

# Leveraging Uniformization and Sparsity for Computation of Continuous Time Dynamic Discrete Choice Games

JASON R. BLEVINS

*The Ohio State University*

July 20, 2024

**Abstract.** Continuous-time formulations of dynamic discrete choice games offer notable computational advantages, particularly in modeling strategic interactions in oligopolistic markets. This paper extends these benefits by addressing computational challenges in order to improve model solution and estimation. We first establish new results on the rates of convergence of the value iteration, policy evaluation, and relative value iteration operators in the model, holding fixed player beliefs. Next, we introduce a new representation of the value function in the model based on uniformization—a technique used in the analysis of continuous time Markov chains—which allows us to draw a direct analogy to discrete time models. Furthermore, we show that uniformization also leads to a stable method to compute the matrix exponential, an operator appearing in the model’s log likelihood function when only discrete time “snapshot” data are available. We also develop a new algorithm that concurrently computes the matrix exponential and its derivatives with respect to model parameters, enhancing computational efficiency. By leveraging the inherent sparsity of the model’s intensity matrix, combined with sparse matrix techniques and precomputed addresses, we show how to significantly speed up computations. These strategies allow researchers to estimate more sophisticated and realistic models of strategic interactions and policy impacts in empirical industrial organization.

**Keywords:** Continuous time, Markov decision processes, dynamic discrete choice, dynamic stochastic games, uniformization, matrix exponential, computational methods, numerical methods.

**JEL Classification:** C63, C73, L13.

## 1. Introduction

Understanding the dynamics of market structure and competition in oligopolistic industries has been a foundational goal of modern empirical industrial organization. Beginning with the seminal works of [Bresnahan and Reiss \(1991\)](#) and [Berry \(1992\)](#), the field has transitioned from static models of market entry to dynamic frameworks that capture the forward-looking behaviors of firms, influenced by methodological advancements from [Aguirregabiria and Mira \(2007\)](#), [Bajari, Benkard, and Levin \(2007\)](#), [Pesendorfer and Schmidt-Dengler \(2008\)](#), and [Pakes, Ostrovsky, and Berry \(2007\)](#). These frameworks have expanded the analytical scope to include market power, entry deterrence, and industry evolution. However, the shift to dynamic strategic interactions introduced significant computational challenges, particularly with discrete time models where solving dynamic equilibrium models becomes complex as the number of players increases. Additionally, the focus on discrete time models with simultaneous moves has been driven by the prevalence of discrete time “snapshot” data. However, in reality often events unfold sequentially—whether observed or not—leading to much different strategic implications.

The introduction of continuous time models in dynamic discrete choice games by [Doraszelski and Judd \(2012\)](#) and [Arcidiacono, Bayer, Blevins, and Ellickson \(2016\)](#) was aimed at addressing the computational limitations of discrete time models. These models also offer a more granular representation of decision-making processes which unfold in a stochastic, sequential manner and align with the nature of strategic interactions in many real-world settings. [Doraszelski and Judd \(2012\)](#) showed that continuous time modeling simplifies the computation of players’ expectations over future states, thereby reducing the computational costs relative to discrete time models. [Arcidiacono, Bayer, Blevins, and Ellickson \(2016\)](#) developed an empirical framework and two-step CCP estimator based on the seminal contributions of [Hotz and Miller \(1993\)](#) and [Hotz, Miller, Sanders, and Smith \(1994\)](#). [Blevins \(2022\)](#) considered identification and estimation of the move arrival rate, which was previously held fixed, and further extended the model to allow for heterogeneity in move arrival rates. [Blevins and Kim \(2024\)](#) later extended the nested pseudo likelihood estimator of [Aguirregabiria and Mira \(2007\)](#) to the case of continuous time games.

Despite these advancements, the empirical application of continuous time models, particularly for estimating dynamic discrete choice games, still faces computational challenges. In particular, when researchers are limited to using snapshot data—a common scenario in many empirical contexts—efficient computation of the matrix exponential becomes critical in models with high-dimensional state spaces. The sparsity of the transition rate matrix (intensity matrix) is a key feature that our methods exploit to improve computational efficiency, not only in computing the matrix exponential but also in facilitating value iteration

or policy iteration. Snapshot data, which provides only periodic observations of the state variables without precise timing of transitions, necessitates the use of log likelihood functions that involve the matrix exponential operator. This underscores the importance of efficient computational methods for the matrix exponential.

This paper contributes to the literature by introducing theoretical and computational innovations to enhance the solution and estimation of continuous time dynamic discrete choice models. We begin with a review of the model and assumptions in Section 2. In Section 3, we begin by establishing the contractivity and modulus of contraction of the Bellman optimality operator for continuous time dynamic discrete choice games when beliefs are held fixed (including single agent models as a special case). Next, inspired by the method of uniformization, introduced by [Jensen \(1953\)](#) for continuous time Markov chain analysis, we develop a new uniform representation of the value function which allows us to draw direct comparisons to the discrete time dynamic discrete choice literature in order to establish rates of convergence for policy iteration, as well as relative value iteration following [Bray \(2019\)](#).

Next, we turn to computation of the log likelihood function with snapshot data in Section 4. We show how to apply uniformization as an efficient method for computing the matrix exponential in a way that is optimized for the sparse structure of transition rate matrices in the model. Accurate and fast evaluation of the matrix exponential is crucial for accurately estimating models using discrete time snapshot data. Additionally, we develop a new algorithm to augment the uniformization algorithm to allow for the simultaneous calculation of derivatives of the matrix exponential with respect to model parameters. This is vital for computationally efficient and accurate estimation using the log likelihood function.

In Section 5 we then focus on specific implementation details for computing the value function and the matrix exponential. Namely, we describe how to use sparse matrix methods in conjunction with precomputed addresses to improve computational efficiency for solving and estimating the model. Finally, we consider areas to apply parallelization.

To demonstrate the practical utility of our computational methods, in Section 6 we conduct Monte Carlo experiments on a dynamic entry-exit model with stochastic market demand. These experiments highlight the efficiency and accuracy of our approach, where we comparing the use of analytical versus numerical Jacobian derivatives in estimating the structural model parameters.

By identifying these computational efficiencies, we aim to enable researchers to more effectively estimate and simulate dynamic discrete choice models in continuous time, paving the way for deeper insights into strategic interactions and policy interventions in oligopolistic markets using more complex and realistic models.

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

We consider a continuous time dynamic discrete choice game involving  $N$  agents, indexed by  $i = 1, \dots, N$ , who operate in an infinite, continuous time horizon indexed by  $t \in [0, \infty)$ .

*State Space* The state space in the model defines all relevant market configurations at any point in time. It encapsulates both exogenous and endogenous variables that influence decision-making and strategic interactions among the agents.

At any instant  $t$ , the state of the system is represented by a state vector  $x$  within a finite state space  $\mathcal{X}$ . This vector includes all payoff-relevant information that is common knowledge among the agents. Each state vector  $x^k \in \mathcal{X}$ , where  $k$  indexes the discrete states, represents a specific combination of attributes and conditions of the market and the agents therein. For analytical tractability and practical implementation, the state space is assumed to be finite, with  $K = |\mathcal{X}|$  denoting the total number of unique states. Furthermore, we will generally prefer to work with the linearized state space  $\mathcal{K} = \{1, \dots, K\}$  and will refer to states by their index  $k$  rather than as a vector  $x$ , as this will allow us to vectorize many expressions.

The components of the state vector can vary depending on the specific application but typically include variables such as agent-specific attributes and market conditions. For example,  $x$  may include each agent's current market entry status, product quality, technology level, or any other characteristic that affects their decisions and payoffs. Usually  $x$  also contains exogenous variables representing the economic environment such as measures of market demand.

*Decisions & Endogenous State Changes* Agents, indexed by  $i$ , make decisions from a set of available actions,  $\mathcal{A} = 0, 1, \dots, J - 1$ , at decision times that occur stochastically. The timing of these decisions is governed by Poisson processes with rate parameters  $\lambda_i$ , which can vary by agent, reflecting the frequency of decision opportunities for each agent. We assume  $\lambda_i < \infty$ , reflecting the cost of attention or other exogenous frictions or bounds on the rate of action.

The outcome of an agent's decision is a deterministic change in the state of the system. Each action  $j$  chosen by agent  $i$  in state  $k$  leads to a transition to a new state, denoted by  $l(i, j, k)$ . This function links each decision with its effect on the state vector.

For example, in a market entry and exit model, an agent's decision to enter in a given state might lead to an increase in the number of active firms within the market, altering the competitive dynamics and, thus, the state vector to reflect this new market configuration.

The rate at which these endogenous state transitions occur is a function of the decision rate  $\lambda_i$  and the conditional choice probabilities  $\sigma_{ijk}$ , which represent the probability of agent  $i$  choosing action  $j$  in state  $k$ . The product  $\lambda_i \sigma_{ijk}$  gives the transition rate for each possible action, integrating the stochastic nature of decision timing with the strategic selection of actions.

We assume that agents in the model are forward-looking, considering the future implications of their current decisions. The decision-making process and resultant state transitions adhere to Markov properties, ensuring that the future evolution of the system depends only on the current state and the decisions made within it.

The strategic decision-making process and the resulting endogenous state changes are central to determining the equilibrium of the model. Agents formulate their strategies based on their beliefs about the actions of others and the dynamics of the environment, seeking to maximize their expected payoffs over time. The equilibrium concept employed is Markov Perfect Equilibrium (MPE), where agents' strategies form a Nash equilibrium in every state of the model, taking into account both the endogenous and exogenous dynamics of state changes. We will define the equilibrium concept below, after describing the remaining components of the model.

*Exogenous State Changes* Exogenous state changes capture the evolution of the environment due to factors outside the control of the agents. In the model, an artificial player referred to as “nature,” indexed by  $i = 0$ , is responsible for exogenous state changes. Nature’s actions, much like those of the agents, result in transitions between states in the state space  $\mathcal{X}$ . However, unlike agent-driven transitions, these changes are not the result of strategic choices but rather represent random events or systematic shifts in the external environment.

Exogenous state changes are modeled using a Markov jump process characterized by an intensity matrix  $Q_0$ . For  $k \neq l$ , the  $(k, l)$  element of this matrix, denoted  $q_{0kl}$ , represents the rates at which the system transitions away from state  $k$  to state  $l$  due to exogenous factors. For example,  $q_{0kl}$  could represent the rate at which a technological advancement occurs, altering the competitive dynamics among firms.

The transition rate  $q_{0kl}$  could be zero, indicating that direct, instantaneous transitions from state  $k$  to state  $l$  do not occur. Consequently, the aggregate rate at which the system transitions away from state  $k$  given by the sum  $\sum_{l \neq k}^K q_{0kl}$ . Accordingly, the  $k$ -th diagonal element,  $q_{0kk}$ , is defined to be minus this rate:

$$(1) \quad q_{0kk} = - \sum_{l=1}^K q_{0kl}.$$

Therefore,  $Q_0$  is itself a proper intensity matrix with zero row sums.

*Payoffs* In the continuous time setting we distinguish between flow payoffs and instantaneous payoffs. Flow payoffs  $u_{ik}$  accrue constantly to agent  $i$  while the model remains in state  $k$ . They are continuous in nature, reflecting the ongoing benefit or cost to the player as long as the state of the environment remains unchanged.

In contrast to the continuous accrual of flow payoffs, instantaneous choice-specific payoffs, denoted by  $c_{ijk}$ , are incurred at the moment player  $i$  chooses action  $j$  in state  $k$ . We assume these payoffs are additively separable into a deterministic component,  $\psi_{ijk}$ , representing the mean payoff associated with choosing action  $j$  in state  $k$ , and a stochastic component,  $\varepsilon_{ijk}$ , which is a choice-specific unobserved payoff shock observed only by the player. This separation allows for the modeling of uncertainty and private information in the decision-making process.

*Assumptions* Before turning to the equilibrium, we state the main assumptions of the model, following Arcidiacono, Bayer, Blevins, and Ellickson (2016) and Blevins (2022):

**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  $\lambda_i$  for each player  $i = 1, \dots, N$  and  $q_{0kl}$  for exogenous state changes from each state  $k$  to  $l \neq k$  due to nature, with  $0 < \lambda_i < \infty$ ,  $0 \leq q_{0kl} < \infty$ , and  $\sum_{l \neq k} q_{0kl} > 0$ .

**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  $c_{ijk} = \psi_{ijk} + \varepsilon_{ijk}$ .

**Assumption 6** (Costless Continuation & 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  $\varepsilon_{ik}$  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  $\mathbb{R}^J$ .

The previously mentioned Assumptions 1–5 lay the groundwork for the model’s structure and dynamics. Assumption 6 clarifies the role of the action  $j = 0$ , designating it as a neutral continuation action that carries no direct cost and does not alter the current state. This particular action can be likened to selecting the “outside option” in demand models, serving as a baseline or reference choice when no other actions within the model are taken. This assumption is needed for identification, ensuring that the choice to remain in the current state is distinct and carries no immediate payoff. Moreover, Assumption 6 stipulates that any action  $j > 0$  must lead to a tangible and potentially observable change in the state, thereby ensuring that all actions within the choice set are substantially different from one another. Assumption 7 describes our assumptions on the idiosyncratic error terms in the model and mirrors similar assumptions used in discrete time models (Aguirregabiria and Mira, 2010).

