

Identification and Estimation of Continuous Time Dynamic Discrete Choice Games*

JASON R. BLEVINS

The Ohio State University

November 14, 2022

Abstract. This paper considers the theoretical, computational, and econometric properties of a class of continuous time dynamic discrete choice games with stochastically sequential moves, introduced by Arcidiacono, Bayer, Blevins, and Ellickson (2016). In a generalized version of the model with heterogeneous move arrival rates, we first re-establish conditions for existence of a Markov perfect equilibrium. Then, we consider nonparametric identification of the model primitives with only discrete time data sampled at a fixed time interval. We consider identification not only of the payoff functions, as previous work has done, but also the move arrival rates. Three canonical models are considered: a single agent renewal model, a dynamic model of entry and exit, and a quality ladder model of oligopoly dynamics. These models are foundational for many applications in applied microeconomics. Through these examples we examine the computational properties of the model and statistical properties of estimators via a series of small- and large-scale Monte Carlo experiments. These experiments shed light on how the parameter estimates behave as one moves from continuous time data to discrete time data of decreasing frequency and on the computational feasibility of the model as the number of firms grows.

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

JEL Classification: C13, C35, C62, C73.

*I am grateful for numerous useful comments from and discussions with seminar participants at Columbia

1. Introduction

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

For many economic models there is not a natural, fixed time interval at which agents make decisions. Despite this, it has been standard practice for applied researchers to calibrate the decision interval in their empirical model to the sampling interval of the data. However, allowing agents to make decisions asynchronously at (possibly unknown) stochastic points in continuous time (which may be unknown and stochastic) can be both more natural and easier computationally. Even in cases where there is a compelling reason to model decisions as simultaneous and occurring at fixed time intervals, there is in general no reason that the decision interval should coincide exactly with the data sampling interval. Continuous time models also have the benefit of being invariant to the interval at which observations are recorded, while standard dynamic discrete choice models have different functional forms when applied to different time intervals.

Given the practical and conceptual benefits continuous time models, the main goal of this paper is to develop new results on identification of a generalized version of the ABBE model which allows for firm- and state-specific move arrival rates. First, existence of Markov perfect equilibrium and a linear characterization of the value function are

University, Indiana University, Northwestern University, Texas A&M University, the University of British Columbia, the University of Chicago, the University of Iowa, the University of Michigan, and the University of Montreal, as well as 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, the 2015 Econometric Society World Congress and the 2019 International Association for Applied Econometrics conference. This work builds on previous joint research with Peter Arcidiacono, Patrick Bayer, and Paul Ellickson and has benefited tremendously from our many discussions together.

established in the model. Next, identification conditions for the model in general and in the context of three canonical example models are examined. Specifically, we consider identification of the rate of move arrivals which was assumed to be known in previous work. Finally, the computational and econometric properties are explored via a series of Monte Carlo experiments. A point of particular interest in our experiments is how estimates behave as one moves from ideal continuous time sampling to discrete time data sampled at longer intervals.

Modeling economic processes in continuous time dates back at least several decades to work in time series econometrics by Phillips (1972, 1973), Sims (1971), Geweke (1978), and Geweke, Marshall, and Zarkin (1986) and work on longitudinal models by Heckman and Singer (1986). Despite this early work on continuous time models, discrete time models became the de facto standard for dynamic discrete choice and now have a long, successful history in structural applied microeconomics starting with the pioneering work of Gotz and McCall (1980), Miller (1984), Pakes (1986), Rust (1987), and Wolpin (1984). A recent series of papers (Aguirregabiria and Mira, 2007; Bajari, Benkard, and Levin, 2007; Pakes, Ostrovsky, and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008) have shown how to extend two-step estimation techniques, originally developed by Hotz and Miller (1993) and Hotz, Miller, Sanders, and Smith (1994) in the context of single-agent dynamics, to more complex multi-agent settings. The computation of multi-agent models remains formidable, despite a growing number of methods for solving for equilibria (Pakes and McGuire, 1994, 2001; Doraszelski and Satterthwaite, 2010).

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

the degree of heterogeneity in applied work using these methods.

Because of these limitations some authors have recently considered continuous time models which more closely reflect the nature and timing of actions by agents in the models while also reducing the computational burden. [Doraszelski and Judd \(2012\)](#) showed that continuous-time dynamic games have desirable computational properties, significantly decreasing the computational burden required to evaluate the Bellman operator, which can be used to compute equilibria. ABBE demonstrated the empirical tractability of continuous-time games, particularly for applications in industrial organization. They proposed an econometric model which retains the computational advantages of continuous time models while incorporating many familiar 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. Continuous time models have since been used in many applications including [Takahashi \(2015\)](#) to movie theaters, [Deng and Mela \(2018\)](#) to TV viewership and advertising, [Nevskaya and Albuquerque \(2019\)](#) to online games, [Agarwal, Ashlagi, Rees, Somaini, and Waldinger \(2021\)](#) to allocation of donor kidneys, [Jeziorski \(2022\)](#) to the U.S. radio industry, [Schiraldi, Smith, and Takahashi \(2012\)](#) to supermarkets in the U.K., [Lee, Roberts, and Sweeting \(2012\)](#) to baseball tickets, [Cosman \(2017\)](#) to bars in Chicago, [Mazur \(2017\)](#) to the U.S. airline industry, [Kim \(2021\)](#) to the U.S. retail banking industry, and [Qin, Vitorino, and John \(2022\)](#)

to airline networks in China.

The remainder of this paper is organized as follows. In [Section 2](#), we review a generalized version of the ABBE model that permits additional heterogeneity in the form of player-specific discount rates and move arrival rates that may vary by player and state. We establish a linear representation of the value function in terms of CCPs as well as the existence of a Markov perfect equilibrium in the more general model. We then develop new identification results for the model in [Section 3](#). We use two canonical examples throughout the paper to illustrate our results: a single agent renewal model based on [Rust \(1987\)](#) and a 2×2 entry model similar to example models used by [Aguirregabiria and Mira \(2007\)](#), [Pesendorfer and Schmidt-Dengler \(2008\)](#), and others. Although our running examples are intentionally quite simple, to better illustrate the main ideas, in [Section 4](#) we introduce a third example: a quality ladder model of oligopoly dynamics with heterogeneous firms based on the model of [Ericson and Pakes \(1995\)](#).¹ Finally, in [Section 5](#) we examine the computational and econometric properties via a series of Monte Carlo experiments. [Section 6](#) concludes.

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

We consider infinite horizon discrete games in continuous time indexed by $t \in [0, \infty)$ with N players indexed by $i = 1, \dots, N$. We introduce a heterogeneous generalization of the ABBE model, where players may have different discount rates and where the move arrival rates may differ by player and across states. After formalizing the components of the structural model, we establish a linear representation of the value function in terms of conditional choice probabilities, as in ABBE and [Pesendorfer and Schmidt-Dengler \(2008\)](#), as well as existence of a Markov perfect equilibrium in the more general model. We conclude the section with a comparison of discrete- and continuous-time models.

¹As another example, [Blevins and Kim \(2019\)](#) specify a continuous-time version of the dynamic entry-exit model of [Aguirregabiria and Mira \(2007\)](#).

2.1. State Space

At any instant, the payoff-relevant market conditions that are common knowledge to all players can be summarized by a state vector x , which is a member of a finite state space \mathcal{X} with $K \equiv |\mathcal{X}| < \infty$. Each element $x \in \mathcal{X}$ represents a possible state of the market and contains information about the market structure (e.g., which players are active, the quality of each player) and market conditions (e.g., demographic and geographic characteristics, input prices). The states $x \in \mathcal{X}$ are typically represented as vectors of real numbers in a finite-dimensional Euclidean space \mathbb{R}^L . 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. For example, $x = (x_1, \dots, x_N, d)$ where the components x_i are player-specific states, such as incumbency status or the number of stores operated by a chain, and d is an exogenous market characteristic, such as population or the level of demand.

Because the state space is finite there is an equivalent *encoded* state space representation $\mathcal{K} = \{1, \dots, K\}$. Although \mathcal{X} is the most natural way to interpret the state, using \mathcal{K} is convenient because it allows us to vectorize payoffs, value functions, and other quantities.

Renewal Example. *As an example, consider a continuous-time version of the single-agent renewal model of Rust (1987). The state of the model is the accumulated mileage of a bus engine $x \in \{x_1, x_2, \dots, x_K\}$, so the state can more simply be represented as an integer index $k \in \mathcal{K} = \{1, \dots, K\}$.*

2 × 2 Entry Example. *As a second example, consider a simple two-firm entry game with a binary exogenous state variable. Each firm $i \in \{1, 2\}$ has two actions $j \in \{0, 1\}$. The choice $j = 1$ represents a switching choice: enter if inactive, or exit if active. On the other hand, the choice $j = 0$ represents a continuation choice: remain active or remain inactive. The exogenous state represents the level of demand, which can either be high or low. The state vector x_k has three components: x_{1k} and x_{2k} are activity indicators for firms 1 and 2 and the level of demand is represented by*

$d_k \in \{L, H\}$. Therefore, in vector form the state space is

$$\mathcal{X} = \{(0, 0, L), (1, 0, L), (0, 1, L), (1, 1, L), (0, 0, H), (1, 0, H), (0, 1, H), (1, 1, H)\}.$$

In encoded form, the state space is simply

$$\mathcal{X} = \{1, 2, 3, 4, 5, 6, 7, 8\}.$$

2.2. Decisions & Endogenous State Changes

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. Each player has J actions represented by the choice set $\mathcal{J} = \{0, 1, 2, \dots, J-1\}$. When the model is in state k , the holding time until the next move by player i is exponentially distributed with rate parameter λ_{ik} . In other words, decision times for player i in state k occur according to a time-homogeneous Poisson process with rate λ_{ik} . We assume these processes are independent across players and the rates λ_{ik} are finite for all i and k , reflecting the fact that monitoring the state and making decisions is costly, so continuous monitoring ($\lambda_{ik} = \infty$) is infeasible. Let h_{ijk} denote the hazard rate at which player i takes action j in state k such that the overall rate of decision in state k is $\lambda_{ik} = \sum_{j=0}^{J-1} h_{ijk}$. The choice-specific hazards are determined endogenously in the model through the equilibrium dynamic payoff maximization problems of players, which we discuss in detail below. When player i chooses action j , the state jumps immediately and deterministically from k the continuation state denoted by $l(i, j, k)$.

