

Estimation of Dynamic Discrete Choice Models in Continuous Time*

Peter Arcidiacono¹, Patrick Bayer¹,
Jason R. Blevins², and Paul B. Ellickson³

¹*Duke University and NBER*

²*Ohio State University*

³*Simon Graduate School of Business, University of Rochester*

October 2, 2012

Abstract

This paper provides a method for estimating large-scale dynamic discrete choice models within a continuous time framework. An advantage of our model is that state changes occur sequentially, rather than simultaneously, avoiding a substantial curse of dimensionality that arises in multi-agent settings. Eliminating this computational bottleneck is the key to providing a seamless link between estimating the model and performing post-estimation counterfactuals. While recently developed two-step estimation techniques have made it possible to estimate large-scale problems, solving for equilibria remains computationally challenging. By modeling decisions in continuous time, we are able to take advantage of the recent advances in estimation while preserving a tight link between estimation and policy experiments. We address the most commonly encountered situation in empirical work in which only discrete-time data are available and the actual sequence of events that occur between two points in time is unobserved. We apply our techniques to examine the effects of Walmart's entry into the retail grocery industry, showing that even the threat of entry by Walmart has a substantial effect on market structure.

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

JEL Classification: C13, C35, L11, L13.

*We thank the attendees of the 2009 Cowles Foundation Conference on Structural Microeconomics, the 2010 cemmap Conference on Matching and Sorting, and the 2012 NBER/NSF/CEME Conference on the Econometrics of Dynamic Games as well as seminar participants at Chicago (Booth), Duke, Ohio State, Penn State, Johns Hopkins, Toronto, Western Ontario, Washington University (Olin), Wisconsin, and Yale for useful comments. Timothy Schwuchow provided excellent research assistance.

1 Introduction

Empirical models of single-agent dynamic discrete choice (DDC) problems have a rich history in structural applied microeconometrics, starting with the pioneering work of [Gotz and McCall \(1980\)](#), [Miller \(1984\)](#), [Pakes \(1986\)](#), [Rust \(1987\)](#), and [Wolpin \(1984\)](#). Because dynamic decision problems are naturally high-dimensional, the empirical DDC literature has been accompanied from the outset by a parallel methodological literature aimed at reducing the computational burden of both estimating and computing these models.

The computational challenges raised by the dimensionality of these problems are even greater in the context of strategic games, where the simultaneous actions of competing players introduces a further curse of dimensionality in computing expectations over rivals' actions. In particular, in order to solve for optimal policies, one must calculate players' expectations over all combinations of actions of their rivals. The cost of computing these expectations grows exponentially in the number of players, making it difficult or impossible to compute the equilibrium in many economic environments.¹

An alternative way to model strategic decision-making that does not suffer from this curse of dimensionality is to characterize players' moves as stochastic sequential decisions. This can be achieved by assuming a stochastically determined order of moves in discrete time ([Doraszelski and Judd, 2007](#)) or by setting the game in continuous time, with independent, competing stochastic processes controlling when players are able to act.² With either approach, because only a single player moves at any point in time, there is no need to compute expectations over all of the possible combinations of decisions that the other players might make simultaneously.

In this paper, we develop a particular characterization of a dynamic game in continuous time that not only eliminates the curse of dimensionality associated with simultaneous move games, but also links naturally with the existing DDC literature. The key feature of our approach is that players face a standard discrete choice problem when stochastically given the opportunity to make a decision. That is, when players have the chance to act, they observe the realization of a random component of the payoff for each available action and select the best available option.

By letting the time and source of each event be determined by the outcome of a collection

¹These limitations have led some to suggest alternatives to the Markov perfect equilibrium concept in which firms condition on long run averages (regarding rivals' states) instead of current information ([Weintraub, Benkard, and Van Roy, 2008](#)).

²[Doraszelski and Judd \(2012\)](#) exploit the structure of continuous time to break the curse of dimensionality associated with the calculation of expectations over rival actions. Players in their model make simultaneous, continuous decisions that control the hazard rate of state changes (e.g., choose an investment hazard which results stochastically in a discrete productivity gain). Because state changes occur only one agent at a time, the dimension of expectations over rival actions grows linearly in the number of players, rather than exponentially, resulting in computation times that are orders of magnitude faster than those of discrete time.

of independent, competing Poisson processes, for which the rate parameters may vary, one does not need to make further assumptions about the order of moves or the number of events that occur in any given period, as one must do in a sequential-move discrete-time model with fixed period lengths. The critical features of our model are that both the identity of the next agent to move and the number of moves that occur in any fixed observation interval are random, with no further restrictions on either the order of moves or the number of moves that can occur in any fixed time period.

Our formulation of the dynamic game in continuous time naturally inherits many features of the standard discrete choice framework and, as a result, many of the insights and tools commonly used in discrete time settings are directly applicable. Importantly, it is possible to extend the two-step CCP (conditional choice probability) methods that have been developed for estimating dynamic games in discrete time to the continuous time setting, resulting in a further reduction in the computational burden and making it possible to estimate and compute especially rich, high-dimensional dynamic games.³

CCP estimation applied to our formulation of a dynamic game in continuous time has several important advantages that carry over from the discrete time literature. Most directly, CCP estimation eliminates the need to compute the full solution of the model for estimation. In most empirical studies, the equilibrium will only need to be computed a handful of times to perform the counterfactual analyses conducted in the paper. In addition, it is straightforward to account for unobserved heterogeneity with our framework by extending the methods of [Arcidiacono and Miller \(2011\)](#). We demonstrate both of these advantages in our empirical application, applying the methods to a high dimensional problem while incorporating unobserved heterogeneity, an important feature of the institutional setting.

But CCP estimation has advantages in continuous time beyond those studied in the discrete time literature. Namely, it is easier to satisfy the finite dependence property of [Arcidiacono and Miller \(2011\)](#) and [Altuğ and Miller \(1998\)](#), whereby only a handful of conditional choice probabilities are needed to express the future utility term. This occurs because the inversion theorem of [Hotz and Miller \(1993\)](#) yields a mapping between differences in value functions and conditional choice probabilities in the continuous time setting, as opposed to differences in *conditional* value functions and conditional choice probabili-

³A recent series of papers ([Aguirregabiria and Mira, 2007](#), [Bajari, Benkard, and Levin, 2007](#), [Pesendorfer and Schmidt-Dengler, 2007](#), [Pakes, Ostrovsky, and Berry, 2007](#)) have shown how to extend two-step estimation techniques, originally developed by [Hotz and Miller \(1993\)](#) and [Hotz, Miller, Sanders, and Smith \(1994\)](#) in the context of single-agent dynamics, to more complex multi-agent settings. The computation of these models remains formidable, despite a growing number of methods for solving for equilibria ([Pakes and McGuire, 1994, 2001](#), [Doraszelski and Satterthwaite, 2010](#)). Two-step estimation of dynamic discrete games was originally proposed by [Rust \(1994\)](#). Rust recommended substituting non-parametric estimates of rivals' reaction functions into each player's dynamic optimization problem, turning a complex equilibrium solution into a collection of simpler games against nature.

ties in the discrete time setting. As we illustrate in [Section 3](#), working with value functions rather than conditional value functions also has the added benefit of limiting the need to estimate conditional choice probabilities for rare events, a prevalent feature in many empirical settings.

We demonstrate the advantages of our formulation of the dynamic game in continuous time with an empirical application that analyzes the entry and exit, expansion and contraction of grocery chain stores in urban markets throughout the United States from 1994–2006. Specifically, we model the decisions of whether to operate grocery stores in a market and at what scale (i.e., number of stores) for Walmart and up to seven competing chains as well as the single-store entry decisions of several dozen potential fringe entrants. Each geographic market is characterized by observed features—most importantly, the level and growth rate of population—as well as unobserved heterogeneity that affects the relative profitability of Walmart, chain, and fringe stores in that market.

This characterization of the problem results in a dynamic game that has a rich error structure (due to the unobserved heterogeneity) and an enormous number of states.⁴ We estimate the model using CCP methods and solve counterfactually for the equilibrium under several scenarios designed to measure how Walmart’s entry into the retail grocery industry (which took place in a number of markets over the study period) affects the profitability and decision-making of rival chain and fringe firms.

The estimates imply that there is considerable heterogeneity in the relative profitability of chain and fringe stores in markets throughout the country. Interestingly, the impact of Walmart is also heterogeneous across markets, with the greatest impact typically felt by the previously dominant players in the region. In chain-dominated metropolitan areas, for example, the entry of Walmart typically leads to the exit of one or more chains firms altogether and the contraction of stores by those firms that remain in the market.

Our estimates show that the mere threat of entry by Walmart has a substantial effect on industry structure. For example, many of the markets we study are forecasted to have more fringe firms when Walmart is allowed to enter, but hasn’t entered yet, than when Walmart is prohibited from entering at all. This occurs because chain firms would expand more rapidly if they were certain Walmart was not going to enter. This chain expansion results in fringe firms contracting.⁵

⁴The number of distinct states is over 287 billion.

⁵While our paper is the first to estimate the effects of Walmart on both chain and on single stores, in part due to complications associated with the large state space, others have examined the effect of Walmart on supermarkets and in the market for discount stores. [Ellickson and Grieco \(2013\)](#) examine the impact of Walmart on the supermarket industry using descriptive methods from the treatment effects literature, while [Beresteanu, Ellickson, and Misra \(2007\)](#) develop a dynamic structural model of retail competition. Walmart’s previous experience in the discount industry has been analyzed by [Jia \(2008\)](#), [Holmes \(2011\)](#), and [Ellickson, Houghton, and Timmins \(2010\)](#).

The paper is structured as follows. [Section 2](#) introduces our model in a simple single-agent context in order to build intuition. [Section 3](#) develops an alternative CCP representation of the value function which will facilitate two-step estimation of the model. [Section 4](#) extends the model to the multi-agent setting. Concrete and canonical examples are provided in both the single- and multi-agent cases. [Section 5](#) develops our estimators, including both full-solution and two-step approaches, and discusses issues associated with time aggregation. [Section 6](#) introduces and describes the results of our empirical analysis of the market structure of grocery store chains in geographically separate U.S. markets. [Section 7](#) concludes.

