Identification and Estimation of Continuous Time Dynamic Discrete Choice Games*

JASON R. BLEVINS

The Ohio State University

November 2, 2016
Preliminary and Incomplete Draft

Abstract. We consider the theoretical and econometric properties of a recently proposed class of continuous time dynamic games with discrete choices. This class of models has computationally-desirable properties and includes widely-used entry, exit, and investment games that underlie many applications in empirical industrial organization. First, we provide more general conditions for existence of equilibrium which allow move arrival rates to differ across firms and model states. Next, we return to the question of nonparametric identification of the intensity matrix and the structural primitives of the model when only discrete-time data are available. We show that the conditions for identification are easily verifiable in representative example models. Finally, we examine both the computational properties of the model and the finite sample performance of a two-step semiparametric estimator for the model through a series of Monte Carlo experiments in a simple single agent renewal model and a dynamic oligopoly quality ladder model.

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

JEL Classification: C13, C35, C62, C73.

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

*I am grateful to seminar participants at Columbia University, Northwestern University, 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 joint research with Peter Arcidiacono, Patrick Bayer, and Paul Ellickson and has benefited tremendously from our many discussions together.

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, provide new identification results, 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.

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) advocated 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; Pesendorfer and Schmidt-Dengler, 2007; Pakes, Ostrovsky, and Berry, 2007) 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 very 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. It has subsequently been used in several other 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.

For completeness, we briefly review the ABBE model in Section 2. We then extend their equilibrium existence result to cases where the move arrival rates may differ across firms and states. We investigate additional aspects of identification and estimation of the ABBE model. 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.

Finally, we also 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. In solving for an equilibrium, relative to a discrete time game, the computing time is many orders of magnitude smaller. For the continuous time model even, the proposed two-step estimator is again orders of magnitude faster than the full solution estimator.

2. A Stochastically Sequential Continuous Time Dynamic Discrete Game

2.1. Markov Jump Processes Representation

We consider infinite horizon games in continuous time with N players indexed by $i = 1, \ldots, N$. Before discussing the full structural model, we first focus on certain features of the reduced form 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 τ before transitioning to some other state $X_{t+\tau}$. The length of time τ 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 τ_n, with n denoting the n-th jump.

Jumps occur according to a Poisson process and the holding times between jumps are therefore exponentially distributed. The probability density function (pdf) of the

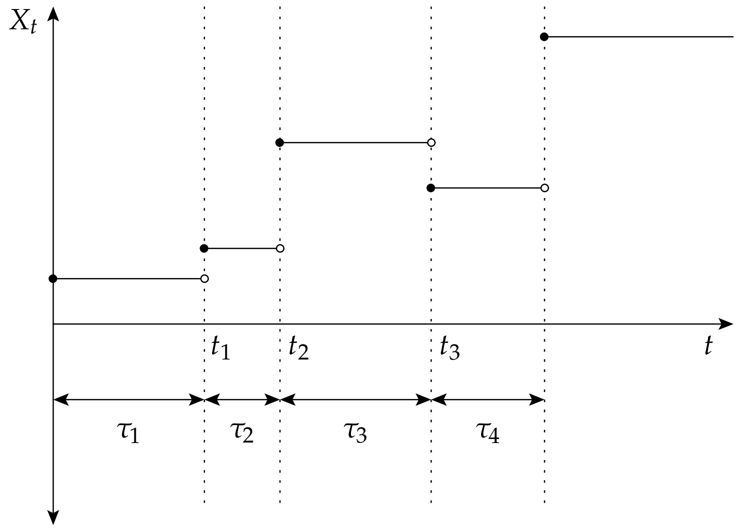

FIGURE 1. Markov jump process

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

exponential distribution with rate parameter $\lambda > 0$ is

$$g(\tau; \lambda) = \begin{cases} \lambda \, \mathrm{e}^{-\lambda \tau}, & \tau \geq 0, \\ 0, & \tau < 0, \end{cases}$$

and the cumulative distribution function (cdf) is

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

The mean of this distribution is $1/\lambda$, the inverse of the rate parameter, λ, and the variance is $1/\lambda^2$.

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

Finally, 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 Δ given that it was initially in state k. Let $P(\Delta) = (P_{kl}(\Delta))$ denote the corresponding $K \times K$ matrix of these probabilities, which we refer to as the *transition matrix*. $P(\Delta)$ is the unique solution to the system of ordinary differential equations

$$(1) \quad P'(\Delta) = QP(\Delta), \quad P(0) = I,$$

frequently referred to as the forward equations. It follows that

$$(2) \quad P(\Delta) = \exp(\Delta Q) = \sum_{j=0}^{\infty} \frac{(\Delta Q)^j}{j!}.$$

This quantity is the matrix exponential, the matrix analog of the scalar exponential $\exp(x)$ for $x \in \mathbb{R}$.

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 λ_i for $i = 1, \dots, N$, then distribution of the minimum wait time 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 1. *Suppose $\tau_i \sim \text{Expo}(\lambda_i)$, for $i = 1, \dots, N$, are independent and define $\tau \equiv \min_i \tau_i$ and $\iota \equiv \arg\min_i \tau_i$. Then*

$$\tau \sim \text{Expo}(\lambda_1 + \cdots + \lambda_N)$$

and

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

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 (τ, ι) 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, λ_{ik}. Conditional on player i making some action, the probability that action j is made is h_{ijk}/λ_{ik}.