*Strategies, Value Functions, and Best Responses* A strategy for a player in this model is a rule that specifies the action the player will take in every possible state of the world. More formally, a strategy for player  $i$ , denoted  $\delta_i$ , maps each state  $k \in \mathcal{K}$  and a vector of choice-specific shocks  $\varepsilon_{ik} \in \mathbb{R}^J$  to an action in the choice set  $\mathcal{A}$ . Any policy  $\delta_i$  determines associated choice probabilities

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

for all choices  $j$  and states  $k$ .

To determine an optimal strategy, we must consider player  $i$ ’s beliefs  $\varsigma_i$  about the actions of other players, reflecting the strategic interdependence inherent in the game. Let  $\varsigma_{im}$  represent the beliefs held by player  $i$  about the choices made by rival  $m$ , characterized by a set of probabilities across  $J$  actions and  $K$  states for each possible state  $k$  and action  $j$ . Given beliefs

$$(3) \quad \varsigma_i \equiv (\varsigma_{i,1}, \dots, \varsigma_{i,i-1}, \varsigma_{i,i+1}, \dots, \varsigma_{i,N})$$

about each rival, a best response for player  $i$  is a strategy that maximizes the player’s expected utility, given the strategies (or equivalently, the choice probabilities) of all other players.

Based on the beliefs  $\varsigma_i$  that player  $i$  holds, we can describe the value function as a vector of length  $K$ , denoted  $V_i(\varsigma_i) = (V_{i1}(\varsigma_i), \dots, V_{iK}(\varsigma_i))^\top$ . Here, each element,  $V_{ik}(\varsigma_i)$ , represents the expected present value of future rewards beginning in state  $k$  and making optimal decisions thereafter, considering the beliefs  $\varsigma_i$ . For a brief time span  $\tau$ , following Assumption 3, the likelihood of an event with rate  $\lambda_i$  happening is  $\lambda_i\tau$ . With  $\rho_i$  as the

discount rate, the discount factor over this interval is  $1/(1 + \rho_i\tau)$ . Consequently, during the interval  $\tau$ , the expected present value while in state  $k$  simplifies to:

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

The term  $o(\tau)$  accounts for the diminishingly small probabilities of multiple events occurring within  $\tau$ , which are negligible as  $\tau$  approaches zero. After simplification and as  $\tau$  tends towards zero, we derive the recursive formula for  $V_{ik}$  under player  $i$ 's beliefs  $\varsigma_i$  as follows:

$$(4) \quad V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{0kl} V_{il} + \sum_{m \neq i} \lambda_m \sum_j s_{imjk} V_{i,l(m,j,k)} + \lambda_i \mathbb{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \right\}}{\rho_i + \sum_{l \neq k} q_{0kl} + \sum_m \lambda_m}.$$

This representation shows the balance between the immediate flow payoff in state  $k$ , the expected values following exogenous and endogenous state transitions, and the anticipated outcomes of player  $i$ 's decisions, all adjusted for the probability of occurrence within a short interval and discounted appropriately. The expectation is over the joint distribution of  $\varepsilon_{ik} = (\varepsilon_{i,0,k}, \dots, \varepsilon_{i,J-1,k})^\top$ .

**Definition.** Let  $T_{i,\varsigma_i}$  denote the *Bellman optimality operator* for player  $i$  with beliefs  $\varsigma_i$ ,

$$V_i = T_{i,\varsigma_i} V_i,$$

defined by stacking (4) across states  $k = 1, \dots, K$ .

Mathematically, a strategy  $\delta_i(k, \varepsilon_{ik})$  is a best response if it assigns an action  $j$  that maximizes the expected utility in state  $k$ , given the shock  $\varepsilon_{ik}$  and the player's beliefs about others' strategies  $\varsigma_i$ :

$$(5) \quad \delta_i(k, \varepsilon_{ik}) = j \iff \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)}(\varsigma_i) \geq \psi_{ij'k} + \varepsilon_{ij'k} + V_{i,l(i,j',k)}(\varsigma_i) \quad \forall j' \in \mathcal{J}.$$

In this paper we focus on the concept of Markov perfect equilibrium (MPE) as defined for continuous-time dynamic games by Arcidiacono, Bayer, Blevins, and Ellickson (2016). They also demonstrate the existence of equilibrium, providing a basis for applying this definition in the present framework. Although the probability of simultaneous moves is zero in this framework, which effectively eliminates one source of multiplicity, Blevins and Kim (2024) showed that in general there may still be multiple equilibria.



### 2.1. Continuous Time Markov Chains

The reduced form of the model is a finite state continuous time Markov chain, also known as a Markov jump process. The process is a stochastic process  $X(t)$  for  $t \in [0, \infty)$ , with values in the state space  $\mathcal{X}$ . Observing this process from an arbitrary starting point  $t$  in state  $X(t)$ , it persists in this state for a randomly determined duration  $\tau$ , subsequently transitioning to a different state  $X(t + \tau)$ . This duration,  $\tau$ , is known as the holding time, and the process's evolution over time can be described as a piecewise-constant, right-continuous function of time characterized by jumps at intervals governed by a Poisson process, making the holding times exponentially distributed. We provide here a brief overview of key results on Markov jump processes. For a comprehensive discussion, refer to [Karlin and Taylor \(1975, Section 4.8\)](#), [Tijms \(2003, Chapter 4\)](#), or [Chung \(1967, part II\)](#).

We will represent such a processes using its *intensity matrix* or *infinitesimal generator matrix*, which succinctly captures the process's dynamics. Consider a generic intensity matrix  $Q$ , defined as follows:

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

where, for  $k \neq l$ ,

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

represents the instantaneous rate at which transitions occur from state  $k$  to state  $l$ . The diagonal elements,  $q_{kk}$ , are set to  $-\sum_{l \neq k} q_{kl}$ , ensuring that the sum of each row is zero. The rate parameter for the exponential distribution of holding times before exiting state  $k$  is given by  $-q_{kk}$ , equating to the aggregate of off-diagonal rates of transition out of state  $k$ . When the process exits state  $k$ , it transitions to a different state  $l \neq k$  with the probability ratio  $q_{kl} / \sum_{l \neq k} q_{kl} = -q_{kl} / q_{kk}$ .

Often we need to calculate the probabilities of transitioning from one state to another over a specific time interval,  $\Delta$ . For example, we need to calculate these probabilities to estimate model parameters using snapshots of data sampled at uniformly-spaced intervals of length  $\Delta$  (e.g., one year or one month). This leads us to the concept of the *transition probability matrix*  $P(\Delta)$ , which contains the probabilities of moving from each state  $k$  to every other state  $l$  after time  $\Delta$ .

For a continuous time Markov chain with intensity matrix  $Q$ , the transition probability

matrix  $P(\Delta)$  for a time interval  $\Delta$  is given by the matrix exponential of  $\Delta Q$ , denoted  $\exp(\Delta Q)$ , defined by its Taylor series expansion:

$$(6) \quad P(\Delta) = \exp(\Delta Q) = \sum_{j=0}^{\infty} \frac{(\Delta Q)^j}{j!} = I + \Delta Q + \frac{(\Delta Q)^2}{2!} + \frac{(\Delta Q)^3}{3!} + \dots,$$

where  $I$  is the identity matrix. The  $(k, l)$  element of  $P(\Delta)$ , denoted as  $p_{ij}(\Delta)$ , represents the probability that the process transitions from state  $k$  to state  $l$  over the interval  $\Delta$ .

Despite its theoretical clarity, directly calculating the matrix exponential can be challenging, particularly for large matrices. Computing the transition probabilities via (6) directly is not ideal for several reasons. First, it requires the computation of an infinite series, which, in practice, must be truncated to a finite number of terms. Determining the proper truncation point is important. Computing matrix powers can be computationally intensive and numerically unstable, especially for matrices that are not well-conditioned. In many applications, such as continuous time dynamic discrete games, the intensity matrices are sparse. The direct approach does not exploit this sparsity in computing powers of  $P$ , leading to unnecessary computational work and higher storage requirements.

Various numerical methods have been developed to approximate the matrix exponential, such as Padé approximations and scaling and squaring methods. In their classic article, (Moler and Loan, 1978) discuss “nineteen dubious ways” to compute the matrix exponential. However, the uniformization method we use in this paper, despite being introduced much earlier by Jensen (1953) was not included. (Moler and Loan, 1978) considered methods for computing  $\exp(A)$ , the exponential of a general matrix  $A$ . Uniformization is tailored for the special case of computing the exponential of a generator matrix  $Q$  of a continuous time Markov chain. This distinction is important and, as it turns out, uniformization is especially effective in the even more special case where we need to calculate the action of the matrix exponential on a probability vector  $v$ , such as  $\exp(\Delta Q)v$ . In particular, it will be easy to exploit sparsity in this situation, which is a central feature of the models we consider in this paper.

## 2.2. Example: A Single-Agent Renewal Model

Here we introduce a continuous-time single-agent renewal model. Inspired by Rust (1987), we consider the decision-making process of a bus company manager. The manager faces a choice between replacing a bus engine and continuing with the existing one. This scenario is modeled using a state variable representing the accumulated mileage of the bus engine, expressed as  $k$  in a discrete set  $\mathcal{K} = \{1, \dots, K\}$ .

The decisions of the manager (player  $i = 1$ ) are binary. The manager may continue

without replacement ( $j = 0$ ), after which the state remains unchanged, or replace the engine ( $j = 1$ ), after which the mileage resets to the initial state  $k = 1$ .

The transitions for these decisions are represented by the function:

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

The hazard rate for replacing the engine in state  $k$  is denoted by  $h_{1k}$ , while the non-replacement decision rate  $h_{0k}$  complements it to maintain a constant overall decision rate  $\lambda$ , such that  $h_{0k} = \lambda - h_{1k}$ .

The mileage accrues according to a simple transition process, with a rate parameter  $\gamma$  governing the increase to the next mileage state (one state ahead). If we let  $Q_0 = (q_{kl})$  denote the intensity matrix for nature (player  $i = 0$ ), where the elements are  $q_{kl}$ , then this can be represented as:

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

The operating cost for a bus with mileage  $k$  is  $\beta k$ , where  $\beta$  is a negative constant reflecting the cost per mileage unit. Upon deciding to continue, no immediate cost is incurred, but engine replacement entails a fixed cost  $\mu$ .

The value function  $V_k$  in this model can be expressed recursively as

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

where  $\rho$  is the discount rate, and  $\varepsilon_{ijk}$  denotes iid shocks received following choice  $j$ .

The dynamics of the model are described by the aggregate intensity matrix  $Q$ , which is the sum of two matrices:  $Q_0$ , representing nature-induced transitions, and  $Q_1$ , representing agent-induced transitions.

To make things concrete, let us consider a simple example with  $K = 5$  mileage states. The matrix  $Q_0$ , which represents the mileage transition due to natural increase, is structured

as follows for  $K = 5$ :

$$(7) \quad Q_0 = \begin{bmatrix} -\gamma & \gamma & 0 & 0 & 0 \\ 0 & -\gamma & \gamma & 0 & 0 \\ 0 & 0 & -\gamma & \gamma & 0 \\ 0 & 0 & 0 & -\gamma & \gamma \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Here,  $\gamma$  represents the rate at which the mileage of a bus increases by exactly one unit. The diagonal elements are negative (indicating the rate of leaving the current state), and the immediately off-diagonal elements are positive, indicating the probability of moving to the next higher mileage state. The last row has zeros because it is assumed there are no further mileage increases beyond state 5 in this model setup.

The matrix  $Q_1$ , which represents transitions induced by the agent's decision to replace the bus engine, is

$$(8) \quad Q_1 = \begin{bmatrix} 0 & 0 & \dots & 0 \\ \lambda\sigma_{12} & -\lambda\sigma_{12} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \lambda\sigma_{1K} & 0 & \dots & -\lambda\sigma_{1K} \end{bmatrix},$$

where  $\sigma_{1k}$  denotes the probability of choosing to replace the engine in state  $k$ . Each non-zero element in  $Q_1$  represents a rate of making a decision to replace the engine, resetting the mileage to 1, from any given state  $k$ .

Adding  $Q_0$  and  $Q_1$  together, we obtain the aggregate matrix  $Q$  for  $K = 5$ :

$$(9) \quad Q = \begin{bmatrix} -\gamma & \gamma & 0 & 0 & 0 \\ \lambda\sigma_{12} & -\lambda\sigma_{12} - \gamma & \gamma & 0 & 0 \\ \lambda\sigma_{13} & 0 & -\lambda\sigma_{13} - \gamma & \gamma & 0 \\ \lambda\sigma_{14} & 0 & 0 & -\lambda\sigma_{14} - \gamma & \gamma \\ \lambda\sigma_{15} & 0 & 0 & 0 & -\lambda\sigma_{15} \end{bmatrix}.$$

The aggregate matrix  $Q$  is clearly sparse for several reasons. Most transitions occur between adjacent states (due to natural increase) or reset to the initial state (due to replacement). There are no transitions to non-adjacent states except for those induced by replacements, which only affect the first column. Many elements of  $Q$  are zero because transitions between non-adjacent states are not possible, except through replacements leading back to state 1.