2 Single-Agent Dynamic Discrete Choice Models

In this section, we introduce a dynamic discrete choice model of single-agent decision-making in continuous time. The single-agent problem provides a simple setting in which to describe the main features of our continuous time framework. We show how these extend directly to multi-agent settings in the following section. We begin this section by laying out the notation and structure of the model in a general context. We then introduce an example—the classic bus engine (capital) replacement model of [Rust \(1987\)](#)—to illustrate how to apply our model in a familiar setting.

Consider a dynamic single-agent decision problem in which time is continuous, indexed by $t \in [0, \infty)$. The state of the model at any time t can be summarized by an element k of some finite state space $\mathcal{X} = \{1, \dots, K\}$. Two competing Poisson processes drive the dynamics of the model. First, a finite-state Markov jump process on \mathcal{X} with a $K \times K$ intensity matrix Q_0 governs moves by nature—exogenous state changes that aren't a result of actions by the agent. The elements of Q_0 are the rates at which particular state transitions occur. Second, a Poisson arrival process with rate parameter λ governs when the agent can move. When a move opportunity arrives, the agent chooses an action j from among J alternatives in a discrete choice set $\mathcal{A} = \{0, \dots, J - 1\}$, conditional on the current state k .

Before describing the agent's problem, we review some properties of finite Markov jump processes, which are the basic building blocks of our model and can be used to characterize both exogenous and endogenous state changes on \mathcal{X} in our model. A finite Markov jump process on \mathcal{X} is a stochastic process X_t indexed by $t \in [0, \infty)$. At any time t , the process remains at X_t for a random time interval τ (the holding time) before transitioning to some new state $X_{t+\tau}$. A sample path of such a process is a piecewise-constant, right-continuous function of time. Jumps occur according to a Poisson process and the holding times between jumps are exponentially distributed.

A finite-state Markov jump processes can be characterized by an *intensity matrix*, which

contains the rate parameters for each possible state transition:

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

For $l \neq k$

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

is the hazard rate for transitions from state k to state l and

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

is the overall rate at which the process leaves state k (and hence, q_{kk} is negative). Transitions out of state k follow an exponential distribution with rate parameter $-q_{kk}$ and, conditional on leaving state k , the process transitions to $l \neq k$ with probability $q_{kl} / \sum_{l' \neq k} q_{kl'}$. For additional details about Markov jump processes see, for example, [Karlin and Taylor \(1975, Section 4.8\)](#).

Returning to the agent's problem, we assume that the agent is forward-looking and discounts future payoffs at rate ρ . Exogenous state changes that the agent cannot control occur according to a Markov jump process with intensity matrix Q_0 . While in state k , the agent receives flow utility u_k . At rate λ the agent makes a decision, choosing an action $j \in \mathcal{A}$ and receiving an instantaneous payoff $\psi_{jk} + \varepsilon_j$, where ψ_{jk} is the mean payoff associated with making choice j in state k and $\varepsilon_j \in \mathbb{R}$ is an instantaneous choice-specific payoff shock.⁶ Let σ_{jk} denote the probability that the agent optimally chooses choice j in state k . The agent's choice may result in a deterministic state change.⁷ Let $l(j, k)$ denote the state that results upon making choice j in state k .

We can now derive the instantaneous Bellman equation, a recursive expression for the value function V_k which gives the present discounted value of all future payoffs obtained from starting in some state k and behaving optimally in future periods. For small time increments h , under the Poisson assumption, the probability of an event with rate λ occurring is λh . Given the discount rate ρ , the discount factor for such increments is $1/(1 + \rho h)$. Thus, for

⁶Although the choice-specific shocks ε_j evolve over time, we omit the t subscript for notational simplicity. For convenience, we also assume the distribution of ε_j is independent of the state, to avoid conditioning on k throughout, but the joint distribution of the vector $\varepsilon = (\varepsilon_0, \dots, \varepsilon_{J-1})^\top$ can, in general, depend on k without additional difficulties.

⁷For expositional simplicity, we focus on the case of deterministic state changes, rather than allowing for a stochastic state transition following each action.

small time increments h the present discounted value of being in state k is

$$V_k = \frac{1}{1 + \rho h} \left[u_k h + \sum_{l \neq k} q_{kl} h V_l + \lambda h \mathbb{E} \max_j \{ \psi_{jk} + \varepsilon_j + V_{l(j,k)} \} + \left(1 - \lambda h - \sum_{l \neq k} q_{kl} h \right) V_k + o(h) \right].$$

Rearranging and letting $h \rightarrow 0$, we obtain the following recursive expression for V_k :

$$V_k = \frac{u_k + \sum_{l \neq k} q_{kl} V_l + \lambda \mathbb{E} \max_j \{ \psi_{jk} + \varepsilon_j + V_{l(j,k)} \}}{\rho + \lambda + \sum_{l \neq k} q_{kl}}. \quad (1)$$

The denominator contains the sum of the discount factor and the rates of all possible state changes. The numerator is composed of the flow payoff for being in state k , the rate-weighted values associated with exogenous state changes, and the expected current and future value obtained when a move arrival occurs in state k . The expectation is with respect to the joint distribution of $\varepsilon = (\varepsilon_0, \dots, \varepsilon_{J-1})^\top$. Alternatively, and perhaps more intuitively, rearranging once again shows that the instantaneous discounted increment to the value function V_k is

$$\rho V_k = u_k + \sum_{l \neq k} q_{kl} (V_l - V_k) + \lambda \mathbb{E} \max_j \{ \psi_{jk} + \varepsilon_j + V_{l(j,k)} - V_k \}. \quad (2)$$

A policy rule is a function $\delta : \mathcal{X} \times \mathbb{R}^J \rightarrow \mathcal{A}$ which assigns to each state k and vector ε an action from \mathcal{A} . The optimal policy rule satisfies the following inequality condition:

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

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

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

Note that the move arrival rate, λ , and the choice probabilities of the agent, σ_{jk} , also imply a Markov jump process on \mathcal{X} with intensity matrix Q_1 , where Q_1 is a function of both λ and σ_{jk} for all j and k . In particular, the hazard rate of action j in state k is simply $\lambda \sigma_{jk}$, the product of the move arrival rate and the choice probability. The choice probability σ_{jk} is thus the proportion of moves in state k , which occur at rate λ , that result in action j . Summing the individual intensity matrices yields the aggregate intensity

matrix $Q = Q_0 + Q_1$ of the compound process, which fully characterizes the state transition dynamics of the model. This simple and intuitive structure is especially important in extending the model to include multiple agents, and in estimation with discrete time data. We discuss both of these extensions in subsequent sections.

2.1 Example: A Single-Agent Renewal Model

Our first example is a simple single-agent renewal model, based on the bus engine replacement problem analyzed by Rust (1987). The state variable represents the accumulated mileage of a bus engine. Let q_{k1} and q_{k2} denote the rates at which one- and two-unit mileage increments occur, respectively. With each move arrival, the agent faces a binary choice: replace the engine ($j = 1$) or continue ($j = 0$). If the agent replaces the engine, the mileage is reset to state $k = 1$ and the agent pays a replacement cost c . The agent faces a cost minimization problem where the flow cost incurred in mileage state k is represented by u_k . The value function for mileage state k is

$$V_k = \frac{u_k + q_{k1}V_{k+1} + q_{k2}V_{k+2} + \lambda \mathbb{E} \max \{V_k + \varepsilon_0, V_0 + c + \varepsilon_1\}}{\rho + q_{k1} + q_{k2} + \lambda}, \quad (3)$$

where, in our general notation from before, the instantaneous payoffs are

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

We will return to this example in the following section, where we discuss a useful CCP representation of the value function.

3 CCP Representation

The primary difference between our framework and traditional discrete time dynamic discrete choice models is that, rather than having state changes and choices made simultaneously at pre-determined intervals, only one event occurs at any given instant (almost surely), with random time intervals between moves. Given that the unobserved preferences evolve stochastically as in the discrete time literature, we are able to show that some of the insights of Hotz and Miller (1993), Altuğ and Miller (1998), and Arcidiacono and Miller (2011) on expressing value functions in terms of conditional choice probabilities (CCPs) also apply here. In fact, as we show below, it is actually much easier to express the value functions in terms of CCPs in the continuous time setting than in discrete time.

We first derive two results that allow us to link value functions across states. The first is

essentially the continuous time analog of Proposition 1 of Hotz and Miller (1993). Namely, using the conditional choice probabilities, we can derive relationships between the value functions associated with any two states as long as both states are feasible from the initial state, should the agent have the right to move.

Let $\sigma_k = (\sigma_{0k}, \dots, \sigma_{J-1,k})^\top$ denote the vector of CCPs in state k . We make the following standard assumptions on the discount rate and the choice-specific shocks (Rust, 1994).

Assumption 1. $\rho > 0$.

Assumption 2. The choice-specific shocks ε are iid over time and across choices with a known joint distribution which is absolutely continuous with respect to Lebesgue measure, has finite first moments, and support equal to \mathbb{R}^J .

Proposition 1. *If Assumptions 1 and 2 hold, then there exists a function $\Gamma^1(j, j', \sigma_k)$ such that for all $j, j' \in \mathcal{A}$,*

$$V_{l(j,k)} = V_{l(j',k)} + \psi_{j'k} - \psi_{jk} + \Gamma^1(j, j', \sigma_k). \quad (4)$$

Proof. See Appendix A. ■

Thus, in our continuous time model the value function can be separated from the choice in each state. This is similar to the result of Proposition 1 of Hotz and Miller (1993), which applied to *conditional* value functions: the value of making a particular choice conditional on behaving optimally in the future. Because only one event occurs at any given instant, we are able to develop relationships between the value functions themselves.

The second result establishes a similar CCP representation for the final term in the Bellman equation.