In the notation of the lemma, we can let ι denote a player-action pair (i, j) and τ denote the holding time until the next action by any player occurs in state k. The distribution of τ 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: $Q = Q_0 + Q_1 + \ldots, Q_N$. We now introduce two examples which illustrate how these individual intensity matrices are determined. We will return to these examples throughout the paper to illustrate other assumptions and results.

Example 1 (Single-Agent Renewal)**.** Consider the restrictions on the Q matrix implied by a continuous-time version of the single-agent renewal model of Rust (1987) as described by ABBE. 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 γ_1 and γ_2:

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

At at rate λ (assumed to be constant across states), the manager of the bus company decides whether or not to replace the engine. Let σ_{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 τ_{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 ($2 \times 2 \times 2$ Entry)**.** 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 $x_3 \in \{\mathrm{L}, \mathrm{H}\}$. Therefore, in vector form the state space is

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

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

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

The state transition hazards can be characterized by an 8×8 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

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

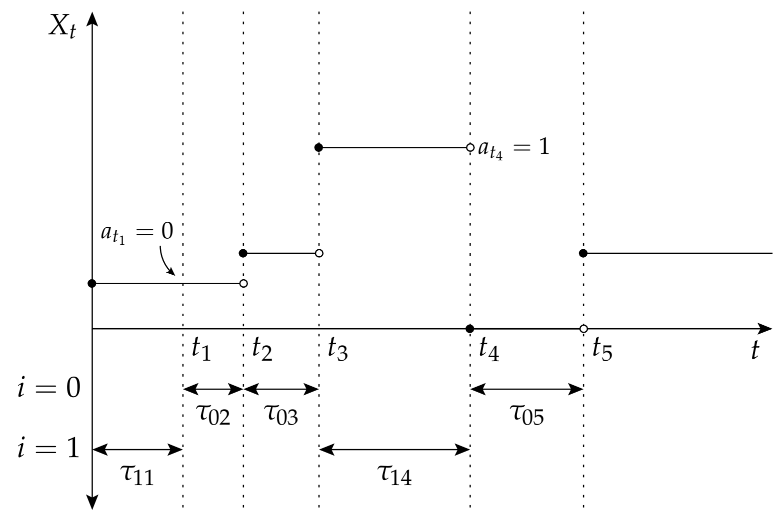

FIGURE 2. Single agent model: a representative sample path where t_n, τ_{in}, and a_{in} denote, respectively, the time, inter-arrival time, and action corresponding to n-th event. Moves by the agent are denoted by $i = 1$ while $i = 0$ denotes a state change (a move by nature).

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.

2.2. *State Changes and Player Actions*

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

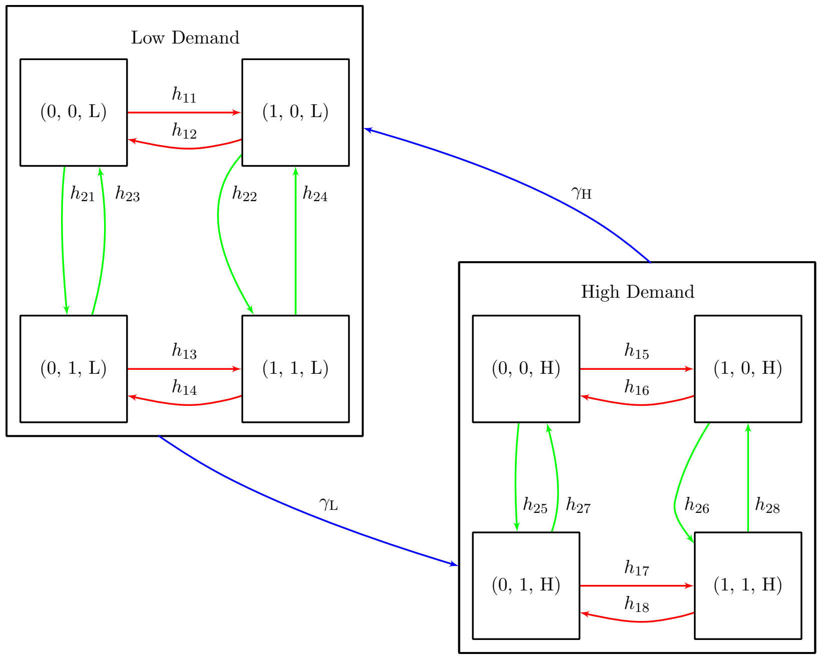

FIGURE 3. Two Player Entry Game with Exogenous Demand State

of length Δ 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 $P(\Delta)$. This matrix will be used to relate the model to the data for estimation.

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 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 λ_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 λ_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 $\tilde{q}_{kk} = -\lambda_k + \lambda_k \sigma_{kk} = -\sum_{l \neq k} q_{kl} = q_{kk}$.

Now, suppose Q characterizes state changes resulting from the actions of a single player. To take the model literally, by choosing $\lambda_1, \ldots, \lambda_K$ we are in effect choosing the rate at which that player makes choices. However, by including in the choice set the possibility of inaction this is an innocuous normalization that allows one to give the model a random utility model interpretation. There is a close parallel in differentiated products demand models where the researcher chooses the market size (analogous to the rate λ_k) but allows individuals to choose an "outside good" instead of an actual product (analogous to allowing players to choose to do nothing).

2.4. *Structural Model and Basic Assumptions*

Now we formalize the underlying structural model which gives rise to the reduced form finite Markov jump process described above. We consider infinite horizon games in continuous time indexed by $t \in [0, \infty)$ with N players indexed by $i = 1, \ldots, N$.