In most economic models, the actions of players only affect the individual components of the overall state vector. For example, when a new firm enters a market it may change the firm-specific activity indicator for that firm but not the level of demand in the market. As we will see below, this leads to sparsity of the continuous time model and helps with identification.

Renewal Example (continued). *There is a single agent in this model, the manager of a bus*

company, so $N = 1$. Suppose the manager decides whether to replace a bus engine ($j = 1$) or not ($j = 0$). Hence, the choice set is $\mathcal{F} = \{0, 1\}$. The hazard rate of engine replacement in state k is h_{1k} (where we drop the i index for simplicity). In practice we may assume the overall rate of decisions constant and equal to λ . This would imply the rate of (unobservable) non-replacement is $h_{0k} = \lambda - h_{1k}$. Even in this case, where the overall rate of decisions is constant, the rates of replacement and non-replacement (specific decisions) are endogenous and vary across states. This is similar to the case of discrete time models, where the sum of conditional choice probabilities is necessarily constant and equal to one while the individual choice probabilities vary across states. The continuous time model allows another degree of flexibility in that the rate of move arrivals can be different from one.

2.3. Exogenous State Changes

The state of the model can also evolve over time in response to exogenous events, which we attribute to an artificial player referred to as nature, indexed by $i = 0$. 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). When the model is in state k , let q_{kl} denote the hazard rate at which transitions to another state $l \neq k$ occur. The rate q_{kl} may be zero if direct transitions from k to l are not possible, or q_{kl} may be some positive but finite value representing the hazard rate of such a transition. Therefore, the overall rate at which the system leaves state k for any other state $l \neq k$ is $\sum_{l \neq k}^K q_{kl}$.

Renewal Example (continued). Suppose the exogenous mileage transition process is characterized by a single rate parameter γ governing one-state-ahead mileage increases. This rate is constant across states for simplicity, so for all $l \neq k$ we have

$$q_{kl} = \begin{cases} \gamma & \text{if } l = k + 1 \\ 0 & \text{otherwise.} \end{cases}$$

2 × 2 Entry Example (continued). In the 2 × 2 entry model, there are two exogenous states:

high demand ($d = H$) and low demand ($d = L$). Suppose nature switches from H to L at rate γ_{HL} and back to H at rate γ_{LH} . Then we have

$$(1) \quad q_{kl} = \begin{cases} \gamma_{HL} & \text{if } d_k = H \text{ and } d_l = L, \\ \gamma_{LH} & \text{if } d_k = L \text{ and } d_l = H, \\ 0 & \text{otherwise.} \end{cases}$$

2.4. Payoffs

In the continuous time setting, we distinguish between the flow payoffs that a player receives while the model remains in state k , denoted u_{ik} , and the instantaneous choice-specific payoffs from making choice j in state k at a decision time t , denoted $c_{ijk}(t)$. The instantaneous payoffs are additively separable as $c_{ijk}(t) = \psi_{ijk} + \varepsilon_{ijk}(t)$, where ψ_{ijk} is the mean payoff and $\varepsilon_{ijk}(t)$ is a choice-specific unobserved payoff. Player i observes the vector $\varepsilon_{ik}(t) \equiv (\varepsilon_{ijk}(t), j = 0, \dots, J-1)$ of choice-specific unobservables before choosing action j . All players and the researcher observe the state k , but only player i observes $\varepsilon_{ik}(t)$.

Remark. Note that in discrete time models, because all actions and state changes resolve simultaneously, the period payoffs are written as functions of the state, the unobservables, and the actions of all players (e.g., $u_i(a_1, \dots, a_N, x_t, \varepsilon_{it})$). In our continuous-time model, the payoffs resulting from competition in the product market accrue as flows u_{ik} in a specific state k while the choice-specific payoffs $c_{ijk}(t)$ accrue at the instant the decision is made.

Renewal Example (continued). *In the renewal model the agent faces a cost minimization problem where the flow utility u_{ik} is the cost of operating a bus with mileage k . For example, if the cost of mileage is β then a parametric flow utility function could be*

$$u_{ik} = -\beta k.$$

Upon continuation, no cost is paid but a fixed amount $\mu > 0$ is paid to replace the engine:

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

Continuation does not change the state, but upon replacement the state resets 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} .

2.5. Assumptions

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

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

Assumption 2 (Discount Rates). The discount rates $\rho_i \in (0, \infty), i = 1, \dots, N$ are known.

Assumption 3 (Move Arrival Times). Move arrival times follow independent Poisson processes with rate parameters λ_{ik} for each player $i = 1, \dots, N$ and state $k = 1, \dots, K$ and q_{kl} for exogenous state changes from each state k to $l \neq k$ due to nature, with $0 \leq \lambda_{ik} < \infty$ and $0 \leq q_{kl} < \infty$.

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

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

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 $j = 0, \dots, J - 1$ and $j' \neq j$.

Assumption 7 (Private Information). The choice-specific shocks $\varepsilon_{ik}(t)$ are iid across players i , states k , and decision times t . The joint distribution F_{ik} is known, is absolutely continuous with respect to Lebesgue measure (with joint density f_{ik}), has finite first moments, and has support equal to \mathbb{R}^J .

Assumptions 1–7 are generalized counterparts of Assumptions 1–4 of ABBE that allow for player heterogeneity and state dependent rates.² Assumptions 1–5 were discussed above. Assumption 6 formalizes that $j = 0$ is a costless continuation action and that all choices are observationally distinct. The first part of Assumption 6 requires that if an inaction decision which does not change the state, denoted $j = 0$, is included in the choice set, then the instantaneous payoff associated with that choice must be zero.³ This is an identifying assumption. 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.

2.6. Strategies, Best Responses, and Value Function

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

$$(2) \quad \sigma_{ijk} = \Pr[\delta_i(k, \varepsilon_{ik}) = j \mid k]$$

for all choices j and states k . Let ζ_{im} denote player i 's beliefs regarding the actions of rival player m , given by a collection of $J \times K$ probabilities for each state k and choice j . Let ζ_{ii}

²Specifically, Assumption 1 is equivalent to Assumption 1 of ABBE, Assumptions 2 and 3 generalize Assumptions 2(a) and 2(b–c) of ABBE, Assumption 4 is equivalent to Assumptions 2(d–e) of ABBE, and Assumptions 5–6 are equivalent to Assumptions 3–4 of ABBE, and Assumption 7 generalizes Assumption 5 of ABBE.

³The 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.

denote player i 's best response probabilities. Then we let $\zeta_i = (\zeta_{i1}, \dots, \zeta_{iN})$ denote player i 's beliefs and best responses. Finally, let $V_{ik}(\zeta_i)$ denote the expected present value for player i being in state k and behaving optimally in the future given beliefs ζ_i . For given beliefs ζ_i , the optimal policy rule for player i satisfies the following inequality condition:

$$(3) \quad \delta_i(k, \varepsilon_{ik}) = j \iff \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)}(\zeta_i) \geq \psi_{ij'k} + \varepsilon_{ij'k} + V_{i,l(i,j',k)}(\zeta_i) \quad \forall j' \in \mathcal{J}.$$

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

Remark. Given the threshold-crossing rule above, if the private information shocks are iid following the Type 1 Extreme Value distribution, the implied choice probabilities have a familiar logistic functional form:

$$(4) \quad \sigma_{ijk} = \frac{\exp(\psi_{ijk} + V_{i,l(i,j,k)}(\zeta_i))}{\sum_{j'} \exp(\psi_{ij'k} + V_{i,l(i,j',k)}(\zeta_i))}.$$

Given beliefs ζ_i held by player i , we can define the value function (here, a K -vector) $V_i(\zeta_i) = (V_{i1}(\zeta_i), \dots, V_{iK}(\zeta_i))^\top$ where the k -th element $V_{ik}(\zeta_i)$ is the present discounted value of all future payoffs obtained when starting in some state k and behaving optimally in future periods given beliefs ζ_i . For a small time increment τ , under Assumption 3 the probability of an event with rate λ_{ik} occurring is $\lambda_{ik}\tau$. Given the discount rate ρ_i , the discount factor for such increments is $1/(1 + \rho_i\tau)$. Thus, for small time increments τ the present discounted value of being in state k is (omitting the dependence on ζ_i for brevity):

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

Rearranging and letting $\tau \rightarrow 0$, we obtain the following recursive expression for V_{ik} for

beliefs ζ_i :

$$(5) \quad V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{kl} V_{il} + \sum_{m \neq i} \lambda_{mk} \sum_j \zeta_{imjk} V_{i,l(m,j,k)} + \lambda_{ik} \mathbf{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \}}{\rho_i + \sum_{l \neq k} q_{kl} + \sum_m \lambda_{mk}}.$$

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 $\varepsilon_{ik} = (\varepsilon_{i0k}, \dots, \varepsilon_{ij,k})^\top$.