Proposition 2. *If Assumptions 1 and 2 hold, then there exists a function $\Gamma^2(j', \sigma_k)$ such that for all $j' \in \mathcal{A}$,*

$$\mathbb{E} \max_j \left\{ \psi_{jk} + \varepsilon_j + V_{l(j,k)} \right\} = V_{l(j',k)} + \psi_{j'k} + \Gamma^2(j', \sigma_k). \quad (5)$$

Proof. See Appendix A. ■

The intuition for Proposition 2 is that we can express the left hand side of (5) relative to $V_{l(j',k)} + \psi_{j'k}$ for an action j' of our choosing, implying that the terms inside the expectation will consist of differences in value functions and instantaneous payoffs. These differences, as established by Proposition 1, can be expressed as functions of conditional choice probabilities.

For a concrete example of these two propositions, consider the case where the ε 's follow the type I extreme value distribution. In this case, closed form expressions exist for both Γ^2 and Γ^1 :

$$\begin{aligned}\Gamma^1(j, j', \sigma_k) &= \ln(\sigma_{jk}) - \ln(\sigma_{j'k}) \\ \Gamma^2(j', \sigma_k) &= -\ln(\sigma_{j'k}) + \gamma\end{aligned}$$

where γ is Euler's constant.

Importantly, [Proposition 1](#) allows us to link value functions across many states. For example, suppose that action 0 is a continuation action which does not change the state, $l(0, k) = k$, and has no instantaneous payoff or cost, $\psi_{0k} = 0$. If in state k the agent is able to move to k' by taking action j' , and is further able to move from k' to k'' by taking action j'' , then it is possible to express $V_{k''}$ as a function of V_k by substituting in the relevant relationships:

$$\begin{aligned}V_k &= V_{k'} + \psi_{j',k} + \Gamma^1(0, j', \sigma_k) \\ &= V_{k''} + \psi_{j'',k'} + \psi_{j',k} + \Gamma^1(0, j'', \sigma_{k'}) + \Gamma^1(0, j', \sigma_k).\end{aligned}$$

By successively linking value functions to other value functions, in many cases it is straightforward to find a chain such that the remaining value functions on the right hand side of (2) can be expressed in terms of V_k and conditional choice probabilities. Then, collecting all terms involving V_k yields an expression for V_k in terms of the flow payoff of state k and the conditional choice probabilities. Since the latter can often be flexibly estimated directly from the data and the former is an economic primitive, it is no longer necessary to solve a dynamic programming problem to obtain the value functions. This is formalized in the following result.

Definition. A state k^* is *attainable from state k* if there exists a sequence of actions from k that result in state k^* .

Proposition 3. *Suppose that 1 and 2 hold and that for a given state k , $j = 0$ is a continuation action with $l(0, k) = k$ and for all states $l \neq k$ with $q_{kl} > 0$ there exists a state k^* that is attainable from both k and l . Then, there exists a function $\Gamma_k(\psi, Q_0, \lambda, \sigma)$ such that*

$$\rho V_k = u_k + \Gamma_k(\psi, Q_0, \lambda, \sigma). \tag{6}$$

Proof. See [Appendix A](#). ■

The function Γ_k for each state may depend on the model primitives ψ , Q_0 , and λ as well as the CCPs, σ . By restating the problem in this way, when the conditional choice

probabilities are available, no fixed point problem needs to be solved in order to obtain the value functions. This can often lead to large computational gains. We now illustrate some examples of how to apply these propositions.

3.1 Example: Inventories

We begin by considering an inventory example. The amount of inventory a firm has is given by $k \in \{0, \dots, K\}$. With each move arrival, the firm may either increase its inventory by one unit or do nothing. The instantaneous cost of increasing inventory is c and the flow cost for holding k units is u_k . At rate q , the firm makes a sale and inventory falls by m units, $m \in \{1, \dots, M\}$ or, if current inventory is less than m , by the current inventory level. The probability that demand is m is $\pi(m)$. Given demand of m , revenues received are $p_{km} = p \min\{k, m\}$. The value function for inventory level k is

$$V_k = \frac{u_k + q \sum_{m=1}^M (V_{l(k,m)} + p_{km})\pi(m) + \lambda \mathbb{E} \max \{V_k + \varepsilon_0, V_{k+1} + c + \varepsilon_1\}}{\rho + q + \lambda}.$$

where $l(k, m) = \max\{k - m, 0\}$. Applying [Proposition 2](#), we can eliminate the value functions in the third term of the numerator:

$$V_k = \frac{u_k + q \sum_{m=1}^M (V_{l(k,m)} + p_{km})\pi(m) + \lambda \Gamma^2(0, \sigma_k)}{\rho + q}.$$

Since k is attainable from $l(k, m)$, we can apply [Proposition 3](#). In particular, we can repeatedly use [Proposition 1](#) to express $V_{l(k,m)}$ in terms of V_k :

$$\begin{aligned} V_{k-l(k,m)} &= V_{k-l(k,m)+1} + c + \Gamma^1(0, 1, \sigma_{k-l(k,m)}) \\ &= V_k + \min\{k, m\}c + \sum_{k'=k-l(k,m)}^{k-1} \Gamma^1(0, 1, \sigma_{k'}). \end{aligned}$$

Hence, the function Γ_k from [Proposition 3](#) is

$$\Gamma_k(\psi, Q_0, \lambda, \sigma) = \sum_{m=1}^M \left(\min\{k, m\}c + \sum_{k'=k-l(k,m)}^{k-1} \Gamma^1(0, 1, \sigma_{k'}) \right) \pi(m) + \lambda \Gamma^2(0, \sigma_k),$$

which yields the expression for V_k in (6). Hence, even though state K cannot be reached by a single decision when $m > 1$, the two value functions can still be linked and expressed as a function of conditional choice probabilities.

3.2 Example: A Single-Agent Renewal Model

Recall the bus engine replacement example of [Section 2.1](#), where the value function was characterized by (3). Applying [Proposition 2](#) eliminates the third term in the numerator:

$$V_k = \frac{u_k + q_{k1}V_{k+1} + q_{k2}V_{k+2} + \lambda\Gamma^2(0, \sigma_k)}{\rho + q_{k1} + q_{k2}}.$$

Although there is no *direct* link between the value function at k and the value functions at $k + 1$ and $k + 2$, it is possible to link the two value functions through the replacement decision. In particular, V_k and V_{k+1} can be expressed as follows:

$$\begin{aligned} V_k &= V_0 + c + \Gamma^1(0, 1, \sigma_k), \\ V_{k+1} &= V_0 + c + \Gamma^1(0, 1, \sigma_{k+1}). \end{aligned}$$

This implies that we can express V_{k+1} in terms of V_k :

$$V_{k+1} = V_k + \Gamma^1(0, 1, \sigma_{k+1}) - \Gamma^1(0, 1, \sigma_k).$$

Using a similar expression for V_{k+2} , we obtain the function Γ_k from [Proposition 3](#):

$$\begin{aligned} \Gamma_k(\psi, Q_0, \lambda, \sigma) &= q_{k1}\Gamma^1(0, 1, \sigma_{k+1}) + q_{k2}\Gamma^1(0, 1, \sigma_{k+2}) \\ &\quad - (q_{k1} + q_{k2})\Gamma^1(0, 1, \sigma_k) + \lambda\Gamma^2(0, \sigma_k). \end{aligned}$$

This example illustrates one of the benefits of continuous time over discrete time when using conditional choice probabilities. Namely, as illustrated by [Arcidiacono and Miller \(2011\)](#), forming renewal problems using CCPs required both expressing the future utility term relative to a particular choice *and* differencing the conditional valuation functions used in estimation. In this example, the future utility terms for both replacing and continuing would be expressed relative to the value of replacing. Hence accurate estimates of the conditional probability of replacing at very low mileages would be needed, but these are low probability events whose estimates will likely depend heavily on the smoothing parameters or functional forms used to mitigate the associated small sample problems. In the present continuous time framework, differencing is not required in order to form the value function in terms of conditional choice probabilities and we only need the replacement probabilities at states k , $k + 1$, and $k + 2$.

4 Dynamic Discrete Games

The potential advantages of modeling decisions using a continuous time framework are particularly applicable to games, where the state space is often enormous. Working in continuous time highlights aspects of strategic interaction that are muted by discrete time (e.g., first-mover advantage) and mitigates unnatural implications that can arise from simultaneity (e.g., ex post regret). In fact, a number of recent papers in the empirical games literature (e.g., Einav, 2010, Schmidt-Dengler, 2006) have adopted a sequential structure for decision-making to accommodate the underlying economic theory associated with their games.

Extending the single-agent model of Section 2 to the case of dynamic discrete games with many players is simply a matter of modifying the intensity matrix governing the state to incorporate players' beliefs regarding the future actions of their rivals. We begin this section by describing the structure of the model, followed by properties of equilibrium strategies and beliefs. We then show how to apply the CCP representation results of Section 3 in the context of dynamic games.

Suppose there are \mathcal{N} players indexed by $i = 1, \dots, \mathcal{N}$. As before, the state space \mathcal{X} is finite with K elements. This is without loss of generality, since each of these elements may be regarded as indices of elements in a higher-dimensional, but finite, space of firm-market-specific state vectors. Player i 's choice set in state k is \mathcal{A}_{ik} . For simplicity, we consider the case where each player has J actions in all states: $\mathcal{A}_{ik} = \{0, \dots, J - 1\}$. We index the remaining model primitives by i , including the flow payoffs in state k , u_{ik} , instantaneous payoffs, ψ_{ijk} , and choice probabilities, σ_{ijk} . Let $l(i, j, k)$ denote the continuation state that arises after player i makes choice j in state k . We assume that players share a common discount rate ρ .

Although it is still sufficient to have only a single jump process on \mathcal{X} , with some intensity matrix Q_0 , to capture moves by nature, there are now \mathcal{N} independent, competing Poisson processes with rate λ generating move arrivals for each of the \mathcal{N} players.⁸ The next event in the model is determined by the earliest arrival of one of these $\mathcal{N} + 1$ processes.