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 $\mathcal{X} \subset \mathbb{R}^L$ with $K = |\mathcal{X}| < \infty$. The states $x \in \mathcal{X}$ are vectors of L real numbers. For example, 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.

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 λ_{ik}. That is, at any point in time and in any state k the holding time τ_i until the next move by player i is exponentially distributed with rate parameter λ_{ik}. We assume these processes are independent across players, that the rates λ_{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. As we argued above, choosing the rate of decisions λ_{ik} is essentially a normalization since the probability of inaction, denoted $j = 0$, will adjust accordingly.

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.

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 ε_{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 ψ_{ijk}. Importantly, the choice-specific payoffs can depend on the market structure because they are indexed by k.

Example 1, Continued. In the renewal model there is a single agent, so $N = 1$, and a binary choice, so the choice set is $\mathcal{J} = \{0, 1\}$. 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 ε_{ijk}.

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 intensity matrix for nature, $Q_0 = (q_{0kl})$, is therefore defined as part of the model. The individual elements q_{0kl} for $k \neq l$ are usually specified by the researcher to depend on a lower-dimensional vector of parameters θ.

Intensity Matrices Note that for each player i the move arrival rates, λ_{ik}, the choice probabilities, σ_{ijk}, and the continuation state function $l(i, \cdot, \cdot)$ imply a Markov jump process on $\mathcal{X}$ with intensity matrix Q_i. Q_i is a function of both λ_{ik} and σ_{ijk} for all j and k. In particular, the rate at which player i plays action j in state k is simply $\lambda\sigma_{ijk}$—the product of the move arrival rate and the state-specific choice probability. Moves by player i in state k occur at rate λ_{ik} and the choice probability σ_{ijk} is thus the proportion of those moves that result in action j. In turn, such moves result in immediate state transitions from k to $l(i, j, k)$. This means that for $j > 0$, the $(k, l(i, j, k))$ entry of Q_i is $\lambda_{ik}\sigma_{ijk}$. The (k, k) diagonal element equals $-\lambda_{ik}(1 - \sigma_{i0k})$, where $1 - \sigma_{i0k}$ is the probability that some action $j > 0$ is played and therefore the probability of leaving state k as a result of an action by player i.

We will use the following convention for working in matrix and vector notation: for scalars indexed by i, j, and k, such as the choice probabilities σ_{ijk}, let

$$\sigma_{ij} = \begin{bmatrix} \sigma_{i,j,1} \\ \vdots \\ \sigma_{i,j,K} \end{bmatrix}, \quad \sigma_i = \begin{bmatrix} \sigma_{i,0} \\ \vdots \\ \sigma_{i,J} \end{bmatrix}, \quad \text{and} \quad \sigma = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_N \end{bmatrix}.$$

Occasionally, we will also use vectors indexed by i and k, in which case

$$\sigma_{ik} = \begin{bmatrix} \sigma_{i,0,k} \\ \vdots \\ \sigma_{i,J,k} \end{bmatrix}.$$

To avoid confusion, we reserve k for indexing states.

Aggregate Dynamics Summing the individual intensity matrices yields the aggregate intensity matrix $Q = \sum_{i=0}^{N} Q_i$ of the compound process. The matrix Q fully characterizes the state transition dynamics of the model for given values of nature's intensity matrix Q_0, the rates of player actions λ, and the choice probabilities σ. The choice probabilities will be determined in equilibrium, but before defining the equilibrium notion we use we must first define payoffs and characterize the dynamic discrete choice problem faced by the players.

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 $i = 1, \dots, N$.

Assumption 3 (Move Arrival Rates)**.** For $i = 1, \dots, N$, $k = 1, \dots, K$, and $l = 1, \dots, K$ the rates of move arrivals and exogenous state changes satisfy $0 < \lambda_{ik} < \infty$ and $0 \leq q_{0kl} < \infty$.

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

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

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

(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 $j = 0, \dots, J$ and $j' \neq j$.

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.

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

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

[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.

ε_{ik} an action from $\mathcal{J}$. For a given policy function δ_i, we can define the conditional choice probabilities

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

for all choices j and states k. Let ς_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 ς_{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 ς_i. For given beliefs ς_i, the optimal policy rule for player i satisfies the following inequality condition:

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

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

Given beliefs ς_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 ς_i. For small time increments h, under the Poisson assumption, the probability of an event with rate λ_{ik} occurring is $\lambda_{ik}h$. Given the discount rate ρ_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 ς_i for brevity):

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

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

(5)
$$V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{0kl} V_{il} + \sum_{m \neq i} \lambda_{mk} \sum_{j=0}^{J} \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_{ik} \operatorname{E} \max_j \{\psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)}\}}{\rho_i + \sum_{l \neq k} q_{0kl} + \sum_{i=1}^{N} \lambda_{ik}}.$$

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 $\varepsilon_{ik} = (\varepsilon_{i0k}, \dots, \varepsilon_{iJk})^\top$. Alternatively, and perhaps more intuitively, rearranging once again shows that the instantaneous discounted increment to the value function V_{ik} for given beliefs ς_i is

$$\rho V_{ik} = u_{ik} + \sum_{l \neq k} q_{0kl} (V_{il} - V_{ik}) + \sum_{m \neq i} \lambda_{mk} \sum_{j=0}^{J} \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_{ik} \operatorname{E} \max_j \{\psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} - V_{ik}\}. \tag{6}$$

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 $N = 1$):

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

Example 2, Continued. In the $2 \times 2 \times 2$ 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 (5) in matrix notation, representing the Bellman operator, which we denote by Γ. Let $\Sigma_i(\sigma_i)$ denote the transition matrix implied by the choice probabilities σ_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 ς_{im}. Let $Q_0 = (q_{0kl})$ denote the matrix of rates of exogenous state transitions and let $\tilde{Q}_0 = Q_0 - \operatorname{diag}(q_{011}, \dots, q_{0KK})$ be the matrix formed by taking Q_0 and replacing the diagonal elements with zeros. Finally, define ς_{ii} to be the best response probabilities under the beliefs in ς_i about rival players.