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

$$V_k = \frac{1}{\rho + \gamma + \lambda} (u_k + \gamma V_{k+1} + \lambda \mathbf{E} \max \{ \varepsilon_{0k} + V_k, -c + \varepsilon_{1k} + V_1 \}).$$

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

$$V_{1k} = \frac{1}{\rho_1 + 1\{d_k = L\}\gamma_{LH} + 1\{d_k = H\}\gamma_{HL} + \lambda_{1k} + \lambda_{2k}} \\ \times \left(u_{1k} + 1\{d_k = L\}\gamma_{LH} V_{1,l(0,H,k)} + 1\{d_k = H\}\gamma_{HL} V_{1,l(0,L,k)} + \lambda_{2k} \zeta_{120k} V_{i,k} \right. \\ \left. + \lambda_{2k} \zeta_{121k} V_{i,l(2,1-x_{k2},k)} + \lambda_{1k} \mathbf{E} \max \left\{ \varepsilon_{i0k} + V_{i,k}, \psi_{i1k} + \varepsilon_{i1k} + V_{i,l(1,1-x_{k1},k)} \right\} \right),$$

where $l(0, H, k)$ and $l(0, L, k)$ are the continuation states when nature switches the level of demand to H and L , respectively, when in state k . ζ_{12jk} is firm 1's belief about firm 2 choosing j .

2.7. Linear Representation of the Value Function

It will be convenient to express the Bellman equation in (5) in matrix notation. Let $\Sigma_m(\zeta_{im})$ denote the transition matrix implied by the beliefs ζ_{im} and the continuation state

function $l(i, \cdot, \cdot)$. That is, the (k, l) element of the matrix $\Sigma_m(\zeta_{im})$ is the probability of transitioning from state k to state l as a result of an action by player m under the beliefs of player i . Let $Q_0 = (q_{kl})$ denote the matrix of rates of exogenous state transitions and let $\tilde{Q}_0 = Q_0 - \text{diag}(q_{11}, \dots, q_{KK})$ be the matrix formed by taking Q_0 and replacing the diagonal elements with zeros.

Then, following (5) we define the operator Γ_i as

$$(6) \quad \Gamma_i(V_i) = D_i \left[u_i + \tilde{Q}_0 V_i + \sum_{m \neq i} L_m \Sigma_m(\zeta_{im}) V_i + L_i \{ \Sigma_i(\zeta_{ii}) V_i + C_i(\zeta_{ii}) \} \right],$$

where D_i is the $K \times K$ diagonal matrix containing the coefficient from (5) for each k , hence $(D_i)_{kk} = 1/(\rho_i + \sum_{l \neq k} q_{kl} + \sum_{m=1}^N \lambda_{mk})$, $L_m = \text{diag}(\lambda_{m1}, \dots, \lambda_{mK})$ is a diagonal matrix containing the move arrival rates for player m , $C_i(\zeta_{ii})$ 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 ζ_{ii} . That is, k -th element of $C_i(\zeta_{ii})$ is $\sum_{j=0}^{J-1} \zeta_{iijk} [\psi_{ijk} + e_{ijk}(\zeta_{ii})]$, where $e_{ijk}(\zeta_{ii})$ is the expected value of ε_{ijk} given that action j is chosen:

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

Hence, for the given beliefs ζ_i , the value function is a fixed point of Γ_i : $V_i = \Gamma_i(V_i)$.

A central result of ABBE (Proposition 2) showed that the differences in choice-specific value functions which appear in the definition of $e_{ijk}(\sigma_i)$ above are functions of the conditional choice probabilities. 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 as well as the estimators of ABBE and [Blevins and Kim \(2019\)](#). It is analogous to a similar result for discrete time games by [Pesendorfer and Schmidt-Dengler \(2008, eq. 6\)](#).

Theorem 1. *If Assumptions 1–7 hold, then for a given collection of beliefs $\zeta = (\zeta_1, \dots, \zeta_N)$, V_i has the following linear representation for each i :*

$$(7) \quad V_i(\zeta) = \Xi_i^{-1}(\zeta) [u_i + L_i C_i(\zeta_i)]$$

where

$$(8) \quad \Xi_i(\zeta) = \rho_i I_K + \sum_{m=1}^N L_m [I_K - \Sigma_m(\zeta_m)] - Q_0$$

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

Proof. See [Appendix A](#).

2.8. Equilibrium

Following the literature, we focus on Markov perfect equilibria.

Definition. A *Markov perfect equilibrium* is a collection of stationary policy rules $\{\delta_i\}_{i=1}^N$ such that (3) holds for all i, k , and ε_{ik} given beliefs $\zeta_i = (\sigma_1, \dots, \sigma_N)$ generated by (2).

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

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

Proof. See [Appendix A](#).

2.9. Continuous Time Markov Jump Processes Representation

The reduced form of the model we consider is a finite state Markov jump process, a stochastic process $X(t)$ indexed by $t \in [0, \infty)$ taking values in \mathcal{X} . If we begin observing




FIGURE 1. Finite State Markov Jump Process

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

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 of $X(t)$ for $t \in [0, \bar{T}]$ is shown along with 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.

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 its *intensity*

matrix or infinitesimal generator matrix. Consider the intensity matrix for nature,

$$Q_0 = \begin{bmatrix} q_{11} & q_{12} & \dots & q_{1K} \\ q_{21} & q_{22} & \dots & q_{2K} \\ \vdots & \vdots & \vdots & \vdots \\ q_{K1} & q_{K2} & \dots & q_{KK} \end{bmatrix}$$

where for $k \neq l$

$$q_{kl} = \lim_{h \rightarrow 0} \frac{\Pr [X(t+h) = l \mid X(t) = k]}{h}$$

is the probability per unit of time that the system transitions from state k to state l and the diagonal elements are $q_{kk} = -\sum_{l \neq k} q_{kl}$ so that the row sums equal zero. The holding times before transitions out of state k follow an exponential distribution with rate parameter $-q_{kk}$, which is the sum of the off-diagonal transition rates. Conditional on leaving state k , the system transitions to state $l \neq k$ with probability $q_{kl} / \sum_{l \neq k} q_{kl} = -q_{kl} / q_{kk}$.

In the case of discrete time data, the times at which actions and state changes occur are not observed by the econometrician. With equispaced data (e.g., annual or quarterly) only the states at the beginning and end of each period of length Δ are observed. Although we cannot know the exact sequence of actions and state changes, the model allows us to determine the likelihood of any particular transition occurring over a time interval of length Δ using the *transition matrix*, which we will denote as $P(\Delta)$.

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 . The transition matrix $P(\Delta) = (P_{kl}(\Delta))$ is the corresponding $K \times K$ matrix of these probabilities. For a finite-state continuous time Markov jump processes, the Kolmogorov forward equations form a system of matrix differential equations characterizing the transition matrix $P(\Delta)$ of

a process with intensity matrix Q (Karlin and Taylor, 1975, 4.8):

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

It follows that the unique solution to this system is

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

The transition matrix is the matrix exponential of the intensity matrix Q scaled by Δ . This is the matrix analog of the scalar exponential $\exp(x)$ for $x \in \mathbb{R}$.⁴

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, relating back to the structural model, note that if the model is in state k then there are $N + 1$ independent competing Poisson processes (or exponential distributions) that may change the state: one for each player and nature. The rates of these processes are λ_{ik} , for $i = 1, \dots, N$, and $\sum_{l \neq k} q_{kl}$ respectively. Therefore, the distribution of the minimum holding time is exponential with rate parameter equal to the sum of the individual rates: $\sum_{i=1}^N \lambda_{ik} + \sum_{l \neq k} q_{kl}$. Notice that this is precisely the rate that appears in the denominator of the Bellman equation in (5). Furthermore, conditional on an event the probability that it is due to process i is $\lambda_{ik} / \left(\sum_{m=1}^N \lambda_{mk} + \sum_{l \neq k} q_{kl} \right)$ for player $i > 0$ (and similarly for nature).

We can think of these $N + 1$ processes as first branching at the player level and then branching again at the action level: conditional on a particular player moving, which action is chosen? Player i plays each action j in state k at rate $h_{ijk} = \lambda_{ik} \sigma_{ijk}$. Since the probabilities σ_{ijk} sum to one, we have $\sum_{j=0}^{J-1} h_{ijk} = \lambda_{ik}$. Therefore, conditional on player i moving the probability that action j is chosen is h_{ijk} / λ_{ik} .

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

⁴Although we cannot calculate the infinite sum (10) exactly, we can compute $\exp(\Delta Q)$ numerically using known algorithms implemented, for example, in the Fortran package Expokit (Sidje, 1998) or the `expm` command in Matlab. A more recent development is the uniformization algorithm of Sherlock (2022).

intensity matrices Q_0, Q_1, \dots, Q_N . Each process corresponds to some player i and the *aggregate intensity matrix* is defined as $Q \equiv Q_0 + Q_1 + \dots, Q_N$.

Renewal Example (continued). Consider the Q matrix implied by the continuous-time single-agent renewal model. The state variable in the model is the total accumulated mileage of a bus engine, $\mathcal{K} = \{1, \dots, K\}$. The exogenous state transition process is characterized by a $K \times K$ intensity matrix Q_0 on \mathcal{K} with one parameter, γ , governing the rate of mileage increases:

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

Let σ_{1k} denote the probability of replacement in state k . The intensity matrix for state changes induced by the the agent is

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




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 aggregate intensity matrix in this case is $Q = Q_0 + Q_1$:

$$(11) \quad Q = \begin{bmatrix} -\gamma & \gamma & 0 & 0 & \cdots & 0 \\ \lambda\sigma_{12} & -\lambda\sigma_{12} - \gamma & \gamma & 0 & \cdots & 0 \\ \lambda\sigma_{13} & 0 & -\lambda\sigma_{13} - \gamma & \gamma & \cdots & 0 \\ \lambda\sigma_{14} & 0 & 0 & -\lambda\sigma_{14} - \gamma & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \lambda\sigma_{1K} & 0 & 0 & 0 & \cdots & -\lambda\sigma_{1K} \end{bmatrix}.$$

A representative sample path generated by this model is shown in Figure 2. Holding 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.