Let ς_i denote player i 's beliefs regarding the actions of rival players, given by a collection of $(\mathcal{N} - 1) \times J \times K$ probabilities ς_{imjk} for each rival player $m \neq i$, state k , and choice j . Applying Bellman's principal of optimality (Bellman, 1957), the value function for an active player i in state k can be defined recursively as⁹

⁸For simplicity, we assume the move arrival rates are equal for each firm. Furthermore, although we do not consider this extension, one can introduce correlation among move arrivals by specifying a multinomial distribution over the possible combinations of simultaneously moving players, which would imply a rate for each of these possible outcomes. However, this generalization comes at the expense of both additional parameters and computational cost.

⁹If a player is not active in state k , then the probability of inaction ($j = 0$) for that player is set to one.

$$V_{ik}(\varsigma_i) = \frac{u_{ik} + \sum_{l \neq k} q_{kl} V_{il}(\varsigma_i) + \sum_{m \neq i} \lambda \sum_j \varsigma_{imjk} V_{i,l(m,j,k)}(\varsigma_i) + \lambda \mathbb{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ij} + V_{i,l(i,j,k)}(\varsigma_i) \right\}}{\rho + \sum_{l \neq k} q_{kl} + \mathcal{N}\lambda}. \quad (7)$$

Following Maskin and Tirole (2001), we focus on Markov perfect equilibria in pure strategies, as is standard in the literature. A *Markov strategy* for player i is a mapping which assigns an action from \mathcal{A}_{ik} to each state $(k, \varepsilon_i) \in \mathcal{X} \times \mathbb{R}^J$. Focusing on Markov strategies eliminates the need to condition on the full history of play.

Given beliefs for each player, $\{\varsigma_i\}$, and a collection of model primitives, a Markov strategy for player i is a *best response* if¹⁰

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

Then, given the distribution of choice-specific shocks, each Markov strategy δ_i implies response probabilities for each choice in each state:

$$\sigma_{ijk} = \Pr[\delta_i(k, \varepsilon_i) = j \mid k]. \quad (8)$$

Definition. A collection of Markov strategies $\{\delta_1, \dots, \delta_{\mathcal{N}}\}$ and beliefs $\{\varsigma_1, \dots, \varsigma_{\mathcal{N}}\}$ is a *Markov perfect equilibrium* if for all i :

1. $\delta_i(k, \varepsilon_i)$ is a best response given beliefs ς_i , for all k and almost every ε_i ;
2. for all players $m \neq i$, the beliefs ς_{mi} are consistent with the best response probabilities implied by δ_i , for each j and k .

Following Milgrom and Weber (1985) and Aguirregabiria and Mira (2007), we can characterize Markov perfect equilibria in probability space, rather than in terms of pure Markov strategies, as a collection of equilibrium best response probabilities $\{\sigma_i\}$ where each probability in σ_i is a best response given beliefs σ_{-i} . Likewise, any such collection of probabilities can be extended to a Markov perfect equilibrium.

In particular, equilibrium conditional choice probabilities are fixed points to the best response probability mapping, which defines a continuous function from $[0, 1]^{\mathcal{N} \times J \times K}$ onto itself. Existence of an equilibrium then follows from Brouwer's Theorem, as established by the following proposition. The proof is reserved for the appendix.

Proposition 4. *A Markov perfect equilibrium exists.*

Proof. See Appendix A. ■

¹⁰In the event of a tie, we assume the action with the smallest index is assigned. Because the distribution of ε_i is continuous under Assumption 2, such ties occur with probability zero.

4.1 CCP Representation

The propositions in [Section 3](#) apply to games as well. Hence, it is possible to eliminate the value functions in the fourth term of the numerator of (7) using [Proposition 2](#):

$$V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{kl} V_{il} + \sum_{m \neq i} \lambda \sum_j \sigma_{mjk} V_{i,l(m,j,k)} + \lambda \Gamma_i^2(0, \sigma_{ik})}{\rho + \sum_{l \neq k} q_{kl} + \mathcal{N}\lambda}. \quad (9)$$

Eliminating the other value functions, however, is problematic as the player may only have control over a portion of the state space. For example, when firms have different numbers of stores, a given firm is only able to choose its own stores, not the stores of its competitors. There are at least two cases where the remaining value functions can be eliminated. The first is the case where there is a terminal choice, such as permanently exiting a market. Since no further choices are made, the value function for the terminal choice does not include other value functions. A concrete example is provided below. The other case is where an action can be taken to reset the system for all players. For example, consider a game that involves technology adoption. By achieving a particular technology level, previous technologies may become obsolete, effectively renewing the states of the other players.

4.2 Example: Multi-Store Entry and Exit

Note that in either the terminal or reset case, there only has to be an attainable scenario where the agent can execute the terminal or reset action. For example, consider a game amongst retailers where firms compete by opening and closing stores. Given a move arrival, a firm can build a store, $j = 1$, do nothing, $j = 0$, or, if the agent has at least one store, close a store, $j = -1$. Once a firm has no stores, it makes no further choices. Let c denote the scrap value of closing a store.

Suppose that firm i has k_i stores and the economy-wide state is $k = (k_1, \dots, k_{\mathcal{N}})$. Let $l^*(i, k, k'_i)$ denote the state that is equal to the initial state k , but where firm i has k'_i stores instead of k_i . Applying [Proposition 1](#) and normalizing the value of zero stores to zero, we can express V_{ik} as:

$$V_{ik} = \sum_{k'_i=1}^{k_i} \Gamma^1(0, -1, \sigma_{i,l^*(i,k,k'_i)}) + k_i c. \quad (10)$$

Since (10) holds for all k , we can use the value of fully exiting to link value functions for any pair of states. Namely, linking the value functions on the right hand side of (9) to

V_{ik} and solving for V_{ik} yields:

$$\begin{aligned} \rho V_{ik} = & u_{ik} + \lambda_i \Gamma_i^2(0, \sigma_{ik}) \\ & + \sum_{m \neq i} \sigma_{m,-1,k} \left(\sum_{k'_i=1}^{k_i} \left[\Gamma_i^1 \left(0, -1, \sigma_{i,l^*(i,l(m,-1,k),k'_i)} \right) - \Gamma_i^1 \left(0, -1, \sigma_{i,l^*(i,k,k'_i)} \right) \right] \right) \\ & + \sum_{m \neq i} \sigma_{m,1,k} \left(\sum_{k'_i=1}^{k_i} \left[\Gamma_i^1 \left(0, -1, \sigma_{i,l^*(i,l(m,1,k),k'_i)} \right) - \Gamma_i^1 \left(0, -1, \sigma_{i,l^*(i,k,k'_i)} \right) \right] \right). \end{aligned}$$

Once again, no fixed point calculation is required to express the full value function, a simplification that is especially powerful in the context of high-dimensional discrete games.

5 Estimation

We now turn to estimation. Methods that solve for the value function directly and use it to obtain the implied choice probabilities for estimation are referred to as full-solution methods. The nested fixed point (NFXP) algorithm of Rust (1987), which uses value function iteration inside of an optimization routine that maximizes the likelihood, is the classic example of a full-solution method. Su and Judd (2012) provide an alternative MPEC (mathematical program with equilibrium constraints) approach which solves the constrained optimization problem directly, bypassing the repeated solution of the dynamic programming problem.

CCP-based estimation methods, on the other hand, are two-step methods pioneered by Hotz and Miller (1993) and Hotz et al. (1994) and later extended by Aguirregabiria and Mira (2002, 2007), Bajari et al. (2007), Pesendorfer and Schmidt-Dengler (2007), Pakes et al. (2007), and Arcidiacono and Miller (2011). The CCPs are estimated in a first step and used to approximate the value function in a closed-form inversion or simulation step. The approximate value function is then used in the likelihood function or the GMM criterion function to estimate the structural parameters.

Full-solution methods have the advantage that the full structure of the model is imposed in estimation. However, these methods can become quite computationally expensive for complex models with many players or a large state space. Many candidate parameter vectors must be evaluated during estimation and, if the value function is costly to compute, even if solving the model once might be feasible, doing so many times may not be. In the presence of multiple equilibria, they also require researchers to make an assumption on the equilibrium selection mechanism and solve for all the equilibria (cf. Bajari, Hong, and Ryan, 2007).¹¹ In addition to allowing the value function to be computed very quickly,

¹¹Since we have not established that there is a unique equilibrium, when discussing the full-solution estimator, we implicitly assume that the equilibrium selection mechanism assigns probability one to the

CCP methods provide an attractive solution to the issue of multiplicity. When the data are generated by a single equilibrium, the resulting likelihood conditions on the equilibrium that is played in the data, bypassing the need to consider other equilibria.

Our model has the advantage of being estimable via either approach. As in Doraszelski and Judd (2012), the use of continuous time breaks one primary curse of dimensionality in that only a single player moves at any particular instant. An attractive and novel feature of our framework is that it is also easily estimable using standard CCP methods. This greatly reduces the computational costs of estimation relative to full-solution methods. Having estimated a large problem with CCP methods, it is then straightforward to use the model for post-estimation exercises, since the computational burden of computing the equilibrium a few times for these purposes is not as great as nesting several such solutions into an estimation routine. In this way, our framework preserves a tight link between the estimated model and that used for post-estimation analysis, something which has proven infeasible for many empirical applications that have been modeled in discrete time.

In the rest of this section we describe the estimation algorithms. We begin with full-solution and two-step methods when continuous time data is observed. Since data is often reported only at discrete intervals, we next show how our methods can be applied to discrete time data. We then extend the methods to incorporate permanent unobserved heterogeneity. Finally, we discuss identification.

5.1 Full-Solution Estimation

We maintain the convention that choice $j = 0$ for each agent is a continuation choice which does not change the state. Then, in state k , the probability of the next state change occurring during an interval of length τ is

$$1 - \exp \left[-\tau \left(\sum_{l \neq k} q_{kl} + \sum_i \lambda \sum_{j \neq 0} \sigma_{ijk} \right) \right]. \quad (11)$$

This is the cdf of the exponential distribution with rate parameter equal to the sum of the exogenous state transition rates and the hazards of the non-continuation actions for each player, where the equilibrium choice probabilities σ_{ijk} are, implicitly, functions of the rates q and λ and the parameters θ .