Then, we can write the Bellman operator Γ_i (for given beliefs ς_i) as

$$\Gamma_i(V_i) = D_i \left[u_i + \tilde{Q}_0 V_i + \sum_{m \neq i} L_m \Sigma_m(\varsigma_{im}) V_i + L_i \{\Sigma_i(\varsigma_{ii}) V_i + C_i(\varsigma_{ii})\} \right], \tag{7}$$

where $D_i = (d_{kl})$ is the $K \times K$ diagonal matrix containing the denominator of (5) for each k, hence $d_{kk} = 1/(\rho_i + \sum_{m=1}^{N} \lambda_{mk} + \sum_{l \neq k} q_{0kl})$, $L_m = \operatorname{diag}(\lambda_{m1}, \dots, \lambda_{mK})$ 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 ς_{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 ε_{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 ς_i is a fixed point of Γ_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. For convenience, we restate their result here as follows.

Lemma 2 (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 $j' \in \mathcal{J}$:*

$$v_{ijk} \equiv (\psi_{ijk} + V_{i,l(i,j,k)}) - (\psi_{ij'k} - V_{i,l(i,j',k)}). \tag{8}$$

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.

Theorem 1. *If Assumptions 1–7 hold, then for a given collection of beliefs $\{\sigma_i\}_{i=1}^N$, V_i has the following linear representation:*

$$V_i(\sigma_i) = \Xi_i^{-1}(Q) \left[u_i + L_i C_i(\sigma_i)\right] \tag{9}$$

where

$$\Xi_i(Q) = \rho_i I_K + \sum_{m=1}^{N} L_m[I_K - \Sigma_m(\sigma_m)] - Q_0 \tag{10}$$

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

Proof. See Appendix A.

2.5. Markov Perfect Equilibrium

Definition. A *Markov perfect equilibrium* is a collection of stationary policy rules $\{\delta_i\}_{i=1}^N$ such that (4) holds for all i, k, and ε_{ik} given beliefs $\varsigma_i = (\sigma_1, \ldots, \sigma_N)$ generated by (3).

The following theorem establishes existence of a Markov perfect equilibrium. ABBE proved that an equilibrium exists under the assumption players share a common move arrival and discount rates and that the move arrival rates do not not vary across states (i.e., $\lambda_{ik} = \lambda$ for all i and k). The following theorem extends their existence result to our more general setting.

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

Proof. See Appendix A.

2.6. Example: A Quality Ladder Model

To illustrate the application to dynamic games we consider a discrete control version of the quality ladder model proposed by Ericson and Pakes (1995). This model is widely used in industrial organization and has been examined extensively by Pakes and McGuire (1994, 2001), Doraszelski and Satterthwaite (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 ω_i is the share of consumers who are aware of firm i's product. Gowrisankaran (1999a) develops a model of endogenous horizontal mergers where ω_i is a capacity level and the product market stage game is Cournot with a given demand curve and cost functions that enforce capacity constraints depending on each firm's ω_i.

2.6.1. 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 $\mathcal{X} = \{1, \ldots, |\mathcal{S}|\}$.[2] Payoff relevant market configurations from the perspective of firm i are then uniquely described by two integers (x, ω_i), where $x \in \mathcal{X}$ denotes the market structure and ω_i is firm i's own quality level.

2.6.2. 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 ε_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. As in Pakes, Gowrisankaran, and McGuire (1993), for some constant ω^* 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 ω 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

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

[2] 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 $k \in \mathcal{X}$.

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 (ω_i, ω_{-i}).

2.6.3. Incumbent Firms

We consider a simple model in which incumbent firms have three choices upon receiving a move arrival. Firms may continue without investing at no cost, they may invest an amount κ in order to increase the quality of their product from ω_i to $\omega_i' = \min\{\omega_i + 1, \bar{\omega}\}$, or they may exit the market and receive some scrap value φ. We denote these choices, respectively, by the choice set $\mathcal{A}_i = \{0, 1, 2\}$. When an incumbent firm exits the market, ω_i jumps deterministically to $\bar{\omega} + 1$. Associated with each choice j is a private shock ε_{ijt}. These shocks are iid over firms, choices, and time and 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 Ω^N. In the general notation introduced above, the instantaneous payoff ψ_{ijk} to firm i from choosing choice j in state k is

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

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

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

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

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

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

2.6.4. *Potential Entrants*

Whenever the number of incumbents is smaller than N, a single potential entrant receives the opportunity to enter at rate λ. 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 η 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 ε^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

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

2.6.5. State Transitions

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

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

We construct the corresponding intensity matrix Q_0 as follows. We map each market structure s to an integer k and map the resulting structure after depreciation s' to an integer k'. The (k, k) element of Q_0 for each eligible state k is $-\gamma$ while the corresponding (k, k') element is γ.