This sparsity is beneficial because it means that in computations one can skip over

the zero elements and focus on the non-zero elements. This reduces the complexity of matrix operations, leading to faster computational times and lower memory usage, especially important in larger systems.

### 2.3. Uniformization of a Continuous Time Markov Chain

Uniformization (Jensen, 1953), also known as randomization, is a long-used technique in the study of continuous time Markov chains, with applications to areas such as queuing theory, operations research, and epidemiology.<sup>1</sup> Recall that a continuous time Markov chain on a finite state space is characterized by transitions between states that occur continuously over time, with the time between transitions following an exponential distribution. The transition rates between states are described by a generator matrix  $Q$ . A discrete time Markov chain involves transitions between states at discrete time steps, with transition probabilities described by a matrix  $P$ .

Uniformization provides a straightforward way to convert a continuous time Markov chain, with potentially varying exit rates, into a comparable discrete time Markov chain subordinate to a Poisson process. Transitions of the discrete time process are not uniformly spaced, rather they are determined by the hitting times of the Poisson process. The rate parameter  $\eta$  of the Poisson process is determined by the highest exit rate of the continuous time process, ensuring that  $\eta$  is sufficient to capture all possible transitions in the continuous-time Markov chain.

A key aspect of the discrete-time Markov chain is the introduction of self-transitions, which permit the process to remain in its current state with a non-zero probability. This modification is necessary to preserve the overall transition rates of the original continuous time Markov chain.

Uniformization involves the following steps:

1. Select a uniform rate  $\eta$ : Choose  $\eta$  larger than the absolute values of the generator matrix  $Q$ 's diagonal elements:

$$\eta > \max_{k=1,\dots,K} |q_{kk}|.$$

This rate equalizes transition rates across states, which can be viewed as following a Poisson process with rate  $\eta$ .

---

<sup>1</sup>Other early contributions on uniformization include Grassmann (1977a,b) and Reibman and Trivedi (1988). See (see e.g., Puterman, 2005, Ch. 11) and van Dijk, van Brummelen, and Boucherie (2018) for more recent overviews.

2. Construct the transition probability matrix  $\Sigma$ : Formulate  $\Sigma$  using  $\eta$  and  $Q$ :

$$\Sigma = I + \frac{Q}{\eta}.$$

$\Sigma$  contains the probabilities of state transitions at each arrival of the Poisson process.

With this uniform representation, we can view potential transitions within the original continuous time Markov chain as being governed by the Poisson process's jump times with rate  $\eta$ . At every jump, the system may shift between states following the transition probabilities in  $\Sigma$ , which includes possible self-transitions.

This approach is particularly useful for computing transient state probabilities for continuous-time Markov chains as well as the transition probabilities which appear in the log likelihood function of such models in settings with discrete time data. The probability of transitioning from state  $k$  to state  $l$  in time  $\Delta$  is given by

$$(10) \quad p_{kl}(\Delta) = e^{-\eta\Delta} \sum_{j=0}^{\infty} \frac{(\eta\Delta)^j}{j!} [\Sigma^j]_{kl},$$

where the sum iteratively calculates the probabilities of transitioning from state  $k$  to state  $l$  in exactly  $j$  jumps. Here,  $[\Sigma^j]_{kl}$  is probability of making such a transition after  $j$  jumps. Each term of the series is weighted by the Poisson probability of exactly  $j$  arrivals within the interval  $\Delta$ , as determined by a Poisson process with rate  $\eta\Delta$ .

**Example 1** (Two-State Model). Consider a simple continuous time Markov chain with two states, represented by the infinitesimal generator matrix

$$Q = \begin{pmatrix} -\alpha & \alpha \\ \beta & -\beta \end{pmatrix}.$$



Here  $\alpha$  is the rate of transitioning from state 1 to state 2 and  $\beta$  is the rate of transitioning from state 2 to state 1. This process is represented in the top panel of Figure 1.

To apply uniformization, we first choose a uniform rate  $\eta > \max\{\alpha, \beta\}$ , ensuring it covers the highest transition rate in  $Q$ . With  $\eta$  selected, we construct the uniform transition probability matrix  $\Sigma$  as follows:

$$\Sigma = I + \frac{Q}{\eta} = \begin{pmatrix} 1 - \frac{\alpha}{\eta} & \frac{\alpha}{\eta} \\ \frac{\beta}{\eta} & 1 - \frac{\beta}{\eta} \end{pmatrix}.$$

This process is represented in the bottom panel of Figure 1. The off-diagonal entries,  $\frac{\alpha}{\eta}$  and  $\frac{\beta}{\eta}$ , are the adjusted probabilities of transitioning between states 1 and 2 under the uniform

Continuous-time Markov chain  $Q$



Discrete-time Markov chain  $\Sigma$   
subordinate to a Poisson process with rate  $\eta$

FIGURE 1. Example of Uniformization

rate  $\eta$ . Because transitions occur at rate  $\eta$ , the overall rates of transitioning between the two states remains unchanged. The diagonal entries of  $\Sigma$ ,  $1 - \frac{\alpha}{\eta}$  and  $1 - \frac{\beta}{\eta}$ , represent the probabilities of self-transitions for states 1 and 2, respectively, ensuring that  $\Sigma$  is a proper row-stochastic matrix. This transformation leverages a Poisson process to maintain the stochastic behavior of the continuous-time process in a discrete-time framework, simplifying analysis and computation while preserving the essential dynamics of the original process.

#### 2.4. Uniformization of the Single Agent Renewal Model

Consider the continuous time Markov chain with five states from the example single agent renewal model with  $Q$  matrix given by (9). The dynamics of this model are characterized by exogenous mileage increments that occur at rate  $\gamma$  and decisions to reset the mileage by replacing the bus engine at rate  $\lambda\sigma_{1k}$  in any given state  $k$ .

To apply the uniformization technique, we select a uniform rate  $\eta > \lambda \max_k \sigma_{1k} + \gamma$ , ensuring it surpasses the highest rate in  $Q$ . We note that since the  $\sigma_{1k}$  for  $k = 1, \dots, K$  are probabilities that are strictly bounded between 0 and 1 under Assumption 7, it suffices to choose

$$\eta = \lambda + \gamma,$$

equal to the rate of player moves plus the rate of state changes by nature. With  $\eta$  selected, the uniform transition probability matrix  $\Sigma$  is constructed as follows:

$$\Sigma = I + \frac{Q}{\eta} = \begin{bmatrix} 1 - \frac{\gamma}{\eta} & \frac{\gamma}{\eta} & 0 & 0 & 0 \\ \frac{\lambda\sigma_{12}}{\eta} & 1 - \frac{\lambda\sigma_{12} + \gamma}{\eta} & \frac{\gamma}{\eta} & 0 & 0 \\ \frac{\lambda\sigma_{13}}{\eta} & 0 & 1 - \frac{\lambda\sigma_{13} + \gamma}{\eta} & \frac{\gamma}{\eta} & 0 \\ \frac{\lambda\sigma_{14}}{\eta} & 0 & 0 & 1 - \frac{\lambda\sigma_{14} + \gamma}{\eta} & \frac{\gamma}{\eta} \\ \frac{\lambda\sigma_{15}}{\eta} & 0 & 0 & 0 & 1 - \frac{\lambda\sigma_{15}}{\eta} \end{bmatrix}.$$

This model is represented in Figure 2, where the top panel depicts the original continuous-time Markov chain and the bottom panel illustrates the discrete-time Markov chain transformed via uniformization. The off-diagonal entries of  $\Sigma$  are the adjusted probabilities of transitioning between states under the uniform rate  $\eta$ , while the diagonal entries represent the probabilities of self-transitions, ensuring  $\Sigma$  is a proper stochastic matrix. This transformation simplifies computation and analysis while maintaining the dynamics of the original process.

### 3. Computing the Value Function

#### 3.1. Value Iteration

We begin by first establishing some properties of the Bellman optimality operator  $T_{i,\varsigma_i}$  defined in (4) above. In particular, we show that holding beliefs fixed, the operator is a contraction mapping with respect to the supremum norm.

**Theorem 1** (Value Iteration). *Suppose Assumptions 1–7 hold. Then for any firm  $i$ , for fixed beliefs  $\varsigma_i$ , the Bellman optimality operator  $T_{i,\varsigma_i}$  defined in (4) is a contraction mapping with respect to the supremum norm with modulus*

$$(11) \quad \underline{\beta}_i \equiv \frac{\bar{\eta}}{\rho_i + \underline{\eta}},$$


where

$$(12) \quad \bar{\eta} \equiv \sum_m \lambda_m + \max_k \sum_{l \neq k} q_{0kl} \quad \text{and} \quad \underline{\eta} \equiv \sum_m \lambda_m + \min_k \sum_{l \neq k} q_{0kl}.$$


Let  $V_i^{(k)} \equiv T_{i,\varsigma_i}^k V_i^{(0)}$  denote the  $k$ -th step of value iteration for any initial guess  $V_i^{(0)} \in \mathbb{R}^K$ . Then,  $\lim_{k \rightarrow \infty} V_i^{(k)} = V_{i,\varsigma_i}^*$ . Furthermore, for all  $k$ ,

$$\|V_i^{(k)} - V_{i,\varsigma_i}^*\|_{\infty} \leq \underline{\beta}_i^k \|V_i^{(0)} - V_{i,\varsigma_i}^*\|_{\infty}$$





Continuous-time Markov chain  $Q$



Discrete-time Markov chain  $\Sigma$  subordinate to a Poisson process with rate  $\eta$

FIGURE 2. Example of Uniformization for the Single Agent Renewal Model

and so the rate of convergence of  $V_i^{(k)}$  to the unique fixed point  $V_{i,\varsigma_i}^*$  is  $O_p(\underline{\beta}_i^k)$ .

The proof of this theorem, and all remaining results, is given in Appendix A.

In games with multiple agents, the beliefs  $\varsigma_i$  represent the strategies of the other agents  $\sigma_{-i}$ . The theorem covers single agent models as a special case, where the beliefs  $\varsigma_i$  are not necessary as there are no rival players to consider.

### 3.2. Uniform Representation of the Value Function

Next, we establish a novel “uniform” representation of the value function. Although the representation is based on the CCP inversion result of Hotz and Miller (1993) as applied to the continuous time setting by Arcidiacono, Bayer, Blevins, and Ellickson (2016), it is a new representation based on the method of uniformization described above in Section 2.3. First, the following Lemma establishes that  $\bar{\eta}$  in (13) is a valid uniformization rate for the model.

**Lemma 1.** *Suppose Assumptions 1–7 hold in the general dynamic discrete choice model. Then  $\bar{\eta}$  defined in (12) is a valid uniformization rate for the continuous time Markov chain defined by the aggregate intensity matrix  $Q$  for the model. That is,*

$$(13) \quad \bar{\eta} > \max_k |q_{kk}|,$$

where  $q_{kk}$  is the  $(k, k)$  diagonal element of  $Q$ .

Next, Theorem 2 establishes the new uniform representation of  $V_i$ . The representation defines a linear operator in terms of equilibrium choice probabilities  $\sigma$  and, as uniformization produces a discrete time Markov chain, the representation allows us to draw direct comparisons to discrete time dynamic discrete choice models.

**Theorem 2.** *Suppose Assumptions 1–7 hold. For equilibrium choice probabilities  $\sigma$ , consisting of equilibrium beliefs  $\varsigma_i = \sigma_{-i}$  and an equilibrium best response policy  $\sigma_i$ , the value function  $V_i$  of the model has a representation in terms of a uniform Bellman policy operator as:*

$$(14) \quad V_i = T_{i,\sigma} V_i \equiv U_i(\sigma) + \bar{\beta}_i \Sigma(\sigma) V_i$$

where

$$(15) \quad U_i(\sigma) \equiv \frac{1}{\rho_i + \bar{\eta}} [u_i + \lambda_i C_i(\sigma)],$$

$$(16) \quad \bar{\beta}_i \equiv \frac{\bar{\eta}}{\rho_i + \bar{\eta}},$$

with  $C_i(\sigma)$  denoting the  $K \times 1$  vector containing the ex-ante expected value of the instantaneous payoff  $c_{ijk} = \psi_{ijk} + \varepsilon_{ijk}$  for each state  $k$  given the choice probabilities  $\sigma$ . The rate  $\bar{\eta}$  is given in Lemma 1.