Differentiating with respect to τ yields the density for the time of the next state change,

equilibrium obtained by a numerical fixed point routine. The computational advantages of continuous time, however, make it easier to explore more complex equilibrium selection mechanisms with non-degenerate weightings.

which is the exponential pdf with the same rate parameter as before:

$$\left(\sum_{l \neq k} q_{kl} + \sum_i \lambda \sum_{j \neq 0} \sigma_{ijk} \right) \exp \left[-\tau \left(\sum_{l \neq k} q_{kl} + \sum_i \lambda \sum_{j \neq 0} \sigma_{ijk} \right) \right]. \quad (12)$$

Conditional on a state change occurring in state k , the probability that the change is due to agent i taking action j is

$$\frac{\lambda \sigma_{ijk}}{\sum_{l \neq k} q_{kl} + \sum_i \lambda \sum_{j \neq 0} \sigma_{ijk}}. \quad (13)$$

Now, define the function

$$g(\tau, k | q, \lambda, \theta) = \exp \left[-\tau \left(\sum_{l \neq k} q_{kl} + \sum_i \lambda \sum_{j \neq 0} \sigma_{ijk} \right) \right], \quad (14)$$

which is the second term from (11). Then, the joint likelihood of the next stage change occurring after an interval of length τ and being the result of player i taking action j is the product of (12) and (13),

$$\lambda \sigma_{ijk} g(\tau, k | q, \lambda, \theta),$$

with the corresponding likelihood of nature moving the state from k to l being

$$q_{kl} g(\tau, k | q, \lambda, \theta).$$

Now consider a sequence of N state changes occurring over a period of length T . Let k_n denote the state prior to the n -th state change, with k_{N+1} denoting the final state. Let t_n denote the time of the n -th event and let τ_n denote the holding time between events, defined as $\tau_n = t_n - t_{n-1}$ for $n \leq N$. For the interval between the last event and the end of the sampling period, we define $\tau_{N+1} = T - t_N$. Let $I_n(l)$ be the indicator for whether the n -th move was a move by nature to state l and, in a slight abuse of notation, let $I_n(i, j)$ be the indicator for whether the n -th move was a move by player i and the choice was j . The maximum likelihood estimates of q , λ , and θ are then the solution to:

$$\left\{ \hat{q}, \hat{\lambda}, \hat{\theta} \right\} = \arg \max_{(q, \lambda, \theta)} \left\{ \sum_{n=1}^N \left[\ln g(\tau_n, k_n | q, \lambda, \theta) + \sum_{l \neq k_n} I_n(l) \ln q_{kl} \right. \right. \\ \left. \left. + \sum_i \sum_{j \neq 0} I_n(i, j) \ln (\lambda \sigma_{ijk}(q, \lambda, \theta)) \right] + \ln g(\tau_{N+1}, k_{N+1} | q, \lambda, \theta) \right\}.$$

Note that embedded in the estimation problem is the solution to a fixed point problem which

is needed in order to obtain the value functions. We have made the dependence of the choice probabilities on the parameters and rates explicit. The last term in the expression is the natural logarithm of one minus the exponential cdf, to account for the fact that another state change was not observed by the end of the sampling period.

5.2 Two-Step Estimation

As discussed in Section 3, it is possible to express differences in continuous time value functions as functions of the conditional choice probabilities. These expressions can sometimes be used in such a way that solving the nested fixed point problem is unnecessary. In this section, we show how two-step methods apply in estimation, linking reduced form hazards to conditional choice probabilities.

5.2.1 Step 1: Estimating the Reduced-Form Hazards

Let $h_{ijk} = \lambda \sigma_{ijk}$ denote the hazard for an active player i choosing action j in state k . In Step 1, one estimates the hazards h_{ijk} nonparametrically. For example, these hazards can be estimated by maximum likelihood by writing the exponential cdf in (11) as a function of h_{ijk} instead of $\lambda \sigma_{ijk}$. Similarly, we can rewrite the function in (14) as

$$\tilde{g}(\tau, k | q, h) = \exp \left[-\tau \left(\sum_{l \neq k} q_{kl} + \sum_i \sum_{j \neq 0} h_{ijk} \right) \right].$$

Then, the maximum likelihood estimates of the hazards q and h are

$$\{\tilde{q}, \tilde{h}\} = \arg \max_{(q, h)} \left\{ \sum_{n=1}^N \left[\ln \tilde{g}(\tau_n, k_n | q, h) + \sum_{l \neq k_n} I_n(l) \ln q_{kl} + \sum_i \sum_{j \neq 0} I_n(i, j) \ln h_{ijk} \right] + \ln \tilde{g}(\tau_{N+1}, k_{N+1} | q, h) \right\}.$$

5.2.2 Step 2: Estimating the Structural Payoff Parameters

In Step 2, we use the estimated reduced-form hazards to estimate the structural payoff parameters θ . Given the estimated hazards and a value of λ , we can estimate the conditional choice probabilities for $j \neq 0$ as \tilde{h}_{ijk}/λ and $1 - \sum_{j \neq 0} \tilde{h}_{ijk}/\lambda$ for $j = 0$. Therefore, when the finite dependence condition holds, we can express the structural conditional choice probabilities as functions of λ , θ , and the estimated hazards and rates so that no fixed-point problem needs to be solved. Let $\check{\sigma}_{ijk}(\tilde{q}, \tilde{h}, \lambda, \theta)$ denote this mapping. Similarly, let $\check{g}(t, k | \tilde{q}, \tilde{h}, \lambda, \theta)$ denote the mapping in (14) with $\check{\sigma}_{ijk}(\tilde{q}, \tilde{h}, \lambda, \theta)$ used in place of σ_{ijk} .

Since we already have estimates of the rates q_{kl} , we focus on estimating λ and θ . The joint likelihood of the next state change occurring after an interval of length τ and being the result of player i taking action j is

$$\lambda \check{\sigma}_{ijk}(\tilde{q}, \tilde{h}, \lambda, \theta) \check{g}(\tau, k | \tilde{q}, \tilde{h}, \lambda, \theta).$$

The second stage estimates are then

$$\{\check{\lambda}, \check{\theta}\} = \arg \max_{(\lambda, \theta)} \left\{ \sum_{n=1}^N \left[\ln \check{g}(\tau_n, k_n | \tilde{q}, \tilde{h}, \lambda, \theta) + \sum_i \sum_{j \neq 0} I_n(i, j) \ln \left(\lambda \check{\sigma}_{ijk}(\tilde{q}, \tilde{h}, \lambda, \theta) \right) \right] + \ln \check{g}(\tau_{N+1}, k_{N+1}, | \tilde{q}, \tilde{h}, \lambda, \theta) \right\}.$$

5.3 Discrete Time Data

Often the exact sequence of events and event times is not observed, rather, the state is only observed at discrete points in time. Here, we consider estimation with discretely-sampled data, focusing in particular on full-solution estimation, but one can easily carry out two-step estimation in an analogous manner.

Let $P_{kl}(\Delta)$ denote the probability that the system has transitioned to state l after a period of length Δ given that it was initially in state k , given the aggregate intensity matrix Q . The corresponding matrix of these probabilities, $P(\Delta) = (P_{kl}(\Delta))$, is the *transition matrix*, which satisfies

$$P(t) = e^{\Delta Q} = \sum_{j=0}^{\infty} \frac{(\Delta Q)^j}{j!}. \quad (15)$$

This quantity is the matrix exponential, the matrix analog of the scalar exponential, which can be computed using one of many known algorithms (cf. Moler and Loan, 1978, Sidje, 1998).¹²

The transition probabilities summarize the relevant information about a pair of observations (t_{n-1}, k_{n-1}) and (t_n, k_n) . That is, $P_{k_{n-1}, k_n}(t_n - t_{n-1})$ is the probability of the process moving from k_{n-1} to k_n after an interval of length $t_n - t_{n-1}$. This includes cases where $k_n = k_{n-1}$, since the transition probabilities account for the case of no jumps at all, as well as all sequences involving any number of jumps to intermediate states before returning to the initial state.

The transition matrix depends on model primitives such as conditional choice probabilities that are themselves functions of q , λ , and θ , which we write as $P(\Delta; q, \lambda, \theta)$. The log

¹²The matrix exponential operator is available, for example, in Matlab, via the `expm` function.

likelihood function for a sample $\{(t_n, k_n)\}_{n=1}^N$ is thus

$$\ln L_N(q, \lambda, \theta) = \sum_{n=1}^N \ln P_{k_{n-1}, k_n}(t_n - t_{n-1}; q, \lambda, \theta).$$

Since the Q matrix can be large, this may seem to introduce a computational curse of dimensionality rivaling that of discrete time models. However, the Q matrix is often very sparse, which substantially reduces the computational burden. Sparse matrix algorithms can be used to compute $P(\Delta)$, which typically require only being able to compute the action of Q on some generic vector v . Since the structure of Q is known, this usually involves very few multiplications relative to the size of the intensity matrix, which is $K \times K$. Furthermore, only at most N rows of $P(\Delta)$ need be calculated to estimate the model, corresponding to the number of observations. Algorithms are available which exploit the sparsity of Q and directly compute the action of $P(\Delta)$ on some vector v , further reducing the computational cost. Since v can be the n -th standard basis vector, one can compute only the necessary rows of $P(\Delta)$.

We now provide some intuition for why discrete-time data will not substantially complicate the problem, which also suggests a natural way to form a simulation-based estimator of $P(\Delta)$. Recall that we are considering stationary models: the hazards do not depend on t . This implies that we can decompose the Markov jump process into two components: a state-independent Poisson process dictating when moves occur, and an embedded Markov chain which dictates where the state moves to.