2.7. 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., $\Delta = 1$). We will consider two examples: a simple entry-exit model and a model where chains compete by opening and closing stores.

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 $\mathcal{J} = \{-J, \dots, J\}$. 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 $\bar{\lambda}$ 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* $1/\bar{\lambda}$ 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 $1/\bar{\lambda}$ 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 $\mathcal{J} = \{-J, \dots, J\}$ for some J which implies on average at most $J/\overline{\lambda}$ openings or closings per year.

In both cases, there is a tradeoff 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

Identification of the model proceeds in two steps. First, we establish that the intensity matrix Q for the continuous time model is identified even with only discrete time data. Second, given Q we show that the structural primitives—the flow payoffs u and instantaneous payoffs ψ—are identified.

The implications of the structural model can be viewed from the bottom up as follows: the structural primitives u and ψ imply value functions V which imply choice probabilities σ. These probabilities and the rates of moves, λ, and state transitions by nature, Q_0, in turn imply a Q matrix. Finally, given the Q matrix for any observation interval Δ, the distribution of observable data is $P(\Delta) = \exp(\Delta Q)$.

The identification problem requires us to work from the top down. The first step of our identification approach addresses the following question: given knowledge of $P(\Delta)$, under what conditions can we uniquely determine the matrix Q? We show that this is possible under very mild conditions on the structure of the game. Then, in the second step we show that knowledge of Q allows us to recover the structural primitives under conditions that are similar to assumptions that are widely-used for discrete time models.

Much of the previous work related to this identification problem seeks conditions on the observable discrete-time transition matrix $P(\Delta)$. We show here that one can exploit the restrictions placed on the intensity matrix by the underlying structural model to obtain weaker conditions for identification that involve more primitive quantities of the model. The conditions for identification of Q are based on a priori restrictions on the Q matrix that

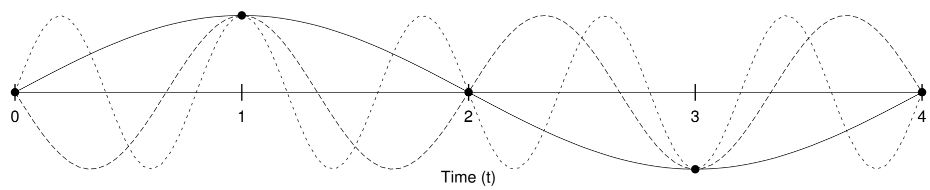

FIGURE 4. The Aliasing Problem in the Frequency Domain

Plots of $f(t; m) = \sin(2\pi m t)$ for $t \in [0, 4]$ and $m = 1/4$ (solid), $m = -3/4$ (long dash), and $m = 5/4$ (short dash).

easily arise from the construction of the model itself. These can be viewed as exclusion restrictions.

Examples of such restrictions are easy to find. For example, if there are player-specific components of the state vector then player i cannot change the state of player j and vice-versa. In an entry-exit model, such a state is incumbency status: players can only exit by their own action; no player can exit 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 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 $P(\Delta)$.

Once the Q matrix is known, so are the matrices Q_i for all $i = 0, \dots, N$ and therefore so are the conditional choice probabilities. Identification of the structural payoffs then proceeds again using the structure of the model (e.g., payoff exclusion restrictions).

3.1. Identification of Q with Discrete Time Data

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. This is known as the *aliasing problem*, which has been studied by many authors in the context of continuous-time multivariate regression models (Sims, 1971; Phillips, 1973; Hansen and Sargent, 1983; Geweke, 1978). The fundamental problem is most obvious in the frequency domain: Figure 4 depicts three sine waves that are observationally equivalent when sampled at intervals of length $\Delta = 1$ but which have distinct frequencies.

In the present model, the concern is that there may be multiple Q matrices which give rise to the same potentially observable transition probability matrix $P(\Delta)$. That is, we seek to establish that if $\exp(\Delta Q) = \exp(\Delta \tilde{Q})$, then $Q = \tilde{Q}$. Below, we discuss this in detail and give formal conditions under which this is the case.

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 δ). 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 $P(\Delta)$. These conditions could provide foundations for model specification tests.

3.1.1. Identification of Discrete Time Models with Discrete Time Data

In addition to the implicit assumptions of discrete time models discussed in Subsubsection 2.7, there is an analogous identification issue for discrete-time models which is hidden by the usual assumption that the frequency of moves is known and equal to the sampling frequency. Suppose to the contrary that there is a fixed move interval of length δ in the model which may be different from the fixed interval Δ at which observations are sampled. In practice, researchers typically assume (implicitly) that $\delta = \Delta$, where $\Delta = 1$ is normalized to be some specific unit of time (e.g., one quarter). This assumption is convenient, but masks the identification problem, which requires that there exist a unique matrix root P_0 of the discrete-time aggregation equation $P_0^{\Delta/\delta} = P(\Delta)$. In general there may be multiple such matrices (Gantmacher, 1959; Singer and Spilerman, 1976), but P_0 is trivially unique under the usual assumption that $\delta = \Delta$.

3.1.2. Identification of Unrestricted Q Matrices

Returning to identification of Q in the context of finite-state Markov jump processes, 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 Δ. 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.

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 (modulo the requirements to 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).[3] 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 Δ (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 Δ_1 and Δ_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.

3.1.3. 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.

[3]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.