In the representation of Theorem 2, we can see that  $U_i(\sigma)$  plays the effective role of the flow utility. It combines the actual flow utility with the rate-weighted expected instantaneous payoff at the agent's next decision time. The effective discount factor is  $\bar{\beta}_i$ , which depends on the hazards of exogenous and endogenous state changes that determine  $\bar{\eta}$ , relative to the discount rate  $\rho_i$ . Under Assumptions 2 and 3,  $\bar{\beta}_i$  is strictly bounded between 0 and 1. Finally, note that due to uniformization the matrix  $\Sigma(\sigma)$  is a stochastic matrix. It becomes the effective transition matrix in this representation.

The vector  $C_i(\sigma)$  is easy to compute in practice. It is a  $k$ -vector such that the  $k$ -th element equals  $\sum_{j=0}^{J-1} \sigma_{ijk} [\psi_{ijk} + e_{ijk}(\sigma)]$ , where  $e_{ijk}(\sigma)$  is the expected value of  $\varepsilon_{ijk}$  given that action  $j$  is chosen:

$$(17) \quad e_{ijk}(\sigma) \equiv \frac{1}{\sigma_{ijk}} \int \varepsilon_{ijk} \cdot \mathbf{1} \left\{ \varepsilon_{ij'k} - \varepsilon_{ijk} \leq \psi_{ijk} - \psi_{ij'k} + V_{l(i,j,k)} - V_{l(i,j',k)} \quad \forall j' \right\} f(\varepsilon_{ik}) d\varepsilon_{ik}.$$

While (17) involves a multivariate integral over the joint distribution of  $\varepsilon_{ik}$ , Aguirregabiria and Mira (2002, 2007) showed that there exist closed-form expressions for the integral for two commonly used error distributions.

1. In the case where the errors  $\varepsilon_{ik}$  are independently and identically distributed according to the type 1 extreme value distribution with a scale parameter  $\sigma_\varepsilon$ , the expression simplifies to

$$e_{ijk}(\sigma) = \gamma_{EM} - \sigma_\varepsilon \ln \sigma_{ijk}$$

where  $\gamma_{EM} \approx 0.577215665$  is the Euler-Mascheroni constant.

2. In a binary choice model with  $J = 2$  and  $\varepsilon_{ik} \sim N(0, \Omega)$ , we have

$$e_{ijk}(\sigma) = \frac{\text{Var}(\varepsilon_{ijk}) - \text{Cov}(\varepsilon_{i0k}, \varepsilon_{i1k}) \varphi(\Phi^{-1}(\sigma_{iijk}))}{\sqrt{\text{Var}(\varepsilon_{i1k} - \varepsilon_{i0k})} \sigma_{iijk}},$$

where  $\Phi$  and  $\varphi$  are the standard normal cumulative distribution function and probability density function, respectively.

The value function for a fixed equilibrium beliefs  $\varsigma_i = \sigma_{-i}$  and best response policy  $\sigma_i$  is the unique fixed point to the uniform Bellman policy operator defined in (14). Given the uniform representation in (14), the theorem also establishes that the modulus of contraction for policy evaluation using  $T_{i,\sigma}$  to find  $V_{i,\sigma}$  is  $\bar{\beta}_i$ .

Importantly, this uniform representation of the value function is not just about converting the continuous time model into a discrete time counterpart. The uniform discrete time Markov chain represents instantaneous transitions and remains sparse like the original continuous time model. This is in contrast to a discrete time representation *over a time interval*  $\Delta$ , or a standard simultaneous-move discrete time model (with period length  $\Delta = 1$ ), in which the transition matrix would become more dense.

Finally, although the present model is different than the model of Doraszelski and Judd (2012), in Appendix B we also show how to apply uniformization in their framework.

### 3.3. A Single Agent Model

In Section 2.4, we showed how to apply uniformization to the aggregate intensity matrix of a simple single agent renewal model. We will return to this idea later, in the context of estimating the model with discrete time data, where we will need to calculate the matrix exponential, for example, following the algorithms we introduce later in Section 4.

Here, we illustrate uniformization in terms of the value function for the model. This allows us to focus on the case of a single agent  $i = 1$  and nature  $i = 0$ . In addition, we will derive the representation in a different way to add intuition. Rather than deriving the uniform representation of the full  $Q$  matrix, as in the proof of Theorem 2, we will first apply uniformization to each individual  $Q_i$  matrix:  $Q_0$  then  $Q_1$ . As before, let  $Q_0$  be the intensity matrix for exogenous state changes due to nature. For  $\eta_0 > \max_k |q_{kk}|$ , we have the uniform representation  $\Sigma_0 = I + \frac{1}{\eta_0}Q_0$  or  $Q_0 = \eta_0(\Sigma_0 - I)$ , where  $\Sigma_0$  is the transition matrix induced by the intensity matrix  $Q_0$  when events by nature occur at rate  $\eta_0$  and we allow self-transitions. We note that  $\eta_1 = \lambda_1$  serves as a uniform rate for  $Q_1$ . For the value function  $V$  and utility function  $u$ , we will drop the  $i$  subscript for simplicity.

Then, we can express the Bellman operator for a policy  $\sigma$  in matrix form as:

$$\begin{aligned} T_\sigma(V) &= \frac{1}{\rho + \eta_0 + \eta_1} \{[u + \eta_1 C(\sigma)] + [\eta_0 \Sigma_0 + \eta_1 \Sigma_1(\sigma)] V\} \\ &= U(\sigma) + \beta \Sigma(\sigma) V \end{aligned}$$

where we define

$$\begin{aligned} U(\sigma) &\equiv \frac{1}{\rho + \eta_0 + \eta_1} [u + \eta_1 C(\sigma)], \\ \beta &\equiv \frac{\eta_0 + \eta_1}{\rho + \eta_0 + \eta_1}, \text{ and} \\ \Sigma(\sigma) &\equiv \frac{\eta_0}{\eta_0 + \eta_1} \Sigma_0 + \frac{\eta_1}{\eta_0 + \eta_1} \Sigma_1(\sigma). \end{aligned}$$

We can see that we have arrived at the same representation as in Theorem 2. To

reach this result, we first defined  $\Sigma_0$  as the discrete-time transition matrix in a uniform representation of the continuous time process for nature determined by  $Q_0$ . Similarly, due to the model assumptions of a common  $\eta_1$  determining decision times for the agent, and the assumption that the choice set contains a continuation choice  $j = 0$ , the implied transition matrix at decision times  $\Sigma_1(\sigma)$  is itself also a discrete time transition matrix resulting from uniformization of the intensity matrix  $Q_1(\sigma)$  associated with the agent’s decision problem.

The combined effective transition matrix is  $\Sigma$ , which is a weighted sum of  $\Sigma_0$ , relating to exogenous transitions due to nature, and  $\Sigma_1(\sigma)$ , the endogenous transitions induced by actions of the agent. The weights on the individual matrices are non-negative and sum to one and therefore can be interpreted as conditional probabilities (i.e., the probability that the next event is an exogenous or endogenous state change).

### 3.4. Newton-Kantorovich Iteration

The idea behind the Newton-Kantorovich method is to solve for the value function  $V$  by iteratively solving the system of equations  $V = T(V)$ . This method is a central part of the NFXP method of Rust (1987) for single agent dynamic discrete choice models in discrete time. In our setting, we could use the same approach either for value iteration for a single agent model or to solve for equilibria locally, either for multiple starting values or for pre-estimated choice probabilities. Generally, the idea is to restate the problem as one of iteratively solving for a zero of the operator  $(I - T)$ . Given  $V^{(k)}$ , the Newton-Kantorovich iteration is

$$(18) \quad V^{(k+1)} = V^{(k)} - [I - T']^{-1} (I - T)V^{(k)}$$

where  $T'$  denotes the Fréchet derivative of the operator  $T$ .

The Newton-Kantorovich method converges quadratically when the initial guess is sufficiently close to the true solution. Therefore, as in Rust (1987), we can use value iteration (contraction iterations) initially until we get within a sufficiently small neighborhood of the solution and then switch to Newton-Kantorovich iterations. This poly-algorithm formed the basis of the NFXP method. It can be used here, for example, to find a unique solution for  $V$  for single agent models, or a unique solution for  $V_i$  in a game for fixed beliefs  $\varsigma_i$ .

### 3.5. Relative Value Iteration

In this section we establish strong convergence results similar to Bray (2019) but in the case continuous time models. The representation derived above by uniformization in (14) maps directly into Bray (2019)’s framework. As such, when the stochastic process induced by  $\Sigma(\sigma)$  is ergodic, relative differences in the value function converge at the rate  $\bar{\beta}_i \gamma_2$ . Here,  $\gamma_2$  is bounded below by the second-largest eigenvalue of  $\Sigma(\sigma)$ , with this eigenvalue being

strictly smaller than one, therefore  $\bar{\beta}_i \gamma_2 < \bar{\beta}_i$ . The following theorem formalizes this.

**Theorem 3** (Relative Value Iteration). *Suppose Assumptions 1–7 hold and let  $\sigma$  denote a profile of equilibrium choice probabilities. Let  $\Sigma(\sigma)$  denote the stochastic transition matrix derived from the uniformization process, as described in Theorem 2, and let  $\bar{\beta}_i$  be the effective discount factor. If  $\Sigma(\sigma)$  is ergodic, then the relative policy evaluation for computing the value function  $V_i$  converges geometrically at a rate given by:*

$$\left| V_i^{(k)} - V_{i,\sigma}^* \right|_{\infty} \leq (\bar{\beta}_i \gamma_2)^k \left| V_i^{(0)} - V_{i,\sigma}^* \right|_{\infty}$$

where  $\gamma_2$  is the second-largest eigenvalue of  $\Sigma(\sigma)$ , and  $V_{i,\sigma}^*$  is the true value function associated with the equilibrium policies  $\sigma$ .

#### 4. Computing the Log Likelihood Function and its First Derivatives

We can consider two scenarios for estimation: continuous time data and discrete time data. We will primarily consider the discrete time case, as it is the most difficult. In this setting, the researcher only observes snapshots of a cross-section of markets at regularly-spaced time intervals of length  $\Delta > 0$ . Let  $M$  denote the number of markets and  $N_m$  the number of snapshots over time observed in market  $m = 1, \dots, M$ . The sample can be represented as

$$\{k_{m,0}, \dots, k_{m,N_m}\}_{m=1}^M.$$

Let  $P(\Delta, \theta) \equiv \exp(\Delta Q(\theta))$  denote the transition matrix over the interval  $\Delta$  and let  $P(\Delta, \theta)_{k,l}$  denote the  $(k, l)$  element of the transition matrix. Let

$$d_{kl} \equiv \sum_{m=1}^M \sum_{n=1}^{N_m} 1 \{k_{m,n-1} = k, k_{m,n} = l\}$$

denote the number of observed transitions from state  $k$  to  $l$ .

The log likelihood function is

$$\begin{aligned}
\ell(\theta) &= \ln \prod_{m=1}^M \prod_{n=1}^{N_m} P(\Delta, \theta)_{k_{m,n-1}, k_{m,n}} \\
&= \ln \prod_{k=1}^K \prod_{l=1}^K P(\Delta, \theta)_{k,l}^{d_{kl}} \\
&= \sum_{k=1}^K \sum_{l=1}^K \ln P(\Delta, \theta)_{k,l}^{d_{kl}} \\
&= \sum_{k=1}^K \sum_{l=1}^K d_{kl} \ln P(\Delta, \theta)_{k,l} \\
(19) \quad &= \sum_{l=1}^K d'_l [\ln P(\Delta, \theta)] e_l
\end{aligned}$$

where  $d_l = [d_{1l}, \dots, d_{Kl}]'$  and  $e_l$  is the  $l$ -th basis vector. In practice, we only need to sum over states  $l$  where  $d_l > 0$ , so evaluating the log likelihood function may involve calculating relatively few columns of  $P(\Delta, \theta)$ .

#### 4.1. Uniformization and the Matrix Exponential

Choose  $\eta > \max_{k=1, \dots, K} |Q_{kk}|$ ,  $\Sigma \equiv Q/\eta + I$ , so that  $\Delta Q = \eta \Delta \Sigma - \eta \Delta I$ .

By properties of the matrix exponential, we note that

$$\begin{aligned}
\exp(\Delta Q) &= \exp(\eta \Delta \Sigma - \eta \Delta I) \\
&= \exp(\eta \Delta \Sigma) \exp(-\eta \Delta I) \\
&= \exp(\eta \Delta \Sigma) e^{-\eta \Delta} I \\
&= e^{-\eta \Delta} \exp(\eta \Delta \Sigma). \\
(20) \quad &= e^{-\eta \Delta} \sum_{j=0}^{\infty} \frac{(\eta \Delta)^j \Sigma^j}{j!}.
\end{aligned}$$

Importantly, all elements of  $\Sigma$  are positive, meaning that this calculation will not suffer from cancellation of alternating positive and negative terms (Goldberg, 1991; Sherlock, 2022). As a result, computations will be much more numerically stable for the uniformization of a rate matrix  $Q$  than for the exponential of a generic matrix  $A$ .