In our setting, the embedded Markov chain associated with moves by agents is a matrix containing the relevant conditional choice probabilities. We can also rewrite the state transition rates for nature in a similar manner. Namely, let \bar{q} denote the maximum sum of these rates across all states:

$$\bar{q} = \max_j \sum_k q_{jk}.$$

If we consider \bar{q} to be the move arrival rate for nature, then at each move opportunity, the probability of transitioning from state k to l due to a move by nature is

$$q'_{kl} = \begin{cases} q_{kl}/\bar{q} & \text{if } k \neq l, \\ 1 - \sum_{k \neq l} q_{kl}/\bar{q} & \text{if } k = l. \end{cases}$$

Given the conditional choice probabilities, σ , and the transition probabilities for nature, q' , we can construct an embedded Markov chain \mathcal{Z} which characterizes the transition probabilities across all states for each arrival of the alternative Poisson process with rate $\bar{q} + \mathcal{N}\lambda$. The transition matrix associated with moving from any state k to any future state

k' in exactly r steps is simply \mathcal{Z}^r . Let a_n denote a vector of length K , which has a one in position k_n , corresponding to the state at observation n , and zeros elsewhere (i.e., the k_n -th standard basis vector). The maximum likelihood estimates given a dataset of discrete observations at intervals of unit length satisfy

$$\{\hat{q}, \hat{\lambda}, \hat{\theta}\} = \arg \max_{(q, \lambda, \theta)} \sum_{n=1}^N \ln \left[\sum_{r=0}^{\infty} \frac{(\bar{q} + \mathcal{N}\lambda)^r \exp(-(\bar{q} + \mathcal{N}\lambda))}{r!} a_n^\top \mathcal{Z}(q, \theta)^r a_{n+1} \right]. \quad (16)$$

The first term in the innermost summation above is the probability of exactly r state changes occurring during a unit interval, under the Poisson distribution. The second term is the probability of the observed state transition, given that there were exactly r moves.

The expression above also suggests a simulation-based estimator. Namely, use the expression inside the sum for the first $R < \infty$ terms. Given an initial guess of λ and \bar{q} , draw from the event distribution conditional on having more than R events. One could then use importance sampling to weight the number of events to avoid redrawing the simulated paths when changing the parameters.

5.4 Unobserved Heterogeneity

Our methods can also be extended to accommodate permanent unobserved heterogeneity using finite mixture distributions. In particular, suppose that N observations are sampled over the interval $[0, T]$ for each of M markets, where each market is one of Z types. Let $\pi_z(k_{m1})$ denote the population probability that market m is of type z conditional on the initial observation.¹³ We can then integrate with respect to the distribution of the unobserved state, so that the maximum likelihood problem becomes

$$\{\hat{q}, \hat{\lambda}, \hat{\theta}, \hat{\pi}\} = \arg \max_{(q, \lambda, \theta, \pi)} \sum_{m=1}^M \ln \left[\sum_{z=1}^Z \pi_z(k_{m1}) \prod_{n=1}^N \ln P_{k_{m,n-1}, k_{mn}}(\tau_{mn}; q, \lambda, \theta, z) \right], \quad (17)$$

where $P(\Delta; q, \lambda, \theta, z)$ is the transition matrix for type z and $\tau_{mn} = t_{mn} - t_{m,n-1}$ is the interval between observations $n-1$ and n for market m .

Although (17) is written for the full-solution case, the methods outlined in [Arcidiacono and Miller \(2011\)](#) apply. They show that the EM algorithm can be used to recover the conditional choice probabilities as part of the maximization problem or in a first stage. The same methods can be applied here, only now it is the reduced form hazards conditional on both the observed and unobserved states that are being recovered.

¹³By letting $\pi_z(k_{m1})$ depend on k_{m1} , we allow for an initial conditions problem.

5.5 Identification

With continuous-time data, identification and estimation of the intensity matrix for finite-state Markov jump processes is straightforward and well-established (Billingsley, 1961). However, when a continuous-time process is only sampled at discrete points in time, the parameters of the underlying continuous-time model may not be identified. In continuous-time models, this is known as the aliasing problem, which has been studied by many authors in the context of continuous-time multivariate regression models (Sims, 1971, Phillips, 1973, Hansen and Sargent, 1983, Geweke, 1978).¹⁴

In the context of finite-state Markov jump processes, the question is whether there a unique matrix Q that leads to the observed transition matrix $P(\Delta)$ when the process is sampled at intervals of length Δ . Singer and Spilerman (1976) and Geweke, Marshall, and Zarkin (1986) discuss the aliasing problem in this context and Singer and Spilerman (1976) provide several sufficient conditions, any of which guarantee that Q is unique, for example, if the eigenvalues of $P(\Delta)$ are distinct, real, and positive, if $\min_k\{P_{kk}(\Delta)\} > 1/2$, or if $\det P(\Delta) > e^{-\pi}$.

Such conditions are useful, but for structural models, many restrictions on Q arise naturally from the underlying model. For multivariate regression models, Phillips (1973) discusses the role of prior information on the intensity matrix and how it can lead to identification. The structural model underlying the Q matrix in our framework provides exactly the sort of prior restrictions needed to mitigate the aliasing problem. The model restricts Q to a lower-dimensional subspace, since it is sparse and must satisfy both within-row and across-row restrictions. Therefore, even if there are multiple matrix solutions to the equation $P(\Delta) = \exp(\Delta Q)$, it is unlikely that two of them simultaneously satisfy the restrictions of the structural model. For example, consider the restrictions on the Q matrix

¹⁴A related issue is the embeddability problem: could the observed discrete-time transition matrix $P(\Delta)$, associated with a time interval of length Δ , have been generated by the proposed data generating process (some continuous-time Markov structure with intensity matrix Q or some discrete-time chain over fixed time periods of length δ). This is a model specification issue, also arising in both discrete time and continuous time: was the data actually generated by a continuous time Markov jump process? We assume throughout that the model is well-specified and therefore, such an intensity matrix Q exists. Singer and Spilerman (1976) provide several necessary conditions for embeddability involving testable conditions on the determinant and eigenvalues of $P(\Delta)$. This problem was first proposed by Elfving (1937). Kingman (1962) derived the set of embeddable processes with $K = 2$ and Johansen (1974) gave an explicit description of the set for $K = 3$.

implied by the simple renewal model of Section 2.1:

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

The number of parameters to be estimated in this matrix is substantially less than if the intensities were allowed to vary across all the non-diagonal elements. Furthermore, the CCPs depend on the rates and payoff parameters, which introduces shape restrictions on σ_{1k} across states k .

Importantly, the question of identification applies equally to discrete-time models, where there is an analogous problem. Suppose there is a fixed move interval of length δ in the model which may be different from the fixed interval Δ at which observations are sampled. In practice, researchers typically assume (implicitly) that $\delta = \Delta$, where $\Delta = 1$ is normalized to be some specific unit of time (e.g., one quarter).¹⁵ This assumption is convenient, but masks the identification problem, which requires that there exist a unique matrix root P_0 of the discrete-time aggregation equation $P_0^{\Delta/\delta} = P(\Delta)$. In general, however, there may be multiple such matrices (Singer and Spilerman, 1976, p. 49). As in our setting, however, valid solutions P_0 must satisfy the restrictions implied by the model. These issues become trivial under the usual assumption that $\delta = \Delta$.

6 Empirical Application

Our empirical application considers the impact of Walmart's entry into the supermarket industry. Walmart gained national prominence in the discount retail segment, building a large network of discount stores radiating out from their corporate headquarters in Bentonville, Arkansas. In the late 1980s, they began opening supercenters, which combined their original dry goods format with a full-line supermarket. In 1994, the first year for which we have data, Walmart owned 97 outlets. However, by 2006, the last year of our data, they had opened 2225 outlets, many of which were conversions of existing discount stores to the new superstore format.

Quantifying the impact of Walmart's entry into groceries requires understanding a bit about the overall structure of the supermarket industry. The grocery industry has long

¹⁵The continuous-time model is more flexible than the discrete-time model in the sense that one can estimate the rate of move arrivals, rather than fixing it at unity. The continuous-time model which most closely mirrors the usual discrete-time assumptions, which are imposed implicitly, is found by setting $\lambda \equiv 1$.

been comprised of two distinct segments: large regional chains and a competitive “fringe” of local stores (the majority of which are sole proprietorships). Most markets have between three and four chain players (with roughly similar market shares) and a long tail of fringe stores (whose number increases monotonically with the size of the market). While Walmart is clearly a disruptive presence, it is unclear which of these segments is most vulnerable to its entry, and how exactly it reshapes the competitive landscape.

6.1 Model

To accommodate the institutional features of the supermarket industry and allow for heterogeneous competitive effects across formats, there are three types of firms in our model: chain firms (who can operate many stores), Walmart (who can also operate many stores), and fringe firms (who can operate at most one store each). With each move arrival, chain stores can open one new store ($j = 1$), do nothing ($j = 0$), or, conditional on having at least one open store, close a store ($j = -1$). A move arrival for an incumbent fringe firm provides an opportunity for the firm to exit. Similarly, move arrivals provide opportunities for potential entrants to enter. In the context of retail competition, a random move arrival process might reflect the stochastic timing of local development projects (e.g., housing tracts and business parks), delays in the zoning and permitting processes, and the random arrival of retailers in other lines of business that have higher valuations for the properties currently occupied by incumbent grocers. All firms have the same move arrival rate, λ , and q_1 and q_{-1} are the rates of moving up and down in population, respectively. The function $l(i, j, k)$ then gives the state conditional on firm i taking action j in state l .

6.1.1 Value Functions

We now provide the general formulation of the value functions and then describe the relevant state variables. For a particular market, the value function for firm i in state k is then given by:

$$V_{ik} = \frac{u_{ik} + \sum_{j \in \{-1, 1\}} q_j V_{i, l(0, j, k)} + \sum_{m \neq i} \lambda \sum_{j \in \{-1, 1\}} \sigma_{mjk} V_{i, l(m, j, k)} + \lambda \mathbf{E} \max_j \left\{ V_{i, l(i, j, k)} + \psi_{ijk} + \varepsilon_{ijk} \right\}}{\rho + \sum_{j \in \{-1, 1\}} q_j + \sum_{m \neq i} \lambda \sum_{j \in \{-1, 1\}} \sigma_{mjk} + \lambda}, \quad (18)$$

where nature is indexed by $i = 0$ and ψ_{ijk} reflects the costs of opening or closing a store.