Our proposed identification strategy is inspired by this work on multivariate regression models, but the ABBE 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 intensity matrix Q in the ABBE model is analogous to the matrix A above and the ABBE model can also be expressed as a system of differential equations as in (1). If Q is an intensity matrix with $q_{kl} \geq 0$ for all $l \neq k$, $q_{kk} \leq 0$, and $\sum_{k=1}^{K} q_{kl} = 0$ for $k = 1, \ldots, K$, 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 ABBE 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. Recall the single-agent renewal model where the aggregate intensity matrix is

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

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 conditional choice probabilities depend on the rates and payoff parameters which introduces shape restrictions on σ_{1k} across states k.

Example 2, Continued. In the $2 \times 2 \times 2$ entry example, the aggregate intensity matrix is

$Q = Q_0 + Q_1 + Q_2$:

$$(12) \quad Q = \left[\begin{array}{cccc|cccc} \cdot & h_{11} & h_{21} & 0 & \gamma_{\mathrm{L}} & 0 & 0 & 0 \\ h_{12} & \cdot & 0 & h_{22} & 0 & \gamma_{\mathrm{L}} & 0 & 0 \\ h_{23} & 0 & \cdot & h_{13} & 0 & 0 & \gamma_{\mathrm{L}} & 0 \\ 0 & h_{24} & h_{14} & \cdot & 0 & 0 & 0 & \gamma_{\mathrm{L}} \\ \hline \gamma_{\mathrm{H}} & 0 & 0 & 0 & \cdot & h_{15} & h_{25} & 0 \\ 0 & \gamma_{\mathrm{H}} & 0 & 0 & h_{16} & \cdot & 0 & h_{26} \\ 0 & 0 & \gamma_{\mathrm{H}} & 0 & h_{27} & 0 & \cdot & h_{17} \\ 0 & 0 & 0 & \gamma_{\mathrm{H}} & 0 & h_{28} & h_{18} & \cdot \end{array}\right],$$

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.

As in these two examples, similar sparse structures arise in even models with large numbers of firms and millions of states, as in the application of ABBE. 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 $R \operatorname{vec}(Q) = r$. For the $K \times K$ matrix $Q = (q_{kl})$, $\operatorname{vec}(Q)$ is the vector obtained by stacking the columns of Q: $\operatorname{vec}(Q) = (q_{11}, q_{21}, \dots, q_{K1}, \dots, q_{1K}, \dots, q_{KK})^\top$.

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 $R \operatorname{vec}(\tilde{Q}) = r$. By the relationship between Q and $\tilde{Q}$ above, we also have $R \operatorname{vec}(Q + UDU^{-1}) = r$. But $R \operatorname{vec}(Q) = r$ and by linearity of the vectorization operator,

$$R \operatorname{vec}(UDU^{-1}) = 0.$$

An equivalent representation is

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

Blevins (2017) shows that when there are at least $\lfloor \frac{K-1}{2} \rfloor$ linear restrictions and R has full rank, then D must be generically zero and therefore the eigenvalues of $\tilde{Q}$ and Q are equal. Since the eigenvectors and all eigenvalues of $\tilde{Q}$ are the same as those of Q, the matrices must be equal and so 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.

Theorem 3. *Suppose the state vector can be written as $x = (x_1, \dots, x_N)$, where each component $x_i \in \mathcal{X}_i$ can only be affected by the action of player i. Suppose that there are $|\mathcal{X}_i| = \kappa$ possible states per player, $J + 1$ actions per player, and $N \geq 2$ players. Suppose that Q has distinct eigenvalues that do not differ by an integer multiple of $2\pi i/\Delta$. Then Q is generically identified*[4] *if the number of choices is no larger than $\frac{1}{N}(\kappa^N - \lfloor(\kappa^N - 1)/2\rfloor - 1)$. A sufficient condition is $J \leq (\kappa^N - 1)/(2N)$.*

Proof. See Appendix A.

To better understand the conclusion of this theorem, we can consider some common cases. First, any non-trivial binary choice game is identified. That is, if $J = 1$ and $\kappa \geq 2$ then Q is identified for any value of N. Similarly, if $J = 2$ (three choices) and $\kappa \geq 3$ (at least three states per player), then Q is identified for any N. Finally, the sufficient condition implies that for any $J \geq 1$ and any $\kappa > 1$, the model is identified for some sufficiently large value of N. Alternatively, for any $N \geq 1$ and $\kappa > 1$, the model is identified as long as J is not too large, where the upper bound on J is exponentially increasing in N. This theorem extends readily to cases with exogenous state variables and where the number of endogenous states per player, κ, may differ across players.

3.1.4. *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 $i = 0, \dots, N$.

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

[4] Generic identification means that Q is identified except possibly for a set of Q matrices having measure zero. In this case, identification only fails where a certain matrix is rank degenerate, which only happens for a measure zero set of possible population Q matrices.

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 in Examples 1 and 2 introduced above where the Q matrices are given in (11) and (12) respectively. 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,

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

3.2. Identification of the Value Functions

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

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. Then we define X_i to be the matrix

$$X_i = \begin{bmatrix} I_K & 0 & \dots & 0 & S_{i1} - I_K \\ 0 & I_K & \dots & 0 & S_{i2} - I_K \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I_K & S_{iJ} - I_K \end{bmatrix}.$$

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 also has full rank.

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*

$$R_i \begin{bmatrix} \psi_{i1} \\ \vdots \\ \psi_{iJ} \\ V_i \end{bmatrix} = r_i$$

and the matrix $\begin{bmatrix} X_i \\ R_i \end{bmatrix}$ has rank $(J+1)K$, then $\psi_{i1}, \dots, \psi_{iJ}$ and V_i are identified.