2 × 2 Entry Example (continued). Let h_{ik} be the hazard of player i switching from active to inactive or vice versa in state k . Let γ_{LH} and γ_{HL} be the rates at which nature switches between

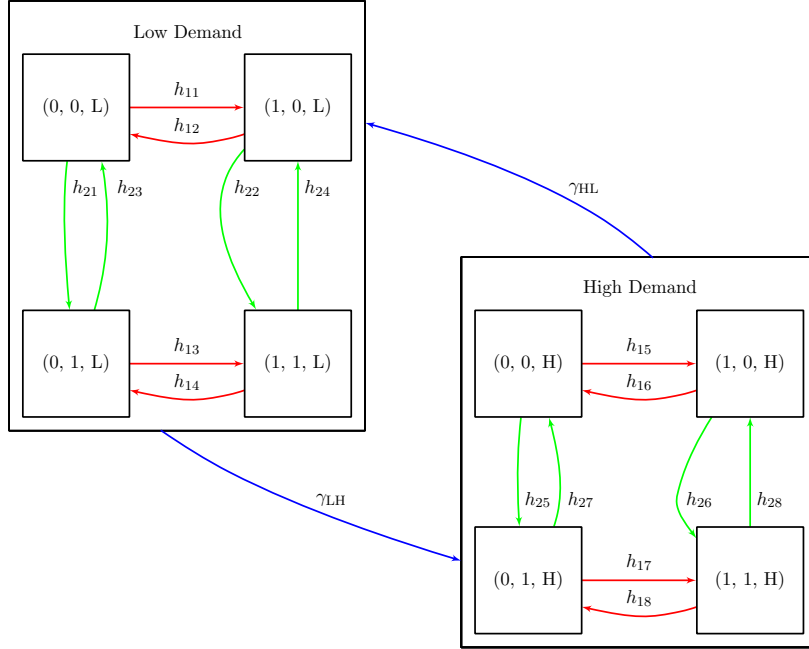


FIGURE 3. Two Player Entry Game with Exogenous Demand State

demand states (i.e., demand moves from low to high at rate γ_{LH}). The aggregate 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

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

The low demand state L corresponds to encoded states $k = 1, \dots, 4$. In this portion of the state

space, firms change the state as follows:

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

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 . The full Q matrix is stated below in (12) in Section 3.

2.10. Comparison with Discrete Time Models

We conclude this section with a few brief remarks on continuous time and discrete time models. First, consider a typical discrete time model in which agents make decisions in unison and where the period between decisions is calibrated to be equal to the sampling period of the data, say one year. In an entry/exit setting where the choice set is $\mathcal{F} = \{0, 1\}$, this implies that there must be exactly one entry or exit per year. For example, entering and leaving within one year is not permitted. In the discretely sampled data, passive actions such as remaining in or out of the market are coded as active decisions (e.g., a choice to not enter or a choice to remain in the market), but in reality they typically represent the absence of an active choice (entry or exit) during the period. Now consider a chain store setting where the choice is the net number of stores to open during the year, the choice set is $\mathcal{F} = \{-J, \dots, J\}$. This implies that there can be at most J openings or closings per year. Hence, J must be chosen by the researcher to be the maximum number of possible stores opened or closed by any chain firm in any period.

Now, consider a continuous time model with a common move arrival rate λ for all players and all states. In the entry/exit setting, the choice set is still $\mathcal{F} = \{0, 1\}$ which implies that there are *on average* $1/\lambda$ entries or exits per year. Multiple entries and exits

are allowed and the model parameters imply a distribution over the number and type of such events. The choice set represents the set of possible *instantaneous* state changes, so the chain store expansion example if we assume that no more than one store is ever opened or closed simultaneously, then we would specify $\mathcal{F} = \{-1, 0, 1\}$. This would imply that *on average* there are at most $1/\lambda$ openings or closings per year. In the continuous time model the rate λ is a free parameter that can adjust to match the data, thus not imposing an ad hoc restriction on the number of actions per unit of time. In other words, the time-aggregated implications of the continuous time model are not functionally different if we change the time period and unrelated to the sampling period of the data.

3. Identification Analysis

Due to the time aggregation problem, our identification analysis proceeds in two main steps corresponding to the reduced form and the structural model primitives. Deriving the implications of the structural model can be viewed as a bottom-up exercise: the structural primitives u and ψ imply value functions V which imply choice probabilities σ . These probabilities along with the rates of moves, λ , and state transitions by nature, Q_0 , in turn imply an intensity matrix Q . Finally, given the Q matrix and a process for sampling data, this implies a data generating process. For example, for a fixed sampling interval Δ the distribution of observable data is $P(\Delta) = \exp(\Delta Q)$.

On the other hand, the identification problem requires us to consider the inverse problem, working from the top down. These steps are represented in [Figure 4](#). If the complete continuous time record is potentially observable, then Q is trivially identified and we can move to identification of the structural model. However, in the case of discrete time data we must first use our knowledge of the data generating process, represented by the transition matrix $P(\Delta)$ for an interval Δ , to derive conditions under which we can uniquely determine the reduced form intensity matrix Q . We will show that this is possible under mild conditions by exploiting the restrictions that the structural model places on the Q matrix.




FIGURE 4. Identification Analysis

Second, with Q in hand we turn to identification of the structural primitives of the model, namely the flow payoffs u and instantaneous payoffs ψ . We show that knowledge of Q allows us to recover these structural primitives with fewer identifying restrictions than are required in discrete time models.

3.1. Identification of Q

With continuous-time data, identification and estimation of the intensity matrix for finite-state Markov jump processes is straightforward and well-established (Billingsley, 1961). However, when a continuous-time process is only sampled at discrete points in time, the parameters of the underlying continuous-time model may not be point identified.⁵ In the present model, the concern is that there may be multiple Q matrices which give rise to the same data generating process, which is the potentially observable transition probability matrix $P(\Delta)$ in the leading case of fixed sampling intervals.⁶

In discrete time settings, there is a similar identification problem that is masked when assuming the unknown frequency of moves is equal to the (known) sampling frequency (Hong, Li, and Wang, 2015). To see this, suppose agents move at intervals of length δ with transition matrix P_0 while the data sampling interval is $\Delta > \delta$. Then the mapping between the data (equispaced observations at length Δ) and the transition matrix is: $P(\Delta) = P_0^{\Delta/\delta}$. In general, there are multiple solutions to this equation (Gantmacher, 1959; Singer and

⁵This is known as the *aliasing problem* and it has been studied extensively in the context of continuous-time systems of stochastic differential equations (Sims, 1971; Phillips, 1973; Hansen and Sargent, 1983, 1991; Geweke, 1978; Kessler and Rahbek, 2004; McCrorie, 2003; Blevins, 2017). See Figure 1 of Blevins (2017) for an illustration in the frequency domain, where the problem is perhaps most obvious.

⁶A 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 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 $P(\Delta)$. We assume throughout that the continuous time model is well-specified and that such an intensity matrix exists.




FIGURE 5. Time aggregation: Two distinct paths which end in the same state at t_2 and begin in the same state at $t_2 - \Delta$ and but differ over intermediate interval of length Δ .

Spilerman, 1976), meaning that identification of P_0 is non-trivial.

To illustrate this issue in the continuous time setting, Figure 5 displays two distinct paths which coincide both before and after an interval of length Δ , but which take different intermediate steps. Consider the possible paths of the process between times $t_2 - \Delta$ and t_2 . The dashed path first moves to a higher state before arriving at the resulting state k_{t_2} , while the dashed and dotted path first moves to a lower state and arrives in k_{t_2} at a later time (but before t_2). There are an infinite number of such paths, but the dynamics of the process over the interval are summarized by the transition matrix $P(\Delta)$.

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

For example, in applications there are typically player-specific components of the state vector where player i is not permitted to change the players-specific state of player j and

vice-versa. In an entry-exit model, such a state is incumbency status: players can enter and exit by their own action, but no player can enter or exit on behalf of another player. Similarly, if the overall state vector has components that are exogenous state variables, such as population, then we know that any state changes involving those variables must be due to nature and not by an action of any other player. This natural structure implies many linear restrictions on the Q matrix. We show that restrictions of this form limit the domain of the mapping $Q \mapsto \exp(\Delta Q) = P(\Delta)$ in such a way as to guarantee an almost surely unique intensity matrix Q for any given discrete time transition matrix $P(\Delta)$.

3.1.1. Identification of Unrestricted Q Matrices

Returning to the general problem of identification of Q , recall that the question is whether there exists a unique matrix Q that leads to the observed transition matrix $P(\Delta) = \exp(\Delta Q)$ when the process is sampled at uniform intervals of length Δ . The matrix logarithm $\ln P(\Delta)$ is not unique in general (see [Gantmacher, 1959](#); [Singer and Spilerman, 1976](#)), so the question amounts to finding suitable conditions under which there is a unique solution.

Previous mathematical treatments have tended to view the relationship $\exp(\Delta Q) = P(\Delta)$ from the perspective of the transition matrix $P(\Delta)$. In such cases there is not an underlying model that generates Q , so Q is the model primitive of interest and is unrestricted (aside from requirement that it must be a valid intensity matrix). As a result, most previous work on the aliasing problem focused on finding sufficient conditions on the matrix $P(\Delta)$ (rather than Q) to guarantee that $\ln P(\Delta)$ is unique. For example, if the eigenvalues of $P(\Delta)$ are distinct, real, and positive, then Q is identified ([Culver, 1966](#)). More generally, [Culver \(1966\)](#) proved that Q is identified if the eigenvalues of $P(\Delta)$ are positive and no elementary divisor (Jordan block) of $P(\Delta)$ belonging to any eigenvalue appears more than once. Other sufficient conditions for identification of Q include $\min_k \{P_{kk}(\Delta)\} > 1/2$ ([Cuthbert, 1972](#)) and $\det P(\Delta) > e^{-\pi}$ ([Cuthbert, 1973](#)). See [Singer and Spilerman \(1976\)](#) for a summary of these results and others.