In many cases, we merely need to calculate the action of the matrix exponential on a

probability vector  $v$ :

$$\begin{aligned}
\exp(\Delta Q)v &= \exp(\eta\Delta\Sigma - \eta\Delta I)v \\
&= \exp(\eta\Delta\Sigma)\exp(-\eta\Delta I)v \\
&= \exp(\eta\Delta\Sigma)e^{-\eta\Delta}Iv \\
&= e^{-\eta\Delta}\exp(\eta\Delta\Sigma)v. \\
(21) \quad &= e^{-\eta\Delta}\sum_{j=0}^{\infty}\frac{(\eta\Delta)^j\Sigma^jv}{j!}.
\end{aligned}$$

Notice that we can compute (21) in a recursive manner as

$$\begin{aligned}
\exp(\Delta Q)v &= e^{-\eta\Delta}\sum_{j=0}^{\infty}\frac{(\eta\Delta)^j\Sigma^jv}{j!} \equiv e^{-\eta\Delta}\sum_{j=0}^{\infty}\nu_j \\
\nu_0 &= v \\
\nu_j &= \frac{\eta\Delta\Sigma}{j}\nu_{j-1}.
\end{aligned}$$

In practice, we need to truncate the infinite series at some finite  $\bar{J}_\varepsilon$  such that the approximation error is smaller than a given tolerance  $\varepsilon$ :

$$\exp(\Delta Q)v = e^{-\eta\Delta}\sum_{j=0}^{\bar{J}_\varepsilon}\frac{(\eta\Delta)^j\Sigma^jv}{j!} + e^{-\eta\Delta}\sum_{j=\bar{J}_\varepsilon+1}^{\infty}\frac{(\eta\Delta)^j\Sigma^jv}{j!}.$$

The truncation introduces approximation error, which must be managed to ensure the resultant probabilities are accurate. Given the nature of the uniform representation with a Poisson process to regularize the intervals between state transitions, we can determine the truncation point,  $\bar{J}_\varepsilon$ , using the tail probabilities of a Poisson distribution (Fox and Glynn, 1988; Reibman and Trivedi, 1988). Importantly,  $\Sigma^jv$  can be interpreted as the state distribution after  $j$  steps in the DTMC, given an initial distribution  $v$ . Therefore each term of the series is itself a probability vector with  $\|\Sigma^jv\|_1 = 1$ . This simplifies calculation of the truncation point since the norm of the second term, representing the amount of probability mass lost due to truncation, can be bounded by  $\varepsilon$  by choosing  $\bar{J}_\varepsilon$  to be the smallest integer such that

$$e^{-\eta\Delta}\sum_{j=\bar{J}_\varepsilon+1}^{\infty}\frac{(\eta\Delta)^j}{j!} = 1 - \text{PoissonCDF}(\bar{J}_\varepsilon; \eta\Delta) < \varepsilon,$$

where  $\text{PoissonCDF}(j; \eta\Delta)$  is the cumulative distribution function of a Poisson distribution



with rate  $\eta\Delta$  evaluated at  $j$ . Conveniently, the truncation point  $\bar{J}_\varepsilon$  is deterministic and can be calculated in advance as

$$\bar{J}_\varepsilon = \text{PoissonCDF}^{-1}(1 - \varepsilon; \eta\Delta).$$

The basic uniformization algorithm for computation of  $\exp(\Delta Q)v$  is presented in pseudocode form as Algorithm 1.

---

**Algorithm 1** Uniformization Algorithm for Matrix Exponential

---

**Require:**  $Q$  (rate matrix),  $\Delta$  (time),  $v$  (initial probability vector),  $\varepsilon$  (tolerance)

**Ensure:**  $\mu$  (resulting probability vector  $\exp(\Delta Q)v$ )

```

1: function EXPMV( $Q, \Delta, v, \varepsilon$ )
2:    $\eta \leftarrow \max(\text{abs}(\text{diag}(Q))) + \varepsilon$ 
3:    $\tilde{\Sigma} \leftarrow \Delta Q + \eta\Delta I$  ▷ Note that  $\tilde{\Sigma} = \eta\Delta\Sigma$ .
4:    $\bar{J}_\varepsilon \leftarrow \text{poissinv}(1 - \varepsilon, \eta\Delta)$ 
5:    $\mu \leftarrow v$ 
6:    $\nu \leftarrow v$ 
7:   for  $j \leftarrow 1$  to  $\bar{J}_\varepsilon$  do
8:      $\nu \leftarrow \tilde{\Sigma}\nu/j$ 
9:      $\mu \leftarrow \mu + \nu$ 
10:  end for
11:   $\mu \leftarrow \exp(-\eta\Delta)\mu$ 
12:  return  $\mu$ 
13: end function

```

---

Sherlock (2022) discusses two important details regarding this algorithm. First, he introduces a method for accurately calculating the required tail probabilities for small  $\varepsilon$  values. Second, he modifies the uniformization algorithm with an adjustment term to account for possible overflow in the calculation of  $\sum_{j=0}^J \frac{(\lambda\Delta)^j}{j!}$ .

#### 4.2. Simultaneous Computation of First Derivatives

To use a gradient-based optimization method to estimate the model, we also need to take the derivative of the log likelihood function with respect to the model parameters  $\theta$ ,  $\frac{\partial \ell}{\partial \theta}$ . Note that the aggregate intensity matrix  $Q$  in our model is a function of the model primitives—the rates of exogenous state changes and the implied hazards of the various endogenous actions that can be taken by the agents in the model. These are in turn functions of the model parameters  $\theta$ . We shall write the matrix as  $Q(\theta)$  to emphasize this. Since  $Q = Q(\theta)$  the derivative  $\frac{\partial Q}{\partial \theta}$  is a tensor. Therefore, for clarity we focus derivatives with respect to a single

scalar component  $\alpha$  of  $\theta$ , which we will denote by  $\frac{\partial \ell}{\partial \alpha}$ . By the chain rule, we have

$$\begin{aligned}
\frac{\partial \ell}{\partial \alpha} &= \frac{\partial}{\partial \alpha} \sum_{k=1}^K \sum_{l=1}^K d_{kl} \ln P(\Delta, \theta)_{k,l} \\
&= \sum_{k=1}^K \sum_{l=1}^K d_{kl} \frac{\partial}{\partial \alpha} \ln P(\Delta, \theta)_{k,l} \\
(22) \quad &= \sum_{k=1}^K \sum_{l=1}^K \frac{d_{kl}}{P(\Delta, \theta)_{k,l}} \frac{\partial P(\Delta, \theta)_{k,l}}{\partial \alpha}
\end{aligned}$$

(23)

Thus, we also need to compute  $\frac{\partial P(\Delta, \theta)}{\partial \alpha}$ .

From (20), we have

$$(24) \quad \frac{\partial P(\Delta, \theta)}{\partial \alpha} = e^{-\eta \Delta} \sum_{j=0}^{\infty} \frac{(\eta \Delta)^j}{j!} \frac{\partial \Sigma^j}{\partial \alpha}.$$

Suppose we can calculate  $\frac{\partial \Sigma}{\partial \alpha}$  analytically. Then we can calculate  $\frac{\partial \Sigma^j}{\partial \alpha}$  recursively by noting that  $\Sigma^j = \Sigma \Sigma^{j-1}$  and therefore, by the product rule,

$$(25) \quad \frac{\partial \Sigma^j}{\partial \alpha} = \frac{\partial \Sigma}{\partial \alpha} \Sigma^{j-1} + \Sigma \frac{\partial \Sigma^{j-1}}{\partial \alpha}.$$

As before, when we only need the action of  $\frac{\partial P(\Delta, \theta)}{\partial \alpha}$  on a vector  $v$ ,

$$(26) \quad \frac{\partial P(\Delta, \theta)}{\partial \alpha} v = e^{-\eta \Delta} \sum_{j=0}^{\infty} \frac{(\eta \Delta)^j}{j!} \frac{\partial \Sigma^j}{\partial \alpha} v,$$

then this will again be computationally simpler. Note that from (25), we have the following recurrence:

$$(27) \quad \frac{\partial \Sigma^j}{\partial \alpha} v = \frac{\partial \Sigma}{\partial \alpha} \Sigma^{j-1} v + \Sigma \frac{\partial \Sigma^{j-1}}{\partial \alpha} v.$$

Substituting (27) into (26), we obtain

$$\begin{aligned}
\frac{\partial P(\Delta, \theta)}{\partial \alpha} v &= e^{-\eta \Delta} \sum_{j=0}^{\infty} \left[ \frac{(\eta \Delta)^j}{j!} \frac{\partial \Sigma}{\partial \alpha} \Sigma^{j-1} v + \frac{(\eta \Delta)^j}{j!} \Sigma \frac{\partial \Sigma^{j-1}}{\partial \alpha} v \right] \\
&= e^{-\eta \Delta} \sum_{j=0}^{\infty} \left[ \frac{\eta \Delta}{j} \frac{\partial \Sigma}{\partial \alpha} \times \underbrace{\frac{(\eta \Delta)^{j-1}}{(j-1)!} \Sigma^{j-1} v}_{=\nu_{j-1}} + \frac{\eta \Delta}{j} \Sigma \times \underbrace{\frac{(\eta \Delta)^{j-1}}{(j-1)!} \frac{\partial \Sigma^{j-1}}{\partial \alpha} v}_{\equiv \delta_{j-1}} \right] \\
&= e^{-\eta \Delta} \sum_{j=0}^{\infty} \frac{\eta \Delta}{j} \left[ \frac{\partial \Sigma}{\partial \alpha} \nu_{j-1} + \Sigma \delta_{j-1} \right] \\
&= e^{-\eta \Delta} \sum_{j=0}^{\infty} \delta_j,
\end{aligned}$$

where the last line follows from (27) and we can calculate  $\nu_j$  and  $\delta_j$  recursively as follows:

$$\begin{aligned}
\nu_0 &= \delta_0 = v, \\
\nu_j &= \frac{\eta \Delta}{j} \Sigma \nu_{j-1} \text{ for } j = 1, 2, \dots \\
\delta_j &= \frac{\eta \Delta}{j} \left[ \frac{\partial \Sigma}{\partial \alpha} \nu_{j-1} + \Sigma \delta_{j-1} \right] \text{ for } j = 1, 2, \dots
\end{aligned}$$

*Remark.* In recursions involving sparse matrix products, matrices stored can be subject to fill-in, increasing the storage required substantially. Importantly, for both  $\nu_j$  and  $\delta_j$  we only need to store  $K \times 1$  vectors. The only matrices stored are the sparse matrices  $\Sigma$  and  $\frac{\partial \Sigma}{\partial \alpha}$ . Each iteration of the recursion only involves three sparse matrix-vector products.

*Remark.* Note that the term  $\nu_j$  is also used for computing  $\exp(\Delta Q)v$ , so there is a benefit to computing the matrix exponential times  $v$  and its derivatives simultaneously.

*Remark.* When  $\theta$  contains multiple parameters, we can compute the derivatives simultaneously in the same loop simply by storing separate vectors  $\delta_{j,\alpha}$  for each component  $\alpha$ . The term  $\nu_j$  is independent of  $\alpha$ .

## 5. Sparse Matrices, Precomputed Addresses, and Parallelization

Sparse matrices are matrices in which most of the elements are zero. Storing and operating on sparse matrices efficiently can be beneficial in models such as ours where sparsity naturally occurs in the value function representation and the rate matrix  $Q$ .

There are several common storage formats for sparse matrices that allow efficient storage and computations. These include coordinate (COO) format, compressed sparse row (CSR)

---

**Algorithm 2** Uniformization Algorithm for Matrix Exponential and Derivatives

---

**Require:**  $Q$  (rate matrix),  $D_\alpha = \frac{\partial \Sigma}{\partial \alpha}$  (derivative),  $\Delta$  (time),  $v$  (initial probability vector),  $\varepsilon$  (tolerance)

**Ensure:**  $\mu$  (result of  $\exp(\Delta Q)v$ ),  $\mu_\alpha$  (result of  $\frac{\partial \exp(\Delta Q)}{\partial \alpha}v$ )

1: **function** EXPMVD( $Q, D, \Delta, v, \varepsilon$ )

2:    $\eta \leftarrow \max(\text{abs}(\text{diag}(Q))) + \varepsilon$

3:    $\tilde{\Sigma} \leftarrow \Delta Q + \eta \Delta I$

4:    $\tilde{D}_\alpha \leftarrow \eta \Delta D_\alpha$

5:    $\bar{J}_\varepsilon \leftarrow \text{poissinv}(1 - \varepsilon, \eta \Delta)$

6:    $\nu \leftarrow v$

7:    $\mu \leftarrow \nu$

8:    $\delta \leftarrow v$

9:    $\mu_\alpha \leftarrow 0_{K \times 1}$

10:   **for**  $j \leftarrow 1$  **to**  $\bar{J}_\varepsilon$  **do**

11:      $\nu \leftarrow \tilde{\Sigma} \nu / j$

12:      $\mu \leftarrow \mu + \nu$

13:      $\delta \leftarrow (\tilde{D}_\alpha \nu + \tilde{\Sigma} \delta) / j$

14:      $\mu_\alpha \leftarrow \mu_\alpha + \delta$

15:   **end for**

16:    $\mu \leftarrow \exp(-\eta \Delta) \times \mu$

17:    $\mu_\alpha \leftarrow \exp(-\eta \Delta) \times \mu_\alpha$

18:   **return**  $\mu, \mu_\alpha$

19: **end function**

---

▷ Note that  $\tilde{\Sigma} = \eta \Delta \Sigma$ .