Proof. For all j, j', and k, the differences v_{ijk} defined in (8) are identified by Lemma 2. Under Assumption 6, for $j' = 0$ these differences can be written as

$$v_{ijk} = \psi_{ijk} + V_{i,l(i,j,k)} - V_{ik}.$$

Collecting these equations all j and k, we can write them in matrix form as

$$\begin{bmatrix} I_K & 0 & \dots & 0 & S_{i1} - I_K \\ 0 & I_K & \dots & 0 & S_{i2} - I_K \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & I_K & S_{iJ} - I_K \end{bmatrix} \begin{bmatrix} \psi_{i,1} \\ \vdots \\ \psi_{i,J} \\ V_i \end{bmatrix} = X_i \begin{bmatrix} \psi_{i,1} \\ \vdots \\ \psi_{i,J} \\ V_i \end{bmatrix} = v_i.$$

Notice that this is a linear system of equations with JK equations and $(J+1)K$ unknowns. We have at least K additional equations in the form of the restrictions R_i. Writing the augmented system, we have

$$\begin{bmatrix} X_i \\ R_i \end{bmatrix} \begin{bmatrix} \psi_i \\ V_i \end{bmatrix} = \begin{bmatrix} v_i \\ r_i \end{bmatrix}.$$

Under the assumptions, there is a unique solution for the unknown vectors ψ_i and V_i. ■

First, we note that only K restrictions for each player are needed (since X_i has JK rows) and therefore the total number of restrictions required is linear in N. This is in contrast to the number of required restrictions in discrete time models, which is exponential in N (Pesendorfer and Schmidt-Dengler, 2007).

Examples of such restrictions are easy to find in most applications. They include states where the value function is known, for example, if $V_{ik} = 0$ when a firm has 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.

Example 1, Continued. In the renewal model, since the replacement cost does not depend on the mileage state we have $\psi_{11k} = c$ for all k. This alone yields $K-1$ restrictions of full rank of the form $\psi_{11k} - \psi_{111} = 0$ for all k.

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. If we further assume symmetry, as is common in applied work, we would have $2K-1$ linearly independent restrictions.

3.3. Identification of the Payoffs

It remains to identify the K-vector of payoffs u_i for each player i. In light of the linear representation in (9),

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

where Ξ_i is the matrix function defined in (10). Under the maintained assumptions and restrictions, V_i and ψ_i are identified for each player. The choice probabilities σ 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, ψ_i, and Q are identified, then the flow payoffs u_i are identified.*

4. Estimation

Here we briefly review the estimator proposed by ABBE and the requisite assumptions for asymptotic normality. The first assumption requires that the data be generated by a single Markov perfect equilibrium and that all players expect the same equilibrium to be played at all instants both in and out of sample. This is an assumption on the data generating process that is commonly used in estimating discrete time models and will allows us to estimate the model even in the presence of multiple equilibria.

Assumption 9. (a) In each market $m = 1, \dots, M$, 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 $P^0(\Delta) = \exp(\Delta Q^0)$.

Let h denote the vector of all distinct hazards appearing in the matrices $Q_0, Q_1, \dots, Q_N$,

$$h = (q_{012}, q_{013}, \dots, q_{0,K-1,K}, \lambda_1\sigma_{111}, \dots, \lambda_1\sigma_{1Jk}, \dots \lambda_N\sigma_{N11}, \dots, \lambda_N\sigma_{NJK}),$$

and let $\mathcal{H} \subset \mathbb{R}^{K(K-1)+NJK}$ denote the space of all possible vectors h. Given h, the matrix Q can be reconstructed. Since the rates $\lambda_1, \dots, \lambda_N$ are known, the corresponding conditional choice probabilities can also be determined.

Let θ denote the finite-dimensional vector of parameters which determine the flow payoffs u and instantaneous payoffs ψ. Then, let $\Lambda : \Theta \times \mathcal{H} \to \mathcal{H} : (\theta, h) \mapsto \Lambda(\theta, h)$ denote the mapping which, given a value of θ 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 θ and h, we reconstruct Q and σ and obtain new value functions using (9). Given the new value functions, we determine the new choice probabilities σ' 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 Λ 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) Θ 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.

Define the pseudo likelihood function

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

where $P(k, l; \Delta, h)$ denotes the (k, l) element of the transition matrix induced by the hazards h. In the case where the model has a unique equilibrium, we can define the maximum likelihood estimator as

$$\begin{aligned} \hat{\theta}_{\text{MLE}} = & \arg\max_{\theta \in \Theta} L_M(\theta, h) \\ & \text{s.t.} \quad h = \Lambda(\theta, h). \end{aligned}$$

A nested fixed point (NFXP) procedure would obtain the equilibrium value of h for each trial value of θ given by an optimization routine (Rust, 1987). A mathematical programming with equilibrium constraints (MPEC) approach would obtain $\hat{\theta}_{\text{MLE}}$ by solving the constrained optimization problem directly (Su and Judd, 2012).

Since it remains unclear whether there is a unique equilibrium or not, we consider a two-step estimation approach under Assumption 9. Suppose we have a $\sqrt{M}$-consistent first stage M-estimator $\hat{h}$ for h^0. ABBE define the feasible two-step pseudo maximum likelihood estimator $\hat{\theta}$ of θ as

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

and show that under the conditions above the estimator is consistent and asymptotically normal. We examine this estimator as well as the more general computational properties of the model in the next section.