Other sufficient conditions for identification of Q involve alternative sampling schemes.

For example, Q can always be identified for some sufficiently small sampling interval 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—restrictions on $P(\Delta)$ —are based on a “top down” approach and are undesirable in cases where Q is generated by an underlying model. The second type of conditions are based on changing how the continuous time process is sampled, which is not possible to change if 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 players, actions, and the resulting state transitions are defined.

3.1.2. Structural Restrictions for Identification of Q

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

Our proposed identification strategy is inspired by this work on multivariate regression models, but our model is different because the Q matrix is known to be an intensity matrix (rather than an arbitrary matrix of regression coefficients) and has a rather sparse structure which is dictated by an underlying structural model. Yet, there are a number of similarities: the present model can also be characterized by a system of differential equations as in

(9), where the intensity matrix Q plays a role similar to the matrix A above. If Q is an valid intensity matrix, then the functions $P(\Delta)$ which solve this system are the transition matrices of continuous-time stationary Markov chains (Chung, 1967, p. 251–257).

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

Renewal Example (continued). *In the single-agent renewal model the aggregate intensity matrix is given in (11) of Section 2. The number of nonzero hazards in this matrix is substantially less than the total number. Of the 20 non-trivial state-to-state transitions, only 8 are permitted: four to nature and four by action of the player. The remaining 12 transitions are not possible in a single step. Nature cannot decrease mileage and can only increase it by one state at a given instant (although multiple state jumps are possible over an interval of time). 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 choice-specific hazards h_{1k} are the products of the overall move arrival rates and the conditional choice probabilities, which introduces shape restrictions on $h_{1k} = \lambda\sigma_{1k}$ across states k .*

2×2 Entry Example (continued). *In the $2 \times 2 \times 2$ entry example, the aggregate intensity*

matrix is $Q = Q_0 + Q_1 + Q_2$:

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

where the diagonal elements have been omitted for simplicity. Some transitions cannot happen at all, such as $(0,1,L)$ to $(1,0,L)$. The remaining transitions can happen only due to the action of one of the firms, but not the other. For example, moving from $(0,0,H)$ to $(1,0,H)$ 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 sparse, which makes identification of Q more likely even with time aggregation since any observationally equivalent Q matrix must have the same sparsity pattern. Finally, given Q we can again separately recover Q_0 , Q_1 , and Q_2 .

Similar sparse structures arise in even models with large numbers of players 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 $R \text{vec}(Q) = r$. For the $K \times K$ matrix $Q = (q_{kl})$, $\text{vec}(Q)$ is the vector obtained by stacking the columns of Q : $\text{vec}(Q) = (q_{11}, q_{21}, \dots, q_{K1}, \dots, q_{1K}, \dots, q_{KK})^\top$.

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 $R \text{vec}(\tilde{Q}) = r$. By the relationship between Q and \tilde{Q} above, we also have $R \text{vec}(Q + UDU^{-1}) = r$. But $R \text{vec}(Q) = r$ and by linearity of the vectorization operator, $R \text{vec}(UDU^{-1}) = 0$. An equivalent representation is

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

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

The following theorem establishes that there are sufficiently many full rank restrictions to identify Q in a broad class of games. This theorem includes exogenous market-specific state variables and shows that such states increase the number of zero restrictions and make identification of Q more likely, as do player-specific state variables.

Theorem 3 (Identification of Q). *Suppose the state vector is $x = (x_0, x_1, \dots, x_N) \in \mathcal{X}_0 \times \mathcal{X}_1 \times \dots \times \mathcal{X}_N$ where the component $x_0 \in \mathcal{X}_0$ is an exogenous market characteristic taking $|\mathcal{X}_0| = K_0$ values and for each $i = 1, \dots, N$ the component x_i is a player-specific state affected only by the action of each player with $|\mathcal{X}_i| = K_1$ possible distinct values. If Q has distinct eigenvalues that do not differ by an integer multiple of $2\pi i / \Delta$, then Q is generically identified when*

$$(13) \quad K_0 K_1^N - K_0 - NJ + \frac{1}{2} \geq 0.$$

The quantity on the left is strictly increasing in K_1 , strictly increasing in K_0 when $K_1 > 1$, and strictly decreasing in J .

Proof. See [Appendix A](#).

Generic identification means that Q is identified with the exception of a measure zero set of population Q matrices (see [Phillips, 1973](#); [Blevins, 2017](#)). The sparsity of Q helps and is increasing in both the number of exogenous states K_0 and player-specific states K_1 , but decreasing in the number of choices J . Therefore, for identification we need either a sufficiently large number of states or a sufficiently small number of choices. Fortunately, in most applications J is small relative to K .

2×2 Entry Example. Our running entry model example is a binary choice game with $N = 2$, $J = 2$, $K_0 = 2$, and $K_1 = 2$, so by [Theorem 3](#) Q is generically identified.

Furthermore, we can see that any binary choice game ($N > 1$ with $J = 2$) with meaningful player-specific states ($K_1 > 1$) is identified, regardless of the number of players or exogenous market states K_0 . The sufficient condition in this case simplifies to $K_0(K_1^N - 1) \geq N - \frac{1}{2}$. When $K_0 \geq 1$ and $K_1 \geq 2$ we have $K_0(K_1^N - 1) \geq 2^N - 1$ which exceeds $N - \frac{1}{2}$ for integers $N > 1$.

3.1.3. Identification of Q_i

Next, we make the following assumption which requires that given the aggregate intensity matrix Q , we can determine the player-specific intensity matrices Q_i .

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

This assumption is obvious in the models we have considered, where players cannot change each other's state variables and where actions by nature can be distinguished from the actions of players. Note also that the diagonal elements are unimportant: if the off-diagonal elements of each Q_i can be identified from Q , then diagonal elements are equal to the negative of the sum of the off-diagonal elements. This assumption can be

verified by inspection of Q in both of our running examples. In the single-agent renewal example Q is given in (11) and for the two-player entry model Q is given in (12). A sufficient condition for Assumption 8 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 ,

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

3.2. Identification of Hazards, Value Functions and Payoffs

We now establish that the value functions, instantaneous payoffs, and utility functions 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, we show below that they are identified.

Importantly, we note that when $j = 0$ is a latent or unobserved continuation action, it is not possible to identify the rates h_{i0k} even with continuous time data, so we cannot immediately treat them as identified quantities.

For simplicity, we consider the case where the choice-specific errors have a type 1 extreme value distribution. Noting that $h_{ijk} = \lambda_{ik}\sigma_{ijk}$ and recalling the choice probabilities in (4), in this case differences in log hazards can be written as

$$\ln h_{ijk} - \ln h_{i0k} = \ln \sigma_{ijk} - \ln \sigma_{i0k} = \psi_{ijk} + V_{i,l(i,j,k)} - V_{ik}.$$

Rearranging, we have

$$\ln h_{ijk} = \ln h_{i0k} + \psi_{ijk} + V_{i,l(i,j,k)} - V_{ik}.$$

The hazards on the left hand side for $j = 1, \dots, J - 1$ are identified from Q , while the quantities on the right hand side are unknowns to be identified.

Stacking equations across states k and choices j gives a linear system with $(J - 1)K$

identified hazards, K unknown hazards, $(J - 1)K$ unknown instantaneous payoffs, and K unknown valuations. The total number of unknowns is $(J + 1)K$. There are $2K$ more unknowns than identified hazards, so identification fails without further restrictions.

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 have,

$$\begin{bmatrix} \ln h_{i1} \\ \vdots \\ \ln h_{i,J-1} \end{bmatrix} = \left[\begin{array}{c|cccc|c} I_K & I_K & 0 & \dots & 0 & S_{i1} - I_K \\ I_K & 0 & I_K & \dots & 0 & S_{i2} - I_K \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ I_K & 0 & 0 & \dots & I_K & S_{i,J-1} - I_K \end{array} \right] \begin{bmatrix} \ln h_{i0} \\ \psi_{i1} \\ \vdots \\ \psi_{i,J-1} \\ V_i \end{bmatrix}.$$

Define X_i to be the $(J - 1)K \times (J + 1)K$ partitioned matrix and let R_i and r_i denote linear restrictions on the unknowns for player i . Let h_i^+ denote the identified hazards for choices $j > 0$ and h_i^0 denote the unidentified hazards for $j = 0$. Then the augmented system is:

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

Under [Assumption 6](#), for any action $j > 0$ in any state k , the continuation state is different from k . Therefore, the diagonal elements of S_{ij} are all zero and $S_{ij} - I_K$ has full rank for each $j > 0$ and these blocks are linearly independent across j . This means that X_i has rank $(J - 1)K$ and so we will need $2K$ additional full-rank restrictions for identification.

Theorem 4. *If for player i there exists a collection of linear restrictions represented by a matrix R_i*

and vector r_i such that

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

and the matrix $\begin{bmatrix} X_i \\ R_i \end{bmatrix}$ has rank $(J + 1)K$, then h_i^0 , ψ_i , and V_i are identified.

