 2,678 | 13 |
| 22 | 25.0 | 8,288,280 | 18.69 | 3.49 | 766 | 5,449 | 23 |
| 24 | 29.0 | 14,250,600 | 20.62 | 3.31 | 1,239 | 5,956 | 43 |
| 26 | 33.0 | 23,560,992 | 22.14 | 3.23 | 2,350 | 9,496 | 69 |
| 28 | 37.0 | 37,657,312 | 24.19 | 3.01 | 3,576 | 14,910 | 110 |
| 30 | 41.0 | 58,433,760 | 25.96 | 2.86 | 5,443 | 23,302 | 172 |

Times are total seconds of CPU time on a dual processor Intel Xeon X5670 system for a single replication of each specification. $N$ denotes the total possible number of players, $M$ denotes the market size, and $K$ denotes the total number of distinct states. For all experiments, the number of quality levels is $\bar{\omega} = 7$ and the discount rate at $\rho = 0.05$.

TABLE 6. Quality ladder model Monte Carlo results: CPU Time (in seconds).