5. Monte Carlo Experiments

5.1. Single-Agent Dynamic Discrete Choice

Here, we generate data according to the simple single player binary choice model of Example 1. The primitives of the model are the payoff (mileage cost) parameter β, the intensity matrix (mileage transition) parameters q_1 and q_2, the reset (engine replacement) cost c, the discount rate ρ, and the move arrival rate λ. We fix $\rho = 0.05$ and focus on estimating $\theta = (\lambda, q_1, q_2, \beta, c)$. We first choose values for θ 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 Δ.

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 θ 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 Δ 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 $\Delta \in \{0.625, 1.25, 2.5, 5.0, 10.0\}$.[5] 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, λ, 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.

[5] 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.

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 (λ, q_1, and q_2) and the parameters of a logistic regression model for the probability of renewal with parameters α. The regressors in our logit model are a constant, the state x, and $\ln x$. Then, we invert the predicted CCPs 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 λ 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 π, 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 β, 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.

5.2. *A Dynamic Discrete Game*

Our second set of Monte Carlo experiments corresponds to the quality ladder model described in Section 2.6. 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_{avg}) grows with the total number of possible players (N). The average quality level of active firms is denoted ω_{avg}. We also report K, the number of states from the perspective of player i—the number of distinct (ω, ω_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 θ. 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 Γ 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 $(0.5, 0.5, 0.5, 6.0, 1.0)$.[6] 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.[7]

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 λ and γ in a first step, but we do not report times because they were less than

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

[7]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.

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 κ, η, and φ via full-solution estimation. For each new trial value of θ, we use the value function at the previous θ as the new starting value for the value function iteration. Finally, we report the time required to estimate the second-stage parameters.[8]

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. 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.[9] 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.

6. Conclusion

This paper presents some additional results on the theoretical and econometric properties of the empirical framework for continuous time dynamic games introduced by Arcidiacono, Bayer, Blevins, and Ellickson (2016) (ABBE). We established equilibrium existence in a more general class of models, we provided additional details and examples regarding nonparametric identification with discrete time data, 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.

[8]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.

[9]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.

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.

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*.

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.

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.

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.

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.

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.

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 (2007). 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.

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.

Su, C.-L. and K. L. Judd (2012). Constrained optimization approaches to estimation of structural models. *Econometrica 80*, 2213–2230.

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.

Sampling	n		q_1	q_2	λ	β	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 Δ denotes the observation interval for discrete-time data. n denotes the average number of observations (continuous-time events or discrete-time intervals) when observing the model on the fixed interval $[0, T]$ with $T = 25,000$.

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

Sampling	n		q_1	q_2	λ	β	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 Δ denotes the observation interval for discrete-time data. The CCPs were estimated in a first step using a bin estimator for continuous-time data and via logistic regression on x and $\ln x$ for estimation with time aggregation. 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.

A. Proofs

Proof of Lemma 1. The result follows directly from the joint distribution function under independence:

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

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

Furthermore,

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

■

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 σ_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).

Then, in matrix form,

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

Collecting 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 on the left hand side is strictly diagonally dominant since the diagonal elements of Q_0 equal the negative of the respective off-diagonal row sums, the elements of each matrix $\Sigma_m(\sigma_m)$ are in $[0,1]$ for all m, and $\rho_i > 0$ by Assumption 2. Therefore, by the Levy-Desplanques theorem, this matrix is nonsingular (Horn and Johnson, 1985, Theorem 6.1.10). Hence, the expression of 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 σ 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 Υ 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. We appeal to Theorem 2 of Blevins (2017) which provides conditions for the generic identification of the intensity matrix of a general finite-state Markov jump process. Under the assumptions on the eigenvalues of Q, we require at least $\lfloor \frac{K-1}{2} \rfloor$ linear restrictions of the form $R\,\mathrm{vec}(Q) = r$ where R has full rank.

In this case, the number of distinct states is $K = \kappa^N$. Furthermore, we will be able to find sufficiently many homogeneous restrictions (i.e., r will be vector of zeros). For each player, there are only J non-zero elements per row of the Q matrix. That is, since there are only J choices, there are at most J possible instantaneous state transitions due to an action of each player. The only other non-zero element per row is the diagonal element and therefore there are $K - NJ - 1 = \kappa^N - NJ - 1$ zeros per row of Q. The order condition needed is that the number of zero restrictions is at least $\lfloor (K-1)/2 \rfloor$. This condition is satisfied when $\kappa^N - NJ - 1 \geq \lfloor (K-1)/2 \rfloor$ or when the number of choices J is such that $J \leq \frac{1}{N}(\kappa^N - \lfloor (\kappa^N - 1)/2 \rfloor - 1)$. The matrix R representing these restrictions has full rank because each zero restriction involves only a single element of Q, so each row of R is a distinct row of the $K^2 \times K^2$ identity matrix.

The sufficient condition follows by rearranging this inequality as $NJ + 1 + \lfloor (\kappa^N - 1)/2 \rfloor \leq \kappa^N$ and then noting that by definition of the floor function, $\lfloor (\kappa^N - 1)/2 \rfloor \leq (\kappa^N - 1)/2$. ■

M	N		q_1	q_{21}	q_{22}	π	λ	β	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		λ	γ	κ	η	φ
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		λ	γ	κ	η	φ
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_{avg}	ω_{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).