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

It is helpful now to consider some examples. If we assume that the instantaneous payoffs are constant across k , as is common in applications, this implies $\psi_{ijk} - \psi_{ijl} = 0$ for all choices $j > 0$ and all states $l \neq k$. This gives $(J - 1)K$ restrictions per player. If there are $J \geq 3$ choices, this is sufficient. When $J = 2$, we still need K additional restrictions. If we further assume that the move arrival rate is constant across states ($\sum_{j=0}^{J-1} h_{ijk} = \sum_{j=0}^{J-1} h_{ijl}$ for all $l \neq k$) then we have $K - 1$ restrictions. In this case, even if $J = 2$ then only 1 additional restriction is are needed.

Finding additional full-rank restrictions is not difficult in most applications. Examples include states where the value function is known, for example, if $V_{ik} = 0$ when a firm has permanently exited. Exclusion restrictions of the form $V_{ik} = V_{ik'}$ are also common, where k and k' are two states that differ only by a rival-specific state and are payoff equivalent to firm i . In all of these cases, the rank condition can be verified by inspection in applications.

Renewal Example (continued). *In the single-agent renewal model, since the replacement cost does not depend on the mileage state we have $\psi_{1k} = c$ for all k . This alone yields $K - 1$ restrictions of full rank of the form $\psi_{1k} - \psi_{11} = 0$ for all k . The linearity of the utility function imposes restrictions on V , and although this does not fit in the linear restriction framework of [Theorem 4](#) it also contributes to identification of ψ and V .*

2 × 2 Entry Example (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.*

Finally, we note that in practice the overall rate of actions can be identified through the nonlinear restrictions imposed by the distributional assumptions on the error term, which imply shape restrictions on the choice probabilities across states. These are difficult to characterize in the linear restriction framework we have used here, but in practice parametric assumptions will aid identification in addition to the linear restrictions considered above.

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 (7),

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

where Ξ_i is the matrix function defined in (8). 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 (Identification of Flow Payoffs). *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 also identified.*

4. A Continuous-Time Quality Ladder Model of Oligopoly Dynamics

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

4.1. 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.⁷ The set of relevant market configurations is thus the set of ordered tuples of length $\bar{\omega} + 1$ whose elements sum to N , denoted $\mathcal{S} = \{(s_1, \dots, s_{\bar{\omega}+1}) : \sum_j s_j = N, s_j \in \mathbb{Z}^*\}$, where \mathbb{Z}^* is the set of nonnegative integers. In this notation, each vector $\omega = (\omega_1, \dots, \omega_N) \in \Omega^N$

⁷In practice, we use the “probability density space” encoding algorithm described in [Gowrisankaran \(1999b\)](#), to map market structure tuples $s \in \mathcal{S}$ to integers $x \in \mathcal{X}$.

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

Each firm also needs to track its own quality, so payoff relevant market configurations from the perspective of firm i are described by a tuple $(\omega_i, s) \in \Omega \times \mathcal{S}$, where ω_i is firm i 's quality level and s is the market configuration. For our implementation, we map the multidimensional space $\Omega \times \mathcal{S}$ to an equivalent one-dimensional state space $\mathcal{X} = \{1, \dots, |\Omega| \times |\mathcal{S}|\}$. so that we can represent quantities in matrix-vector form and we use pre-computed transition addresses to avoid re-computing continuation states.

4.2. Product Market Competition

Again, we follow [Pakes and McGuire \(1994\)](#) in assuming a continuum of consumers with measure $\bar{M} > 0$ 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 define

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

Let $\delta_i(\omega, p)$ 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

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

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

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

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

Caplin and Nalebuff (1991) show that (in this single-product firm setting) there is a unique Bertrand-Nash equilibrium, which is given by the solution to the first order conditions:

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

Given the functional forms above, the first order conditions become

$$-(p_j - c)(1 - s_j) + 1 = 0.$$

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}) .

4.3. 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 $\mathcal{J} = \{0, 1, 2\}$. 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.

For any market-wide state $k \in \mathcal{K}$, let $\omega_k = (\omega_{k1}, \dots, \omega_{kN})$ denote the corresponding market configuration 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.

Each incumbent firm pays a constant flow fixed cost μ while remaining in the market, and receives the flow profits $\pi_{ik} = \pi(\omega_{ki}, \omega_{k,-i})$ associated with product market competition. The value function for an incumbent firm in state k is thus

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

where $\lambda_{ik} = \lambda$ for incumbents and potential entrants and $\lambda_{ik} = 0$ if firm i is not active in state k . Conditional upon moving while in state k , incumbent firms face the maximization problem $\max \{V_{ik} + \varepsilon_{i0}, -\kappa + V_{ik'} + \varepsilon_{i1}, \varphi + \varepsilon_{i2}\}$. The resulting choice probabilities are

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

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

4.4. 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_{ijt}^e . 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{K} 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

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

yielding the conditional entry-choice probabilities

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

4.5. 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.

5. Monte Carlo Experiments

In this section we describe Monte Carlo experiments conducted using the single-agent renewal model and the quality ladder model described in [Section 4](#).

5.1. Maximum Likelihood Estimation

The model can be estimated using maximum likelihood if either the equilibria can be enumerated or there is a unique equilibrium. Since the focus of this paper is identification, rather than developing a new estimator, our Monte Carlo experiments all proceed using the maximum likelihood estimator using value function iteration.⁸ Multiplicity of equilibria is not a concern for the single agent model and appears not to be a major issue in practice for the continuous-time oligopoly model specifications we consider below, although we have not established that there is a unique equilibrium.

With continuous-time data, we have a sample of \bar{N} tuples $(\tau_n, i_n, a_n, k_n, k'_n)$. Each describes a jump or move where, for each observation n : τ_n is the holding time since the previous event, i_n is the player index associated with this event ($i_n = 0$ is nature), a_n is the action taken by player i_n , k_n denotes the state at the time of the event, and k'_n denotes the state immediately after the event. Let $g(\tau; \lambda)$ and $G(\tau; \lambda)$ denote the pdf and cdf of $\text{Expo}(\lambda)$. Now, let $\ell_n(\theta)$ denote the likelihood of observation n given θ :

$$\ell_n(\theta) = \underbrace{g(\tau_n; q(k_n, k_n; \theta))}_{\text{Arrival time}} \left[\underbrace{\frac{q_0(k_n, k_n; \theta)}{q(k_n, k_n; \theta)}}_{\text{Event is jump}} \cdot \underbrace{p(k_n, k'_n; \theta)}_{\text{Transition}} \right]^{1_{\{i_n=0\}}} \\ \times \left[\underbrace{\frac{q_N(k_n, k_n; \theta)}{q(k_n, k_n; \theta)}}_{\text{Event is move}} \cdot \underbrace{\sigma(i_n, a_n, k_n; \theta)}_{\text{CCP}} \right]^{1_{\{i_n>0\}}} .$$

⁸More generally, it is possible that methods proposed for discrete time models, such as the homotopy method (Borkovsky, Doraszelski, and Kryukov, 2010; Besanko, Doraszelski, Kryukov, and Satterthwaite, 2010; Bajari, Hong, Krainer, and Nekipelov, 2010) or recursive lexicographical search (Iskhakov, Rust, and Schjerning, 2016), could be adapted to our model as well, but this is beyond the scope of the present paper.

Here, $q(k, k'; \theta)$ denotes the absolute value of the (k, k') element of the intensity matrix $Q(\theta)$ for given parameters θ . We use $q_0(k, k'; \theta)$ and $q_N(k, k'; \theta)$ similarly to denote the elements of Q_0 and $\sum_{i=1}^N Q_i$ respectively. Finally, $p(k, k'; \theta)$ denotes the probability of a jump from k to k' conditional on a jump occurring. Now the full log-likelihood of the sample of \bar{N} observations on the interval $[0, T]$ is simply

$$\ln L_{\bar{N}}^{\text{CT}}(\theta) = \sum_{n=1}^{\bar{N}} \ln \ell_n(\theta) + \ln [1 - G(T - t_{\bar{N}}, q(k_{\bar{N}}, k_{\bar{N}}; \theta))].$$

The final term is the probability of not observing an event on the interval $(t_{\bar{N}}, T]$.

With discrete-time data sampled at equispaced intervals Δ our sample takes the form of a collection of states $\{k_1, \dots, k_{\bar{N}}\}$ with \bar{N} observations. The likelihood function is simply

$$\ln L_{\bar{N}}^{\text{DT}}(\theta) = \sum_{n=2}^{\bar{N}} \ln P(k_{n-1}, k_n; \Delta, \theta),$$

where $P(k, l; \Delta, \theta)$ denotes the (k, l) element of the transition matrix induced by θ .

Other estimators for the model have been proposed. First, ABBE introduced a two-step PML (pseudo maximum likelihood) estimator, which is similar in spirit to the CCP estimator of [Hotz and Miller \(1993\)](#) for discrete time single agent models. More recently, [Blevins and Kim \(2019\)](#) note that the process of obtaining the two-step PML estimator can be iterated, in the spirit of [Aguirregabiria and Mira \(2007\)](#), to define a continuous time nested pseudo likelihood (CTNPL) estimator for the model, which is more stable than the discrete time counterpart. However, since the focus of this paper is identification of the continuous time model given discrete time data, rather than estimation, we focus on the simpler maximum likelihood estimator. This allows us to focus on the computational properties of the model and to examine how estimates behave when the sampling frequency of the data changes without concerns about two-step estimation error.

5.2. Single Agent Renewal Model

Here, we generate data according to the single agent binary choice model described above. The parameters of the model to be estimated are $\theta = (\lambda, \gamma, \beta, \mu)$ which include the move arrival rate λ , the rate of mileage increase γ , the mileage cost parameter β , and the engine replacement cost μ . We fix the number of mileage states at $K = 90$ and the discount rate at $\rho = 0.05$. The population parameters are $(\lambda, \gamma, \beta, \mu) = (1.0, 1.0, -3.0, -14.0)$. We also report the cost ratio μ/β which, as is common in discrete choice, is more precisely estimated in most specifications than β or μ individually.