▷ Note that  $\tilde{D}_\alpha = \eta \Delta \frac{\partial \Sigma}{\partial \alpha} = \Delta \frac{\partial Q}{\partial \alpha}$ .

format, and compressed sparse column (CSC) format. These are perhaps best explained through a simple example.

Consider again the  $Q$  matrix from the single agent renewal example in (9) with  $K = 5$  states, with specific values for the parameters:

$$Q = \begin{bmatrix} -\gamma & \gamma & 0 & 0 & 0 \\ \lambda\sigma_{12} & -\lambda\sigma_{12} - \gamma & \gamma & 0 & 0 \\ \lambda\sigma_{13} & 0 & -\lambda\sigma_{13} - \gamma & \gamma & 0 \\ \lambda\sigma_{14} & 0 & 0 & -\lambda\sigma_{14} - \gamma & \gamma \\ \lambda\sigma_{15} & 0 & 0 & 0 & -\lambda\sigma_{15} \end{bmatrix} = \begin{bmatrix} -0.2 & 0.2 & 0 & 0 & 0 \\ 0.1 & -0.3 & 0.2 & 0 & 0 \\ 0.4 & 0 & -0.6 & 0.2 & 0 \\ 0.6 & 0 & 0 & -0.8 & 0.2 \\ 0.9 & 0 & 0 & 0 & -0.9 \end{bmatrix}.$$

Note that this matrix has 25 elements, but only 13 nonzero values to store. For matrices this small, because we will also have to store the row and column locations, sparse matrix storage is not more efficient. However, for larger matrices, the savings can be substantial. For the model with  $K = 100$  states, there are  $K \times K = 100 \times 100 = 10,000$  elements in the matrix but only  $3 \times 100 - 2 = 298$  nonzero elements. Using sparse matrix methods will have a significant impact on both storage (i.e., required to hold the matrix in memory) and computation (e.g., the number of operations required to multiply the matrix by a vector). Computational savings occur because only the relatively small number of nonzero elements of the matrix are involved in any given computation.

The COO format represents a sparse matrix by storing lists of the row indices, column indices, and corresponding nonzero values. In this format, each of the arrays has the same length, equal to the number of nonzeros in the matrix. The COO representation of the example  $Q$  matrix is shown in Figure 3.

Values	-0.2	0.2	0.1	-0.3	0.2	0.4	-0.6	0.2	0.6	-0.8	0.2	0.9	-0.9
Row Indices	1	1	2	2	2	3	3	3	4	4	4	5	5
Column Indices	1	2	1	2	3	1	3	4	1	4	5	1	5

FIGURE 3. Sparse  $Q$  Matrix in COO format

The COO format is straightforward, listing each non-zero element of the matrix by its row and column indices along with the corresponding value. This format is particularly useful for initially constructing sparse matrices due to its simplicity and direct way of mapping matrix elements.

The COO format, offers straightforward indexing capabilities. This format allows direct

access to any non-zero element along with its position, making it ideal for applications that require frequent insertion, deletion, or modification of elements where knowing the exact position is crucial. COO format does not inherently optimize row or column access but provides equal facility for accessing data in any order, which is beneficial for random access and indexing.

The CSR format stores the values of each row of the matrix, the column indices of those values, and a separate array of pointers to indicate the index in the value array of where each row begins. The CSC representation of the example  $Q$  is illustrated in Figure 4.




FIGURE 4. Sparse  $Q$  Matrix in CSR Format

CSC format is similar to CSR format, but stores the values of the matrix in column order, along with the row indices of the values. There is a separate array of column pointers indicating the index in the value array where each column begins. The example  $Q$  matrix can be stored in CSC format as illustrated in Figure 5.




FIGURE 5. Sparse  $Q$  Matrix in CSC Format

The choice of storage format impacts both storage efficiency and ease of operations like matrix-vector products. CSC format is preferable for vector-matrix multiplication (where the vector is multiplied on the left of the matrix as in  $v'Q$ ), since the operation involves accumulating contributions from columns of the matrix  $Q$  scaled by corresponding elements of  $v$ . The CSC format stores data by columns, which makes it efficient for accessing all

elements of a particular column sequentially. Likewise, CSR format is faster for matrix-vector multiplication due to the values being stored row-wise. Finally, COO format is generally fastest for indexing—looking up values at specific locations—since it allows direct access to any non-zero element along with its position. COO format is ideal for applications that require frequent insertion, deletion, or modification of elements where knowing the exact position is crucial.

Therefore, it is important to understand the implementation details of your sparse matrix library to understand which operations may be faster, allowing you to potentially reorganize your programs or data structures accordingly. MATLAB, for instance, uses the CSC format internally [Gilbert, Moler, and Schreiber \(1992\)](#). If your sparse matrix library uses the CSC format and you are, say, implementing the log likelihood function using the sparse matrix exponential algorithm, you may wish to arrange the computation as  $v' \exp(\Delta Q)$  if possible—which involves vector-matrix products in the Taylor series—instead of  $\exp(\Delta Q)v$ —which involves a series of matrix-vector products.

Finally, we note that switching formats is possible and may be advantageous for certain operations. For instance, you may want to construct the initial sparse matrix in COO format and then convert to CSR format for subsequent matrix-vector multiplications.

### 5.1. Computing $Q$

In large-scale models involving multiple players, computational efficiency can significantly impact both performance and feasibility. One of the critical strategies we employ involves the use of sparse matrices with precomputed addresses of nonzero elements. These techniques are exceptionally beneficial in dynamic discrete choice games, where the locations of state transitions—stemming from player actions and natural state changes—can be calculated in advance and retained for subsequent computations.

Sparse matrices primarily store non-zero elements, optimizing both storage space and computation time. In our approach, we initially construct a sparse matrix  $Q$  in COO (coordinate) format. This format is particularly advantageous for constructing the matrix as it allows direct placement of the elements (see [Figure 3](#)). For instance, in an entry-exit model, each state’s transitions resulting from firms’ entry and exit decisions, as well as changes induced by nature, are considered. For each potential transition, we record the corresponding row and column indices that connect the current state to the future state.

Precomputing the locations of these non-zero entries represents a significant optimization. By establishing the sparsity pattern of the matrix—which delineates where the non-zero entries are located—we ensure that this structural blueprint can be reused. This reuse is viable as long as the underlying structure of the matrix remains unchanged, even if the structural parameters of the model, and consequently the values within the matrix, do

change. As a result, the computational overhead associated with reassessing the matrix structure each time the model is updated is substantially reduced. The algorithm simply updates the values at these pre-established locations without recalibrating the entire matrix layout.

After constructing the matrix in COO format, we convert it to CSR (compressed sparse row) format. This transformation is designed to enhance the efficiency of subsequent computations, particularly those involving matrix-vector multiplications, which are prevalent in iterative methods used for computing value functions and updating choice probabilities.

When the model parameters are updated, updates to the sparse  $Q$  matrix are confined to the value array. The precomputed column index and row pointer arrays (see Figure 4) remain constant, as the sparsity pattern of the matrix, indicating the possible transitions, does not change. This method is more efficient because it circumvents the computationally expensive task of recalculating the transition matrix structure with each alteration in model parameters.

## 5.2. Computing $V_i$

Once we have a convenient sparse matrix representation of the intensity matrix  $Q$ , the vectorized representation of the value function  $V_i$  allows for efficient computation of both value iteration and Newton-Kantorovich iteration using sparse matrix operations.

In value iteration, for example in (14), the predominant computation is a sparse matrix-vector product  $\Sigma(\sigma)V_i$ , where the sparse matrix  $\Sigma(\sigma)$ , which is related to  $Q$ , represents the transition probabilities between states, and the dense vector  $V_i$  represents the value function for player  $i$ . The product represents the expected continuation value next period, given the current state and transition probabilities. Sparse matrix-vector multiplication (SpMV) operations are common computational kernels in these algorithms, optimized for leveraging the non-zero structure of the matrix to minimize computational load and memory usage.

Similarly, in the Newton-Kantorovich iteration given by (18), the predominant computation involves solving a sparse linear system with a positive but (in general) non-symmetric matrix  $\Sigma(\sigma)$ . Since the conjugate-gradient (CG) method is only applicable to symmetric positive-definite systems, it is not a viable option in this case. Two alternative approaches are QR decomposition and GMRES. QR decomposition is a direct method that robustly computes the solution in a finite number of operations, but calculating the decomposition can be computationally expensive and memory-intensive for large systems. GMRES, on the other hand, is an iterative method that computes successive approximations to the solution. Since its convergence rate depends on factors like the matrix condition number and eigenvalue distribution, the relative speed of QR vs GMRES is system-dependent.

By leveraging the sparsity of the matrices involved in these computations, significant



computational savings can be achieved, enabling efficient solution of large-scale dynamic programming problems. Next, we consider the computational aspects of implementing the log-likelihood function.

### 5.3. Computing $P(\Delta, \theta) = \exp(\Delta Q(\theta))$

The main computational challenge in evaluating the log likelihood function for estimation is the evaluation of the matrix exponential  $P(\Delta, \theta) = \exp(\Delta Q(\theta))$ , where  $Q(\theta)$  is the intensity matrix of the continuous-time Markov chain implied by the structural model. In large-scale models with relatively smaller sample sizes, typically the dataset will contain only observations on a subset of states. When we only need specific rows of the transition matrix  $P(\Delta, \theta)$ , sparse matrix methods can be effectively utilized in the matrix exponential algorithm for computing  $P(\Delta, \theta)v = \exp(\Delta Q(\theta))v$  described in Section 4.1. Since the intensity matrix  $Q(\theta)$  is typically sparse, with non-zero entries only for transitions between states that are directly connected in a single instant, we can use the methods discussed in Section 5.1 to store and update the  $Q(\theta)$  matrix as we search over  $\theta$ . Thus, we can more efficiently compute the transition probabilities compared to dense matrix exponential methods, given the inherent sparsity of our model. Notice that the main computational operation in Algorithms 1 and 2 are again sparse matrix-vector multiplications (SpMV). Therefore, by leveraging sparse matrix methods, the matrix exponential algorithm can be applied to large-scale continuous-time Markov chain models.

### 5.4. Opportunities for Parallelization

In our computational approach, there are several key opportunities for implementing vectorization and parallelization:

1. The value function updates can be parallelized across states, given that the update for each state  $k$  is independent of the others. This form of parallelization is particularly useful in models with a large number of states, allowing for simultaneous updates that drastically reduce computational time. Utilizing multi-threading or distributed computing frameworks can facilitate this parallel processing, ensuring that each processor or core is assigned a subset of the total states for which to compute value updates.
2. SpMV-level parallelization: Sparse matrix-vector multiplications appear in operations such as value iteration and computing matrix exponentials. Libraries such as Intel’s Math Kernel Library (MKL) and NVIDIA’s CUDA offer highly optimized routines for SpMV that exploit the hardware capabilities of CPUs and GPUs, respectively. These

libraries implement advanced parallelization techniques that can significantly speed up the computations by handling multiple operations concurrently across the sparse matrix structure.

3. Optimization-level parallelization: The optimization routine used for estimating model parameters  $\theta$  can be effectively parallelized across diverse starting points or regions of the parameter space. This is especially useful in global optimization algorithms that perform multiple local searches to escape local minima and find a more global solution. By parallelizing these searches, the algorithm can explore various parts of the parameter space simultaneously, enhancing the efficiency and robustness of the optimization process.

The trade-offs between these approaches depend on the specific problem size, hardware resources, and the relative computational costs of different components of the algorithm. State-level parallelization may be more effective for problems with a large number of states, while SpMV-level parallelization may be more effective if a highly optimized library routine is available in your software package or programming language. Optimization-level parallelization can be complementary to the other approaches and may be particularly useful for global optimization routines.

## 6. Monte Carlo Experiments

The model for our Monte Carlo experiments is a simple dynamic entry-exit model among  $N$  firms. The firms operate under stochastically varying market demand conditions. For simplicity, firms are assumed to be symmetric and myopic. Therefore, the state space is defined by the number of active firms and the current state of market demand. The market demand is discretized into several states, allowing us to capture fluctuations in market conditions that can influence firm behavior.

The state space of the model is constructed by considering all possible combinations of firm states and demand levels. For  $N$  firms, this leads to a configuration space where each firm can independently choose to enter (switch from inactive to active) or exit (switch from active to inactive) the market. The number of possible configurations of firm states is  $2^N$ , reflecting all combinations of entry and exit across the firms.