Following standard convention in the empirical entry literature, we assume that if a chain or fringe firm closes all of its stores, then the firm cannot enter again later. Hence, if a chain firm exits, it would be replaced by a new potential chain entrant. For chain and fringe firms, this allows us to replace the value functions on the right-hand side of (18) using Propositions 1 and 2. As a result (and exploiting Proposition 3), the value

function on the left-hand side of (18) can be expressed as a function of the flow payoffs, the move arrival parameters, and the probabilities of making particular decisions. Note that this representation eliminates the need to make any assumptions regarding the maximum number of chain firms, Walmart stores, or fringe firms in the market, since no fixed point problem is solved.

6.1.2 Flow Profits for Chain Firms

We specify per-store revenue for chain firms, u_{ik}^c , as a linear function of population, pop , the number of own stores, s_i , and the number of competing stores of each type (chain, Walmart, and fringe), \tilde{s}_i . Revenues also depend on an unobserved (to the econometrician) characteristic of the market, z , which reflects the tastes of those in the market for particular types of products. Flow profits for a chain firm in state k are

$$u_{ik}^c = s_i (\beta_0^c + \beta_1^c pop - \beta_2^c s_i - \beta_3^c \cdot \tilde{s}_i + \beta_4^c z) - e_{ik}^c,$$

where e_{ik}^c is the flow cost of operating a set of stores:

$$e_{ik}^c = \phi_1^c s_i + \phi_2^c s_i^2 + \phi_3^c s_i^3.$$

A cubic cost function allows there to be regions of increasing and then decreasing returns to scale, so that for each state the optimal value of s_i is finite. Expanding and collecting terms yields

$$u_{ik}^c = (\beta_0^c - \phi_1^c - \beta_3^c \cdot \tilde{s}_i + \beta_4^c z) s_i + \beta_1^c s_i pop - (\beta_2^c + \phi_2^c) s_i^2 - \phi_3^c s_i^3.$$

6.1.3 Flow Profits for Fringe Firms

Revenues for fringe stores have a similar linear form to that of chain stores, though with different coefficients and a different flow cost function, e_{ik}^f . Namely, an operating fringe store has profits given by:¹⁶

$$u_{ik}^f = \beta_0^f + \beta_1^f pop - \beta_3^f \cdot \tilde{s}_i + \beta_4^f z - e_{ik}^f,$$

Fringe competitors often depend on the same suppliers. Hence, there may be some economies of scale present at first. However, at some point competitive influences will drive up costs, suggesting a quadratic cost function in the total number of fringe stores, \bar{s} :

$$e_{ik}^f = \phi_1^f + \phi_2^f \bar{s}^2.$$

¹⁶Note that there is no $\beta_2 s_i$ term as fringe stores can only operate at most one store.

6.2 Data

The data for the supermarket industry come from yearly snapshots of the Trade Dimensions Retail Database, capturing the set of players that are active in September of each year, starting in 1994 and ending in 2006. Trade Dimensions continuously collects information on every supermarket (and many other retailers) operating in the United States for use in their Marketing Guidebook and Market Scope publications and as a standalone, syndicated dataset. The definition of a supermarket used by Trade Dimensions is the government and industry standard: a store selling a full line of food products that grosses at least \$2 million in revenue per year. Store level data on location, revenue and employment is linked to the firm level through a firm identity code, which can also be used to identify the location of the nearest distribution facility.

A firm is deemed to be a chain firm in a market if it has at least 20 stores open nationally and its maximum market share (in terms of number of stores) exceeds 20% in at least one year. In our model, there are seven chain players in each MSA who may or may not be active in the market. If a chain has no stores in a particular period and chooses not to build a store, that chain is replaced by a new potential chain entrant. In our model, each MSA has ten potential fringe firm entrants, so the number of fringe firms is the number of incumbent fringe firms plus ten.

One issue with the data is that supermarket chains may move one of their stores to another location. Hence, exits and entries can be *positively* correlated. Since location within an MSA is beyond the scope of the model, we look at net entry and exit within the period. Hence, we assume that if we see both an entry and exit by the same chain, this pair of moves is equivalent to not moving at all.

Demand for supermarkets is a function of population. The data on market population are interpolated from the decennial censuses of the United States and population is discretized into six categories.¹⁷ Each MSA is assigned to one of three population growth categories based on the change in the population of the MSA over the full sample period.

Table 1 gives descriptive statistics for the sample. On average, there are about two and a half chain firms per market, with 3.7 stores per chain firm on average. Markets contain an average of 13 fringe stores. The number of Walmarts is much smaller, averaging one store per market in the sample. On average, there are 0.277, 0.177, and 0.825 stores built per market within a year by chain firms, Walmart, and fringe firms, respectively. The corresponding figures for store closings are 0.224, 0.002, and 0.908, revealing that Walmart virtually never exits in our sample.

Table 2 looks at entry and exit decisions for chain firms and fringe firms one year

¹⁷The discretization was such that differences in log population between adjacent categories was equal.

before, the year of, and the year after initial entry of Walmart. Here, we see that chain and fringe firms both respond negatively to Walmart. The number of new chain stores falls from 0.311 in the period before Walmart enters to 0.189 in the period after—a 40% drop. Similarly, the number of stores closing increases by over 6.5% from a base level of 0.122. The qualitative patterns for fringe firms are the same, though the effects are muted, suggesting that Walmart’s presence is more detrimental to chain firms than fringe firms.

This table highlights an advantage of using a continuous time model. Note that the numbers of entering and exiting chain stores in the year of Walmart’s initial entry are bracketed by the corresponding values the year before and the year after Walmart’s entry. In markets where chain and fringe stores saw little change in their building patterns, this indicates that Walmart entered later in the period. But when Walmart enters early in the period, exit by chain and fringe stores is more likely to occur within the period. A continuous time model can explain both the increase in exit rates during the entry period as well as the even further increase in exit rates in the period after.

6.3 Estimation

We estimate the model in two steps, first estimating the reduced form hazards that capture the rate of change in the number of stores of each format and population over time, and then estimating the structural parameters of the profit functions, taking the reduced form hazards as given. Throughout, we normalize λ , the arrival rate of moves for firms, to one, implying an average of one move per year.

6.3.1 Step 1: Estimating Reduced-Form Hazards

We estimate the probabilities of opening a store, closing a store (if the firm has at least one store), and doing nothing using a linear-in-parameters multinomial logit sieve, with the parameters varying by firm type (chain, Walmart, and fringe).¹⁸ In particular, let $\tilde{p}_{ij}(k, z, \alpha)$ denote the reduced form probability of firm i making choice j in state (k, z) , which has the form

$$\tilde{p}_{ij}(k, z, \alpha) = \frac{\exp(h_j(k, z, \alpha))}{\sum_{j' \in \mathcal{A}_{ik}} \exp(h_{j'}(k, z, \alpha))},$$

where $h_j(k, z, \alpha)$ is a flexible function of the state variables with finite dimensional parameter vector α . The likelihood of a particular event, choice j by firm i in state k , in a market

¹⁸The variables included in the multinomial logit models are the number of fringe stores and its square, the number of chain stores and its square, the number of Walmarts and its square, and the total number of firms squared, and interactions of each of these variables with population. In addition, we control for city growth type, the unobserved state, and the unobserved state interacted with an indicator for building a new store.

with unobserved state z , with an interval of length τ since the previous event, is

$$\lambda \tilde{p}_{ij}(k, z, \alpha) \exp \left[- \left(\sum_{j' \in \{-1, 1\}} q_{j'} + \sum_m \lambda \sum_{j' \in \mathcal{A}_{mk}} \tilde{p}_{mj'}(k, z, \alpha) \right) \tau \right].$$

Since we have annual data, we simulate possible sequences of events that can happen over the course of each year. As discussed earlier, the structure of our data is such that we observe all events that took place in each year, but do not observe the exact times at which these events occur. Suppose that we observe W events in period n . For periods with $W > 0$, we draw R simulated paths, randomly assigning each observed event to a simulated time. Once we have the likelihood of each simulated sequence of events, we average over these simulated sequences, integrating over move times.

Focusing on a particular observation period n , let k_{n-1} and k_n denote the states at the beginning and end of the period. Let $k_w^{(r)}$ denote the state immediately preceding event w in simulation r , with $w = 1, \dots, W + 1$. The observed state at the beginning of the period is then $k_1^{(r)}$, so $k_1^{(r)} = k_{n-1}$ for each r . Similarly, the terminal state for each path is the observed state at the end of the period, so $k_{W+1}^{(r)} = k_n$ for each r . Let $I_w^{(r)}(i, j)$ be the indicator for whether event w of the r -th simulation was action j taken by firm i . Conditional on knowing the unobserved state z , the simulated likelihood of observation n in market m is

$$\begin{aligned} \tilde{L}_{mn}(q, \lambda, \alpha; z) &= \frac{1}{R} \sum_{r=1}^R \left\{ \prod_{w=1}^W \left[\sum_{j \in \{-1, 1\}} I_w^{(r)}(0, j) q_j + \sum_i \lambda \sum_{j \neq 0} I_w^{(r)}(i, j) \tilde{p}_{ij}(k_w^{(r)}, z, \alpha) \right] \right. \\ &\quad \times \exp \left[- \left(\sum_{j \in \{-1, 1\}} q_j + \sum_i \lambda \sum_{j \neq 0} \tilde{p}_{ij}(k_w^{(r)}, z, \alpha) \right) \tau_w^{(r)} \right] \left. \right\} \\ &\quad \times \exp \left[- \left(\sum_{j \in \{-1, 1\}} q_j + \sum_i \lambda \sum_{j \neq 0} \tilde{p}_{ij}(k_{W+1}^{(r)}, z, \alpha) \right) (1 - t_W^{(r)}) \right], \end{aligned}$$

where W is the number of events that occurred in the n -th interval and $t_W^{(r)}$ is the time of the last simulated move.