M	Sampling		λ	γ	β	μ	μ/β
∞	DGP	True	1.000	1.000	-3.000	-14.000	4.667
50	Continuous	Mean	1.381	0.998	-3.031	-14.277	4.727
		S.D.	0.961	0.012	0.301	1.106	0.279
50	$\Delta = 1.00$	Mean	1.411	1.006	-3.053	-14.303	4.701
		S.D.	1.046	0.012	0.302	1.124	0.286
50	$\Delta = 8.00$	Mean	1.543	1.006	-3.105	-14.508	4.701
		S.D.	1.269	0.013	0.371	1.301	0.359
200	Continuous	Mean	1.129	1.000	-2.974	-13.974	4.705
		S.D.	0.435	0.006	0.177	0.585	0.168
200	$\Delta = 1.00$	Mean	1.123	1.008	-2.995	-13.989	4.678
		S.D.	0.441	0.006	0.179	0.603	0.163
200	$\Delta = 8.00$	Mean	1.276	1.008	-2.997	-14.049	4.696
		S.D.	0.822	0.007	0.208	0.788	0.212
800	Continuous	Mean	1.003	1.001	-3.010	-14.038	4.666
		S.D.	0.092	0.003	0.088	0.294	0.059
800	$\Delta = 1.00$	Mean	0.999	1.009	-3.031	-14.055	4.639
		S.D.	0.098	0.003	0.092	0.299	0.061
800	$\Delta = 8.00$	Mean	1.008	1.009	-3.027	-14.036	4.638
		S.D.	0.188	0.003	0.099	0.321	0.088
3200	Continuous	Mean	0.994	1.000	-2.998	-13.989	4.666
		S.D.	0.055	0.002	0.040	0.132	0.031
3200	$\Delta = 1.00$	Mean	0.988	1.008	-3.018	-13.997	4.639
		S.D.	0.054	0.002	0.040	0.136	0.031
3200	$\Delta = 8.00$	Mean	0.958	1.008	-3.028	-13.996	4.623
		S.D.	0.078	0.002	0.049	0.162	0.042

The mean and standard deviation are reported for 100 replications under several sampling regimes. For each replication, M markets were simulated over a fixed time interval $[0, T]$ with $T = 120$.

TABLE 1. Single Agent Renewal Model Monte Carlo Results

We use the full-solution maximum likelihood approach to estimate the model following Rust (1987). The value functions are obtained through value function iteration for each value of θ in an inner loop to within a tolerance of $\varepsilon = 10^{-6}$ under the supremum norm. We maximize 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 sampling regimes including full continuous-time data and discrete time data sampled at short and long intervals $\Delta = 1.0$ and $\Delta = 8.0$.

We simulate data over the interval $[0, T]$ with $T = 120$ for each of M markets, with M varying from 50 to 3200. We simulated both continuous time data and discrete time data with sampling intervals $\Delta = 1.0$ and $\Delta = 8.0$. For each specification, we report the mean and standard deviation of the parameter estimates over 100 replications in Table 1. All parameters are estimated quite precisely and with little finite-sample bias. The loss of precision from moving away from continuous time data is minimal at $\Delta = 1.0$ and more noticeable at $\Delta = 8.0$.

5.3. Quality Ladder Model

Our second set of Monte Carlo experiments corresponds to the quality ladder model described in Section 4. Table 2 summarizes the model specifications and computational time required for value function iteration. In this table we consider models ranging from $N = 2$ players and $K = 56$ states up to $N = 30$ players and $K = 58,433,760$ states. We hold the number of possible quality levels fixed at $\bar{\omega} = 7$, we set the quality of entrants to $\omega^e = 4$, and we increase the market size (\bar{M}) so that the average number of active players (n_{avg}) grows with the total number of possible players (N). We also report K , the number of states from the perspective of player i , which is the number of distinct (ω_i, ω) combinations in \mathcal{X} .

The final column of Table 2 compares the computational time required (wall clock time) for obtaining the value function across specifications. This step is necessary to either

N	$\bar{\omega}$	K	\bar{M}	n_{avg}	Obtain V
2	7	56	0.40	1.77	0.15
4	7	840	0.60	3.52	0.27
6	7	5,544	0.75	5.27	0.65
8	7	24,024	0.85	7.02	3
10	7	80,080	0.95	8.78	10
12	7	222,768	1.05	10.49	30
14	7	542,640	1.15	12.27	79
16	7	1,193,808	1.20	13.97	199
18	7	2,422,728	1.25	15.75	422
20	7	4,604,600	1.30	17.47	882
22	7	8,288,280	1.35	19.29	1648
24	7	14,250,600	1.40	20.98	2964
26	7	23,560,992	1.45	22.74	6481
28	7	37,657,312	1.50	24.47	10804
30	7	58,433,760	1.55	26.24	17712

N denotes the number of players (including potential entrants), $\bar{\omega}$ denotes the number of quality levels, \bar{M} denotes the market size, K denotes the total number of distinct states, n_{avg} denotes the average number of active firms, and ω_{avg} denotes the average quality level of active firms. Computational times are wall clock times using GNU Fortran 12.2 on a 2019 Mac Pro with a 2.5 GHz 28-Core Intel Xeon W processor.

TABLE 2. Quality Ladder Model Monte Carlo Specifications

generate a dataset or to simulate the model (e.g., to perform counterfactuals). We used value function iteration where the stopping criterion is that the choice probabilities are within a tolerance of $\varepsilon = 10^{-8}$ in the supremum norm.

To put the computational times in perspective, Doraszelski and Judd (2012) noted that it would take about *one year* to just solve for an equilibrium of a comparable⁹ 14-player game using the Pakes-McGuire algorithm. Similar computational times are reported in Doraszelski and Pakes (2007). However, it takes just over *one minute* to solve the continuous-time game with 14 players and 542,640 states. Even in the game with 30 players and over 58 million states, obtaining the value function took under 5 hours. We note that this would be infeasible for full-solution estimation, but when estimating the model using two-step methods, such as in ABBE or Blevins and Kim (2019), one may only need to carry out this step once, after estimation, for simulating a counterfactual. Overall, these computational

⁹The times reported by Doraszelski and Judd (2012) are for a model with $\bar{\omega} = 9$ but with no entry or exit, which for a fixed value of N , is roughly comparable in terms of dimensionality to our model with $\bar{\omega} = 7$, which includes entry and exit.

times suggest that a much larger class of problems can be estimated and simulated in the continuous-time framework.

Table 3 summarizes the results of our Monte Carlo experiments. We estimate the parameters $(\bar{\lambda}, \gamma, \kappa, \eta, \mu)$ where $\bar{\lambda} = N\lambda$ is the sum of move arrival rates across players. We do this so that the games are comparable under the same parameter values across specifications as the number of firms increases. Because we estimate firm fixed costs μ , we set the scrap value received upon exit equal to zero ($\varphi = 0$). The true parameter values, which are also shown in the table, are $(\bar{\lambda}, \gamma, \kappa, \eta, \mu) = (1.0, 0.4, 0.8, 6.0, 0.8)$.

N	K	Sampling		$\bar{\lambda}$	γ	κ	η	μ
		DGP	True	1.000	0.400	0.800	6.000	0.800
2	56	Continuous	Mean	1.000	0.399	0.800	6.048	0.799
			S.D.	0.013	0.010	0.026	0.200	0.020
		$\Delta = 1.0$	Mean	1.063	0.400	0.824	6.063	0.836
			S.D.	0.311	0.007	0.380	0.675	0.135
4	840	Continuous	Mean	0.999	0.398	0.799	6.016	0.799
			S.D.	0.011	0.013	0.024	0.176	0.016
		$\Delta = 1.0$	Mean	1.017	0.400	0.806	6.008	0.810
			S.D.	0.144	0.007	0.222	0.401	0.054
6	5,544	Continuous	Mean	1.000	0.397	0.802	6.036	0.799
			S.D.	0.010	0.016	0.024	0.185	0.017
		$\Delta = 1.0$	Mean	0.999	0.400	0.790	5.997	0.801
			S.D.	0.088	0.006	0.146	0.264	0.030
8	24,024	Continuous	Mean	1.000	0.397	0.797	6.029	0.801
			S.D.	0.010	0.016	0.022	0.191	0.017
		$\Delta = 1.0$	Mean	1.003	0.400	0.798	5.999	0.802
			S.D.	0.077	0.006	0.125	0.222	0.028
10	80,080	Continuous	Mean	0.999	0.398	0.800	6.044	0.800
			S.D.	0.009	0.017	0.021	0.157	0.016
		$\Delta = 1.0$	Mean	0.994	0.400	0.784	5.973	0.800
			S.D.	0.063	0.006	0.108	0.200	0.021

TABLE 3. Quality Ladder Model Monte Carlo Results

We first used samples containing $\bar{N} = 10,000$ continuous time events. In this case we observe the time of each event, the identity of the player, and the action chosen. For each specification we also report results for estimation with discrete time data with a fixed

sampling interval of $\Delta = 1.0$.¹⁰ In this case, we must calculate the matrix exponential of the Q matrix at each trial value of θ . To do so, we use the uniformization algorithm of Sherlock (2022). Because this matrix is high dimensional, but sparse, we adapted the algorithm to use sparse matrix methods and we precomputed the locations of the non-zero elements to improve the computational speed.

For each replication we used simulated annealing (Goffe, Ferrier, and Rogers, 1992, 1994) to maximize the log-likelihood function¹¹ and used $\varepsilon = 10^{-10}$ as the tolerance for value function iteration. Because of the time required to complete many replications of each specification, we limit our consideration to models up to $N = 10$ players and $K = 80,080$ states for the Monte Carlo experiments. Each replication involves an extensive global parameter search and each parameter evaluation solves a full solution problem for accuracy.¹² Although this is computationally costly, it allows us to focus on identification, computation, and estimation under time aggregation in a setting without additional tuning parameters and two-step estimation error.