Given  $D$  discrete demand states, the total number of states in our model is given by  $K = 2^N \times D$ . Each state thus corresponds to a unique combination of firm activities and a particular level of demand, allowing us to represent the complete system as a matrix of transitions between these states.

The transition rates between states are governed by a set of parameters and the firms' decisions, which are influenced by current market conditions and the firms' entry status.

The decision process for each firm is modeled as a logit probability where the probability of being active in the market is a function of the state of the model.

Let the current state of the model be  $k$  and let  $l = l(i, j = 1, k)$  be the state resulting from the entry of firm  $i$ . Then the probability of entering for firm  $i$  is given by: These probabilities are defined as:

$$p_{k,l} = \frac{1}{1 + \exp(-(\theta_{\text{EC}} + \theta_{\text{RN}} \times n_{\text{active}} + \theta_{\text{D}} \times \text{demand}))}$$

where  $\theta_{\text{EC}}$  is the entry cost,  $\theta_{\text{RN}}$  represents competitive effect from the number of currently active firms ( $n_{\text{active}}$ ), and  $\theta_{\text{D}}$  represents the profitability of the market level of demand (demand).

The model’s dynamics are implemented through a continuous-time Markov process, where each possible state transition is associated with a rate derived from the firms’ decision probabilities. The transition matrix,  $Q$ , is constructed as a sparse matrix due to the majority of transitions being zero (i.e., many firm configurations do not directly transition between each other). The sparsity of this matrix is leveraged computationally to enhance the efficiency of operations such as matrix exponentiation, which is fundamental for calculating transition probabilities over time.

This matrix exponentiation,  $\exp(\Delta Q)$ , where  $\Delta$  represents a time step, is computed using a uniformization technique that is particularly suited for sparse matrices. This approach not only facilitates the calculation of the matrix exponential but also allows for the simultaneous computation of derivatives with respect to the parameters  $(\theta_{\text{EC}}, \theta_{\text{RN}}, \theta_{\text{D}}, \lambda, \gamma)$ , which are essential for gradient-based optimization methods used in parameter estimation.

We first choose parameters for the model and simulate data. For estimation, we use the observed data to fit the model parameters via Maximum Likelihood supported by the derivatives of the likelihood function with respect to the model parameters.

For our Monte Carlo experiment, we specify  $N = 5$  players,  $D = 5$  demand states, with 1,000 observations observed at discrete intervals of length  $\Delta = 1.0$ . We carried out 101 Monte Carlo replications and report the mean, median, standard deviation, root mean squared error, mean bias, and median bias of the estimates. We used L-BFGS-B to optimize the log likelihood function for discrete time data given in (19). We make an initial pass over each simulated dataset in order to pre-compute the transition counts  $d_{kl}$ .

In Tables 1 and 2, we present the estimation results for the case where we use the analytical Jacobian of the log likelihood function to estimate the model parameters. Specifically, Table 1 displays the results for small samples of 1,000 observations while the results in Table 2 are based on large samples of 4,000 observations. The Jacobian is given in (22) and the derivative of the matrix exponential with respect to the structural parameters are computed

following Algorithm 2. Computational times, in seconds, are reported along with the total number of log likelihood function evaluations. We can see that although there is some small finite sample bias in Table 1, it disappears in the large sample results in Table 2. By comparing the RMSE values, we observe approximately the expected parametric rate of convergence.

By comparison, we report results using numerical Jacobian estimates in Table 3. For brevity, we only report the large sample results. We can see that, although in this simple model, the numerical Jacobian estimates appear to be accurate, they are much more expensive to calculate. The computational times are approximately 4 times longer on average when calculating numerical derivatives and they require about 6 times as many function evaluations on average. We expect that in more complex, dynamic models, the numerical Jacobian estimates may be less accurate, providing another advantage to calculating analytical derivatives using Algorithm 2.

Parameter	True Value	Mean	Median	S.D.	RMSE	Mean Bias	Median Bias
$\theta_{EC}$	-0.500	-0.472	-0.466	0.119	0.122	0.028	0.034
$\theta_{RN}$	-0.050	-0.066	-0.062	0.049	0.051	-0.016	-0.012
$\theta_D$	0.100	0.100	0.098	0.028	0.028	-0.000	-0.002
$\lambda$	1.000	0.995	0.997	0.036	0.037	-0.005	-0.003
$\gamma$	0.300	0.301	0.302	0.017	0.017	0.001	0.002
Time (s)		77.234	78.620	11.222			
Func. Eval.		18.446	19.000	1.375			

TABLE 1. Monte Carlo Results: Analytical Jacobian, 1,000 Obs.

Parameter	True Value	Mean	Median	S.D.	RMSE	Mean Bias	Median Bias
$\theta_{EC}$	-0.500	-0.504	-0.501	0.055	0.055	-0.004	-0.001
$\theta_{RN}$	-0.050	-0.050	-0.053	0.023	0.023	-0.000	-0.003
$\theta_D$	0.100	0.101	0.101	0.014	0.014	0.001	0.001
$\lambda$	1.000	1.000	1.000	0.018	0.018	-0.000	-0.000
$\gamma$	0.300	0.300	0.300	0.009	0.009	0.000	0.000
Time (s)		80.907	80.161	10.036			
Func. Eval.		18.891	19.000	1.266			

TABLE 2. Monte Carlo Results: Analytical Jacobian, 4,000 Obs.

## 7. Conclusion

This paper has presented computational strategies that enhance the efficiency and efficacy of solving and estimating continuous-time dynamic discrete choice models. By integrating

Parameter	True Value	Mean	Median	S.D.	RMSE	Mean Bias	Median Bias
$\theta_{EC}$	-0.500	-0.504	-0.501	0.055	0.055	-0.004	-0.001
$\theta_{RN}$	-0.050	-0.050	-0.053	0.023	0.023	-0.000	-0.003
$\theta_D$	0.100	0.101	0.101	0.014	0.014	0.001	0.001
$\lambda$	1.000	1.000	1.000	0.018	0.018	-0.000	-0.000
$\gamma$	0.300	0.300	0.300	0.009	0.009	0.000	0.000
Time (s)		314.706	312.590	39.642			
Func. Eval.		113.287	114.000	7.518			

TABLE 3. Monte Carlo Results: Numerical Jacobian, 4,000 Obs.

uniformization techniques and sparse matrix operations, we have streamlined the computation of value functions and matrix exponentials in our setting and improved memory use. These enhancements are crucial for handling large-scale dynamic models effectively.

Our Monte Carlo experiments illustrate the practical utility of these computational techniques in a dynamic entry-exit model under varying market demand conditions. Specifically, the use of analytical derivatives of the log likelihood function, compared to numerical derivatives, significantly reduces computational demands. This advantage is particularly notable in scenarios where only discrete-time data are available, which significantly complicates the maximization of the log likelihood function relative to the case of continuous time data.

Moreover, we have established the contractivity and rate of convergence of value iteration, policy evaluation, and relative value iteration. These findings are pertinent not only to single-agent models but also to games with fixed beliefs, such as in the case of two-step estimation with pre-estimated equilibrium choice probabilities. The computational techniques described here are well-suited for modern computational architectures and can be enhanced by parallelization, thereby increasing the speed and scalability of simulations. This advancement opens new avenues in the field of empirical industrial organization, facilitating the exploration of more complex models and deepening our understanding of strategic interactions and policy impacts.

## A. Proofs

### A.1. Proof of Theorem 1

Let  $V_i, V'_i \in \mathbb{R}^K$  denote two arbitrary value functions for player  $i$ . Then, from (4) we have

$$\begin{aligned}
& \|T_{i,\varsigma_i} V_i - T_{i,\varsigma_i} V'_i\|_\infty \\
&= \sup_k \left| T_{i,\varsigma_i} V_{i,k} - T_{i,\varsigma_i} V'_{i,k} \right| \\
&= \sup_k \left| \frac{u_{ik} + \sum_{l \neq k} q_{0kl} V_{il} + \sum_{m \neq i} \lambda_m \sum_j \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_i \mathbf{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \}}{\rho_i + \sum_{l \neq k} q_{0kl} + \sum_m \lambda_m} \right. \\
&\quad \left. - \frac{u_{ik} + \sum_{l \neq k} q_{0kl} V'_{il} + \sum_{m \neq i} \lambda_m \sum_j \varsigma_{imjk} V'_{i,l(m,j,k)} + \lambda_i \mathbf{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V'_{i,l(i,j,k)} \}}{\rho_i + \sum_{l \neq k} q_{0kl} + \sum_m \lambda_m} \right| \\
&= \sup_k \left| \frac{1}{\rho_i + \sum_{l \neq k} q_{kl} + \sum_m \lambda_m} \left( \sum_{l \neq k} q_{kl} (V_{il} - V'_{il}) + \sum_{m \neq i} \lambda_m \sum_j \varsigma_{imjk} (V_{i,l(m,j,k)} - V'_{i,l(m,j,k)}) \right) \right. \\
&\quad \left. + \lambda_i \left[ \mathbf{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \} - \mathbf{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V'_{i,l(i,j,k)} \} \right] \right| \\
&\leq \sup_k \frac{1}{\rho_i + \sum_{l \neq k} q_{kl} + \sum_m \lambda_m} \left( \sum_{l \neq k} q_{kl} |V_{il} - V'_{il}| + \sum_{m \neq i} \lambda_m \sum_j \varsigma_{imjk} |V_{i,l(m,j,k)} - V'_{i,l(m,j,k)}| \right. \\
&\quad \left. + \lambda_i \max_j |V_{i,l(i,j,k)} - V'_{i,l(i,j,k)}| \right) \\
&\leq \sup_k \frac{\sum_{l \neq k} q_{kl} + \sum_m \lambda_m}{\rho_i + \sum_{l \neq k} q_{kl} + \sum_m \lambda_m} |V_{ik} - V'_{ik}| \\
&\leq \frac{\bar{\eta}}{\rho_i + \underline{\eta}} \sup_k |V_{ik} - V'_{ik}| \\
&= \underline{\beta}_i \|V_i - V'_i\|_\infty
\end{aligned}$$

The first equality follows by definition of supremum norm, the second from (4), and the third from collecting a common denominator, noting that  $u_{ik}$  cancels out, and then collecting terms. The first inequality follows from the triangle inequality and, for the final term, using the fact that  $|\sup_k f(k) - \sup_k g(k)| \leq \sup_k |f(k) - g(k)|$  for two functions  $f$  and  $g$ . We

drop the expectation because the random variables  $\varepsilon_{ijk}$  cancel out after subtracting:

$$\begin{aligned}
& \lambda_i \left[ \mathbb{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \right\} - \mathbb{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ijk} + V'_{i,l(i,j,k)} \right\} \right] \\
& \leq \lambda_i \left[ \mathbb{E} \max_j \left\{ \left( \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \right) - \left( \psi_{ijk} + \varepsilon_{ijk} + V'_{i,l(i,j,k)} \right) \right\} \right] \\
& \leq \lambda_i \max_j \left\{ V_{i,l(i,j,k)} - V'_{i,l(i,j,k)} \right\} \\
& \leq \lambda_i \sup_k \left\{ V_{i,k} - V'_{i,k} \right\}
\end{aligned}$$

The last line here follows from  $\max_j |V_{i,l(i,j,k)} - V'_{i,l(i,j,k)}| \leq \sup_k |V_{ik} - V'_{ik}|$ , since the choice  $j$  simply determines the continuation state and the same choice is applied to both value functions. This leads us to the second inequality in the main derivation after noting that the beliefs are each probabilities in  $[0, 1]$ . The final inequality follows from  $\sup_k |f(k)g(k)| \leq \sup_k |f(k)| \sup_k g(k)$  and  $\sup_k |f_1(k)/f_2(k)| \leq \sup_k |f_1(k)| / \inf_k f_2(k)$  and by definition of  $\bar{\eta}$ . The last equality follows from the definition of  $\underline{\beta}_i$  in (11).

### A.2. Proof of Lemma 1

First, note that the aggregate  $Q$  matrix has diagonal elements of the form

$$q_{kk} = - \sum_{m=1,\dots,N} \lambda_m \sum_{j>0} \sigma_{m,j,l(m,j,k)} - \sum_{l \neq k} q_{0kl}.$$

This follows from the first part of Assumption 6 and (1). Since the  $\sigma_{m,j,l(m,j,k)}$  are probabilities that sum to one across choices, but here choice  $j = 0$  is excluded, from Assumption 7 it follows that  $\sum_m \lambda_m > \sum_m \lambda_m \sum_{j>0} \sigma_{m,j,l(m,j,k)}$ . Adding  $\max_k \sum_{l \neq k} q_{0kl}$  to both sides and comparing to the absolute value of  $q_{kk}$  implies the result.

### A.3. Proof of Theorem 2

Recall the Bellman equation from (4). Multiplying both sides by the sum of rates in the denominator yields

$$\left[ \rho_i + \sum_{l \neq k} q_{kl} + \sum_m \lambda_m \right] V_{ik} = u_{ik} + \sum_{l \neq k} q_{kl} V_{il} + \sum_{m \neq i} \lambda_m \sum_j s_{imjk} V_{i,l(m,j,k)} + \lambda_i \mathbb{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} \right\}.$$

Collecting common terms yields a representation in terms of the instantaneous increment to the value function:

$$\rho_i V_{ik} = u_{ik} + \sum_{l \neq k} q_{kl} (V_{il} - V_{ik}) + \sum_{m \neq i} \lambda_m \sum_j s_{imjk} (V_{i,l(m,j,k)} - V_{ik}) + \lambda_i \mathbb{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} - V_{ik} \right\}.$$

First, recall that since the rows of  $Q_0$  sum to zero, we have  $q_{kk} = -\sum_{l \neq k} q_{kl}$ . Therefore,

$$\sum_{l \neq k} q_{kl} (V_{il} - V_{ik}) = \sum_{l \neq k} q_{kl} V_{il} - \sum_{l \neq k} q_{kl} V_{ik} = \sum_{l \neq k} q_{kl} V_{il} + q_{kk} V_{ik} = \sum_{l=1}^K [Q_0]_{kl} V_{il}$$

which is the inner product of the  $k$ -th row of  $Q_0$  with the vector  $V_i$ .

Next, consider the term for rival firm  $m$ . Recall that the choice probabilities sum to one across choices  $j$  and that choice  $j = 0$  is a continuation choice such that  $l(m, 0, k) = k$ :

$$\begin{aligned} \lambda_m \sum_j \varsigma_{imjk} (V_{i,l(m,j,k)} - V_{ik}) &= \lambda_m \sum_{j>0} \varsigma_{imjk} V_{i,l(m,j,k)} + \lambda_m \varsigma_{i,m,0,k} V_{ik} - \lambda_m V_{ik} \\ &= \lambda_m \sum_{j>0} \varsigma_{imjk} V_{i,l(m,j,k)} - \lambda_m \sum_{j>0} \varsigma_{imjk} V_{ik} \\ &= \sum_{l=1}^K [Q_m]_{kl} V_{il}, \end{aligned}$$

which, in equilibrium, is equal to row  $k$  of the matrix  $Q_m$  multiplied by the vector  $V_i$ .

Next, we show that we can write the expectation of the maximum term in terms of choice probabilities, following Arcidiacono, Bayer, Blevins, and Ellickson (2016) and Aguirregabiria and Mira (2002, 2007). Using the definition of  $C_i(\sigma)$  from the statement of the theorem, note that we can write the final term, related to the agent's own optimization, as

$$\lambda_i \mathbb{E} \max_j \{ \psi_{ijk} + \varepsilon_{ijk} + V_{i,l(i,j,k)} - V_{ik} \} = \lambda_i [C_i(\sigma)]_k + \sum_{l=1}^K [Q_i]_{kl} V_{il}.$$

Note that the second term is the product of the  $k$ -th row of  $Q_i$  with the vector  $V_i$ .

Combining these results, we note that we can write the vectorized form of the value function as

$$\rho_i V_i = u_i + Q_0 V_i + \sum_{m \neq i} Q_m V_i + \lambda_i C_i(\sigma) + Q_i V_i.$$

Noting that  $Q = Q_0 + \sum_{m=1}^M Q_m$ , we can write the vectorized form of the value function more simply as

$$(28) \quad \rho_i V_i = u_i + \lambda_i C_i(\sigma) + Q V_i.$$

Note that this is not a contraction. We can, however, now restate the equation so that it is a vectorized version of (4), by subtracting the diagonal of  $Q$ . Let  $\tilde{Q} = Q - \text{diag}(\text{diag } Q)$  denote  $Q$  with the diagonal entries removed. Let  $D = \text{diag}\left(\frac{1}{\rho_i - q_{11}}, \dots, \frac{1}{\rho_i - q_{KK}}\right)$  denote the



diagonal matrix whose elements are equal to the reciprocal of the discount factor,  $\rho_i$ , plus the sum of the exit rates from each state,  $-q_{KK}$ . Then we have

$$V_i = D \left[ u_i + \lambda_i C_i(\sigma) + \tilde{Q} V_i \right].$$

Recall that uniformization involves adjusting the arrival rates by adding self-transitions as necessary in each state  $k$ . The uniformization of  $Q$  is given by  $\Sigma = I + \frac{1}{\eta} Q$ , where  $\eta$  is a positive scalar such that  $\eta \geq \max_k |Q_{kk}|$ . As established in Lemma 1,  $\bar{\eta}$  defined in (16) is a valid uniform rate. Therefore, we may write  $Q = \bar{\eta}(\Sigma - I)$ . To emphasize the dependence of  $\Sigma$  on the policies, through the  $Q_i$  matrices, we write it below as  $\Sigma(\sigma)$ .

Finally, to derive the stated uniform representation of  $V_i$ , note that from (28) we can write

$$\rho_i V_i = u_i + \lambda_i C_i(\sigma) + \bar{\eta}(\Sigma(\sigma) - I) V_i$$

and therefore,

$$\begin{aligned} V_i &= \frac{1}{\rho_i + \bar{\eta}} [u_i + \lambda_i C_i(\sigma)] + \frac{\bar{\eta}}{\rho_i + \bar{\eta}} \Sigma(\sigma) V_i \\ &\equiv U_i(\sigma) + \bar{\beta}_i \Sigma(\sigma) V_i \end{aligned}$$

where we define  $U_i$  and  $\bar{\beta}_i$  as in the proposition.

## B. Uniformization of Value Iteration in Doraszelski and Judd (2012)

Value iteration in Doraszelski and Judd (2012), starting from (12):

The value function is

$$(29) \quad V^i(\omega) = \frac{1}{\rho + \varphi(X^i(\omega), X^{-i}(\omega), \omega)} \left[ \pi(X^i(\omega), X^{-i}(\omega), \omega) \right. \\ \left. + \varphi(X^i(\omega), X^{-i}(\omega), \omega) \mathbb{E}_{\omega'} \left\{ \Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \omega') + V^i(\omega') \mid \omega, X^i(\omega), X^{-1}(\omega) \right\} \right],$$

where  $\pi$  is the flow payoff,  $\varphi(X^i(\omega), X^{-i}(\omega), \omega)$  is the rate at which the next state change occurs,  $\Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \omega')$  is the instantaneous payoff that accrues when the state changes, and  $\mathbb{E}_{\omega'}$  is the expectation with respect to the induced transition distribution. We note that

$$\frac{1}{\rho + \varphi(X^i(\omega), X^{-i}(\omega), \omega)} \varphi(X^i(\omega), X^{-i}(\omega), \omega) < 1,$$

so (29) is contractive holding fixed  $X^{-i}(\omega)$ .

First, we will rearrange (29). We multiply both sides by the denominator of the leading fraction  $\frac{1}{\rho + \varphi(X^i(\omega), X^{-i}(\omega), \omega)}$  and collect terms:

$$\begin{aligned} \rho V^i(\omega) &= \pi(X^i(\omega), X^{-i}(\omega), \omega) \\ &+ \varphi(X^i(\omega), X^{-i}(\omega), \omega) \mathbb{E}_{\omega'} \left\{ \Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \omega') + V^i(\omega') \mid \omega, X^i(\omega), X^{-1}(\omega) \right\} \\ &\quad - \varphi(X^i(\omega), X^{-i}(\omega), \omega) V^i(\omega). \end{aligned}$$

We now apply the uniformization method. Choose  $\nu > \max_{\omega} \sup_{x^i} \varphi(x^i, X^{-i}(\omega), \omega)$ . We will introduce self-transitions so that in state  $\omega$ , at rate  $\varphi(X^i(\omega), X^{-i}(\omega), \omega)$ , we transition according to  $\Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \omega')$  as before, but at rate  $\nu - \varphi(X^i(\omega), X^{-i}(\omega), \omega)$  we will have a self-transition that does not change the state. This does not change the value function, but yields a uniform representation with a constant rate  $\nu$  of transitions across all states  $\omega$ . To see this, we add  $\nu V^i(\omega)$  to both sides of the previous equation:

$$\begin{aligned} (\rho + \nu) V^i(\omega) &= \pi(X^i(\omega), X^{-i}(\omega), \omega) \\ &+ \varphi(X^i(\omega), X^{-i}(\omega), \omega) \mathbb{E}_{\omega'} \left\{ \Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \omega') + V^i(\omega') \mid \omega, X^i(\omega), X^{-1}(\omega) \right\} \\ &\quad + \left( \nu - \varphi(X^i(\omega), X^{-i}(\omega), \omega) \right) V^i(\omega). \end{aligned}$$

Finally, solving for  $V^i$  gives

$$\begin{aligned} V^i(\omega) &= \frac{1}{\rho + \nu} \pi(X^i(\omega), X^{-i}(\omega), \omega) \\ &+ \frac{\varphi(X^i(\omega), X^{-i}(\omega), \omega)}{\rho + \nu} \mathbb{E}_{\omega'} \left\{ \Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \omega') + V^i(\omega') \mid \omega, X^i(\omega), X^{-1}(\omega) \right\} \\ &\quad + \frac{\nu - \varphi(X^i(\omega), X^{-i}(\omega), \omega)}{\rho + \nu} V^i(\omega). \end{aligned}$$

Alternatively, we can think of this as an expectation over the state  $\tilde{\omega}'$  where

$$\tilde{\omega}' = \begin{cases} \omega' & \text{with probability } \frac{\varphi(X^i(\omega), X^{-i}(\omega), \omega)}{\nu}, \\ \omega & \text{with probability } 1 - \frac{\varphi(X^i(\omega), X^{-i}(\omega), \omega)}{\nu}, \end{cases}$$

so that

$$\begin{aligned} V^i(\omega) &= \frac{1}{\rho + \nu} \pi(X^i(\omega), X^{-i}(\omega), \omega) \\ &\quad + \frac{\nu}{\rho + \nu} \mathbb{E}_{\tilde{\omega}'} \left\{ \Phi^i(X^i(\omega), X^{-i}(\omega), \omega, \tilde{\omega}') + V^i(\omega') \mid \omega, X^i(\omega), X^{-1}(\omega) \right\}, \end{aligned}$$

and where we define self-transitions to be costless:

$$\Phi^i \left( X^i(\omega), X^{-i}(\omega), \omega, \omega \right) = 0.$$

## References

- Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica* 70, 1519–1543.
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53.
- Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: a survey. *Journal of Econometrics* 156, 38–67.
- 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.
- Berry, S. T. (1992). Estimation of a model of entry in the airline industry. *Econometrica* 60, 889–917.
- Blevins, J. R. (2022). Identification and estimation of continuous time dynamic discrete choice games. Working paper, Ohio State University.
- Blevins, J. R. and M. Kim (2024). Nested pseudo likelihood estimation of continuous-time dynamic discrete games. *Journal of Econometrics* 238, 105576.
- Bray, R. L. (2019). Strong convergence and dynamic economic models. *Quantitative Economics* 10, 43–65.
- Bresnahan, T. F. and P. C. Reiss (1991). Empirical models of discrete games. *Journal of Econometrics* 48, 57–81.
- Chung, K. L. (1967). *Markov Chains with Stationary Transition Probabilities*. Berlin: Springer.
- Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics* 3, 53–93.

- Fox, B. L. and P. W. Glynn (1988). Computing Poisson probabilities. *Communications of the ACM* 31, 440–445.
- Gilbert, J. R., C. Moler, and R. Schreiber (1992). Sparse matrices in MATLAB: Design and implementation. *SIAM Journal on Matrix Analysis and Applications* 13(1), 333–356.
- Goldberg, D. (1991). What every computer scientist should know about floating-point arithmetic. *ACM Computing Surveys* 23(1), 5–48.
- Grassmann, W. (1977a). Transient solutions in Markovian queues. *European Journal of Operational Research* 1, 396–402.
- Grassmann, W. K. (1977b). Transient solutions in Markovian queueing systems. *Computers & Operations Research* 4, 47–53.
- 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.
- Jensen, A. (1953). Markoff chains as an aid in the study of Markoff processes. *Scandinavian Actuarial Journal* 1953, 87–91.
- Karlin, S. and H. M. Taylor (1975). *A First Course in Stochastic Processes* (Second ed.). San Diego, CA: Academic Press.
- Moler, C. and C. V. Loan (1978). Nineteen dubious ways to compute the exponential of a matrix. *SIAM Review* 20, 801–836.
- 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.
- Puterman, M. L. (2005). *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Second ed.). Wiley-Interscience.
- Reibman, A. and K. Trivedi (1988). Numerical transient analysis of Markov models. *Computers & Operations Research* 15, 19–36.

- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55, 999–1013.
- Sherlock, C. (2022). Direct statistical inference for finite Markov jump processes via the matrix exponential. *Computational Statistics* 36, 2863–2887.
- Tijms, H. C. (2003). *A First Course in Stochastic Models* (Second ed.). Wiley.
- van Dijk, N. M., S. P. J. van Brummelen, and R. J. Boucherie (2018). Uniformization: Basics, extensions and applications. *Performance Evaluation* 118, 8–32.