While, the estimates are accurate and precise in all cases, we can see that the precision is decreased (standard errors are increased) due to the information lost with only discretely sampled data. Although the standard errors are larger than those with continuous time data, they are still reasonably small. Under this parameterization, we can see that more information is lost when the number of players is small. In this case the rate of move arrivals for each player is $\lambda = \bar{\lambda}/N$, so although the overall average number of events over a given interval is the same in both models, the equilibrium choice probabilities are such that firms choose to entry and exit more frequently when N is smaller. Hence, more information is lost over an interval $\Delta = 1.0$ relative to continuous time sampling, where

¹⁰To see the effects of varying the discrete time sampling interval Δ , please refer to the single agent model Monte Carlo experiments in the previous section.

¹¹For simulated annealing, we set the initial temperature to 0.01. We used an exponential decay schedule with parameter 0.70. The initial stepsizes were (1.0, 1.0, 1.0, 3.0, 1.0). The period for temperature reductions was 20 and dwell time between step size adjustments was 10. The step size adjustment factor was 2.0 and the function value tolerance, considering the previous three best values, was 10^{-3} . This resulted in about 15,000–20,000 log likelihood function evaluations per replication.

¹²To ease the computational burden, we store up to 100 previous value functions and associated parameter values. Then for each trial value of θ , we search for the closest (in Euclidean distance) previous parameter values and use the associated value function as the starting value for value function iteration.

the standard errors are relatively constant across N . This is not a general conclusion, but a feature of this particular data generating process.

6. Conclusion

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

A. Proofs

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 [ABBE](#), 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).

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

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

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

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

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

Proof of Theorem 2. Define the mapping $Y : [0, 1]^{N \times J \times K} \rightarrow [0, 1]^{N \times J \times K}$ by stacking best response probabilities:

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

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

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

zero as the eigenvalue. In this case, the number of required restrictions on Q is reduced to $\lfloor (K-1)/2 \rfloor$ because we know Q has at least one real eigenvalue.

Under the assumptions the number of distinct states in the model is $K \equiv K_0 \prod_{i=1}^N K_i$. Therefore, we will require at least $\lfloor \frac{K-1}{2} \rfloor$ linear restrictions of the form $R \text{vec}(Q) = r$ where R has full rank. We proceed by showing that the present model admits an intensity matrix Q with a known sparsity pattern and so we can use the locations of zeros as homogeneous restrictions, where r will be a vector of zeros.

Recall that each player has J choices, but $j = 0$ is a continuation choice. This results in $J-1$ non-zero off-diagonal elements per row of Q per player. There are at most $K_0 - 1$ non-zero off-diagonal elements due to exogenous state changes by nature. The only other non-zero elements of each row are the diagonal elements and therefore there are at least $K - N(J-1) - (K_0 - 1) - 1 = K_0 K_1^N - N(J-1) - K_0$ zeros per row of Q . The order condition we need to show is that the *total* number of zero restrictions is at least $\lfloor (K-1)/2 \rfloor$. For simplicity, it will suffice to show that there are $K/2 \geq \lfloor (K-1)/2 \rfloor$ restrictions. Summing across rows, this condition is satisfied when $(K_0 K_1^N)(K_0 K_1^N - N(J-1) - K_0) \geq K_0 K_1^N / 2$. Simplifying yields the sufficient condition in (13).

The derivative of the left-hand-side of (13) with respect to K_0 is $K_1^N - 1$. This value is always non-negative, since $K_1 \geq 1$, and is strictly positive when $K_1 > 1$. The derivative with respect to K_1 is $NK_0 K_1^{N-1}$. This value is always strictly positive since $K_0 \geq 1$ and $K_1 \geq 1$. Finally, the derivative with respect to J is $-N$. ■

References

- Agarwal, N., I. Ashlagi, M. Rees, P. Somaini, and D. Waldinger (2021). Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica* 87, 37–76.
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53.

- Arcidiacono, P., P. Bayer, J. R. Blevins, and P. B. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *Review of Economic Studies* 83, 889–931.
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370.
- Bajari, P., H. Hong, J. Krainer, and D. Nekipelov (2010). Computing equilibria in static games of incomplete information using the all-solution homotopy.
- Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica* 78, 453–508.
- Billingsley, P. (1961). *Statistical Inference for Markov Processes*. Chicago: University of Chicago Press.
- Blevins, J. R. (2017). Identifying restrictions for finite parameter continuous time models with discrete time data. *Econometric Theory* 33, 739–754.
- Blevins, J. R. and M. Kim (2019). Nested pseudo likelihood estimation of continuous-time dynamic discrete games. Working paper, The Ohio State University.
- Borkovsky, R. N., U. Doraszelski, and Y. Kryukov (2010). A user’s guide to solving dynamic stochastic games using the homotopy method. *Operations Research* 58(4), 1116–1132.
- Byrd, R. H., P. Lu, and J. Nocedal (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific and Statistical Computing* 16, 1190–1208.
- Caplin, A. and B. Nalebuff (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica* 59, 25–59.
- Chung, K. L. (1967). *Markov Chains with Stationary Transition Probabilities*. Berlin: Springer.

- Cosman, J. (2017). 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.
- Deng, Y. and C. F. Mela (2018). TV viewing and advertising targeting. *Journal of Marketing Research* 55, 99–118.
- 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.
- Goffe, W. L., G. D. Ferrier, and J. Rogers (1992). Simulated annealing: An initial application in econometrics. *Computational Economics* 5, 1572–9974.
- Goffe, W. L., G. D. Ferrier, and J. Rogers (1994). Global optimization of statistical functions with simulated annealing. *Journal of Econometrics* 60, 65–99.
- 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.
- Hansen, L. P. and T. J. Sargent (1991). Identification of continuous time rational expectations models from discrete time data. In *Rational Expectations Econometrics*, Chapter 9. Boulder, CO: Westview Press.
- Heckman, J. J. and B. Singer (1986). Econometric analysis of longitudinal data. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics*, Volume 3, Chapter 29. Amsterdam: North Holland.
- Hong, H., W. Li, and B. Wang (2015). Estimation of dynamic discrete models from time aggregated data. *Journal of Econometrics* 188(2), 435–446.

- Horn, R. A. and C. R. Johnson (1985). *Matrix Analysis*. New York: Cambridge University Press.
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–529.
- Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61, 265–289.
- Iskhakov, F., J. Rust, and B. Schjerning (2016). Recursive lexicographical search: Finding all Markov perfect equilibria of finite state directional dynamic games. *Review of Economic Studies* 83, 658–703.
- Jeziorski, P. (2022). Empirical model of dynamic merger enforcement—choosing ownership caps in U.S. radio. *Management Science*.
- Johansen, S. (1974). Some results on the imbedding problem for finite Markov chains. *Journal of the London Mathematical Society* s2-8, 345–351.
- Karlin, S. and H. M. Taylor (1975). *A First Course in Stochastic Processes* (Second ed.). San Diego, CA: Academic Press.
- Kessler, M. and A. Rahbek (2004). Identification and inference for multivariate cointegrated and ergodic Gaussian diffusions. *Statistical Inference for Stochastic Processes* 7, 137–151.
- Kim, M. (2021). Does the internet replace brick-and-mortar bank branches? Working paper, Ohio State University.
- 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.

- Mazur, J. (2017). Can stricter bankruptcy laws discipline capital investment? Evidence from the U.S. airline industry. Working paper, Purdue University.
- McCrorie, J. R. (2003). The problem of aliasing in identifying finite parameter continuous time stochastic models. *Acta Applicandae Mathematicae* 79, 9–16.
- McFadden, D. L. (1974). Conditional logit analysis of qualitative choice analysis. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press.
- Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political Economy* 92(6), 1086–1120.
- Nevskaya, Y. and P. Albuquerque (2019). How should firms manage excessive product use? A continuous-time demand model to test reward schedules, notifications, and time limits. *Journal of Marketing Research* 56, 379–400.
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* 54, 755–784.
- Pakes, A., G. Gowrisankaran, and P. McGuire (1993). Implementing the Pakes-McGuire algorithm for computing Markov perfect equilibria in Gauss. Working paper, Harvard University.
- Pakes, A. and P. McGuire (1994). Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *The RAND Journal of Economics* 25, 555–589.
- Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric Markov perfect equilibrium, and the ‘curse’ of dimensionality. *Econometrica* 69, 1261–1281.
- Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics* 38, 373–399.

- Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75, 901–928.
- Phillips, P. C. B. (1972). The structural estimation of a stochastic differential equation system. *Econometrica* 40, 1021–1041.
- Phillips, P. C. B. (1973). The problem of identification in finite parameter continuous time models. *Journal of Econometrics* 1, 351–362.
- Qin, M., M. A. Vitorino, and G. John (2022). Planes, trains, and co-opetition: Evidence from china. Working paper, INSEAD.
- 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.
- Sherlock, C. (2022). Direct statistical inference for finite Markov jump processes via the matrix exponential. *Computational Statistics* 36, 2863–2887.
- Sidje, R. B. (1998). Expokit: A software package for computing matrix exponentials. *ACM Transactions on Mathematical Software* 24, 130–156.
- Sims, C. A. (1971). Discrete approximations to continuous time distributed lags in econometrics. *Econometrica* 39, 545–563.
- Singer, B. and S. Spilerman (1976). The representation of social processes by Markov models. *The American Journal of Sociology* 82(1), 1–54.
- Takahashi, Y. (2015). Estimating a war of attrition: The case of the US movie theater industry. *American Economic Review* 105, 2204–2241.

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.