Since z is unobserved, we estimate the reduced form hazards using mixture distributions. Higher value of the unobserved state may make it easier or harder to operate as a chain, fringe, or Walmart store respectively. We discretize the standard normal distribution into five points and then estimate the population probabilities of being at each of these points. Note that there is an initial conditions problem here, so we allow the prior probability of being in a particular unobserved state to depend on the first period state variables, similar to Keane and Wolpin (1997) and Arcidiacono, Sieg, and Sloan (2007). In particular, we

specify the prior probabilities as following an ordered logit that depends on the number of chain stores, the number of Walmarts, and the number of fringe stores, all interacted with population, and the city growth type.

Let $\pi_z(k_1)$ be the probability of the unobserved state being z given that the observed state was k_1 for the first observation. With M markets and N periods each, integrating with respect to the distribution of the unobserved market states yields

$$\{\tilde{q}, \tilde{\lambda}, \tilde{\alpha}, \tilde{\pi}\} = \arg \max_{(q, \lambda, \alpha, \pi)} \sum_{m=1}^M \ln \left(\sum_z \pi_z(k_{m1}) \prod_{n=1}^N \tilde{L}_{mn}(q, \lambda, \alpha; z) \right).$$

The first stage estimates then give both the reduced form hazards, which are subsequently used in the second stage to form the value functions, as well as the conditional probability of each location being in each of the unobserved states, as in [Arcidiacono and Miller \(2011\)](#).

6.3.2 Step 2: Estimating the Structural Parameters

In Step 2, we take the probabilities of being in each unobserved state and the reduced-form hazards from Step 1 as given. We then separately estimate the structural parameters for chain and fringe stores. Let $P_m(z)$ denote the probability of MSA m being in unobserved state z given the data. Using Bayes' rule, we have

$$P_m(z) = \frac{\pi_z(k_{m1}) \prod_{n=1}^N \tilde{L}_{mn}(\tilde{q}, \tilde{\lambda}, \tilde{\alpha}; z)}{\sum_{z'} \pi_{z'}(k_{m1}) \prod_{n=1}^N \tilde{L}_{mn}(\tilde{q}, \tilde{\lambda}, \tilde{\alpha}; z')}.$$

These probabilities are then used as weights in the likelihood function for Step 2.

Next, using [Proposition 3](#), we can express the value function in (18) as a function of the structural parameters, θ , and the reduced form hazards from the first stage, $\tilde{p}_{ij}(k, z, \tilde{\alpha})$. Thus, let $\tilde{\sigma}_{ij}(k, z, \theta)$ denote the implied probability that firm i takes action j in state (k, z) , for a given value of the second stage parameters.

Define $\check{L}_{mn}(\theta; \tilde{q}, \tilde{\lambda}, \tilde{\alpha}; z)$ as

$$\begin{aligned} \check{L}_{mn}(\theta; \tilde{q}, \tilde{\lambda}, \tilde{\alpha}, z) &= \frac{1}{R} \sum_{r=1}^R \left\{ \prod_{w=1}^W \left[\sum_{j \in \{-1, 1\}} I_w^{(r)}(0, j) q_j + \sum_i \lambda \sum_{j \neq 0} I_w^{(r)}(i, j) \tilde{\sigma}_{ij}(k_w^{(r)}, z, \theta) \right] \right. \\ &\quad \times \exp \left[- \left(\sum_{j \in \{-1, 1\}} q_j + \sum_i \lambda \sum_{j \neq 0} \tilde{\sigma}_{ij}(k_w^{(r)}, z, \theta) \right) \tau_w^{(r)} \right] \left. \right\} \\ &\quad \times \exp \left[- \left(\sum_{j \in \{-1, 1\}} q_j + \sum_i \lambda \sum_{j \neq 0} \tilde{\sigma}_{ij}(k_{W+1}^{(r)}, z, \theta) \right) (1 - t_W^{(r)}) \right], \end{aligned}$$

where we use the same simulation draws as in Step 1, but replace each \tilde{p}_{ij} with $\tilde{\sigma}_{ij}$, which

is a function of θ . Then, the second stage estimates are

$$\{\check{\theta}\} = \arg \max_{\theta} \sum_{m=1}^M \sum_z P_m(z) \sum_{n=1}^N \ln \check{L}_{mn}(\theta; \tilde{q}, \tilde{\lambda}, \tilde{\alpha}, z).$$

6.4 Results

The structural parameter estimates for chain stores are presented in [Table 3](#). While the competitive effects of other chain stores and fringe stores are significant, they are dominated by the effects of Walmart. In terms of flow profit, one Walmart is equivalent to between six and seven fringe or chain stores. With the negative effects of number of stores squared, operating more stores becomes increasingly costly, effectively bounding the size of chain firms. Markets with higher values of the unobserved state face lower building costs for chain firms and lower diminishing returns for increasing chain size, but the costs of entering the market are higher. Other coefficients are as expected—population increases profits and the costs of building stores is substantial, with even higher costs incurred for entering a market.

Results for fringe firms are presented in [Table 4](#). Due perhaps to the importance of distribution networks (which rely on achieving a minimal local scale), having more fringe competitors raises profits at first, with competitive effects dominating as the number of fringe competitors increases. Both chain competitors and Walmarts lower profits but the effects of Walmart on fringe profits are substantially smaller than the effects on chain store profits. Population again has a positive effect on profits and there are significant store building costs. Similar to chain stores, higher values of the unobserved state lower store building costs and lessen the competitive impact from fringe competitors. However, this latter effect is smaller for fringe stores.

6.5 Counterfactuals

We consider two counterfactual experiments in which Walmart is prohibited from entering the markets in which it was not present at the beginning of our sample. We compare the resulting dynamics in these markets, both in the short run and long run, under our estimated structural parameters and under two counterfactual simulation designs. These experiments are designed to separate the ex-post direct competitive effects of Walmart’s entry on chain and fringe firms from the ex-ante effect of firms’ expectations regarding Walmart’s entry decisions. In the first experiment, Walmart is prohibited from entering but rival chain and fringe firms are not aware of this, and thus expect Walmart to behave as usual (which most likely means entering soon). In the second experiment, Walmart is again prohibited from entering, but this is common knowledge to rival chain and fringe firms, who then rationally

behave as if Walmart does not exist.

In the first case, we do not solve for counterfactual conditional choice probabilities for fringe and chain firms, but rather, assume that firms will behave according to their estimated policies, fully expecting Walmart to enter, while we artificially prohibit Walmart from acting. This out-of-equilibrium design shows that even without an actual Walmart store in any of these markets, the mere expected threat of entry by Walmart reduces entry by other firms. Interestingly, the effect of this threat is quite different across markets, both in terms of the level of the unobserved heterogeneity state and by market structure.

Our second simulation involves a counterfactual equilibrium in the usual sense, where we use the estimated structural parameters to obtain a new equilibrium between chain and fringe firms and simulate the evolution of these markets in a world where Walmart does not enter the retail grocery market at all.¹⁹ Firms in this experiment start from the same initial conditions observed at the beginning of our sample, but with the knowledge that Walmart will not enter, and thus operate without either the direct competitive effect or the threat of entry in expectation. Importantly, in neither case do we need Walmart’s structural parameters to carry out the simulations—only their conditional choice probabilities.

For the second experiment, we already have estimates of u_{ik} and q_j . With $\lambda = 1$, we solve for the value functions with Walmart eliminated, \ddot{V}_{ik} , using a suitably adapted version of equation (18). In the counterfactual equilibrium, we have

$$\ddot{V}_{ik} = \frac{u_{ik} + \sum_{j \in \{-1,1\}} q_j \ddot{V}_{i,l(0,j,k)} + \sum_{m \neq i} \lambda \sum_{j \in \{-1,1\}} \ddot{\sigma}_{mjk} \ddot{V}_{i,l(m,j,k)} + \lambda \mathbf{E} \max_j \left\{ \ddot{V}_{i,l(i,j,k)} + \psi_{ijk} + \varepsilon_{ijk} \right\}}{\rho + \sum_{j \in \{-1,1\}} q_j + \sum_{m \neq i} \lambda \sum_{j \in \{-1,1\}} \ddot{\sigma}'_{ijk} + \lambda}, \quad (19)$$

where the sum over rival firms no longer includes Walmart and where the values $\ddot{\sigma}_{ijk}$ are the counterfactual equilibrium conditional choice probabilities.

We simulate counterfactual outcomes over a wide range of markets in which Walmart has no presence at the beginning of our dataset. We first simulate the market evolution using the estimated conditional choice probabilities, given that Walmart enters as usual, and that firms have rational expectations about its entry behavior. We simulate 10,000 paths for each market in this way and report the average number of stores for each type, as well as the number of chain firms in the initial period and 5, 10, and 20 years in the future.

As seen in the simulations in Table 5, and the parameter estimates in Tables 3 and 4,

¹⁹Since we are simulating individual markets, none of which are large in all dimensions simultaneously, we are able to bound the state spaces for our simulations so that the value functions can be stored in memory. As such, we set appropriate bounds on the number of stores per chain and fringe firms for each simulation, while ensuring that the bounds are neither binding, nor near to binding, in our simulations. For example, although there would be over 287 billion states in the full model (in order to accommodate all markets we observe), the number of states in our simulations range from 20 million (Merced, CA, which only has a few small chains) to 183 million (Boulder, CO, which has several large chains).

the unobserved state appears to relate strongly to the direction of the competitive effect on fringe firms. In markets where the unobserved state is negative, such as Bend and Merced, Walmart tends to compete more strongly with fringe firms, to the benefit of chain firms, which are able to expand more quickly as a result. The opposite happens in markets with positive values of the unobserved state, such as Boulder and Erie. In these markets, Walmart actually tends to benefit fringe firms by competing more directly with chain firms instead. When the unobserved state is zero, the competitive effect is more balanced, as we see in Santa Cruz and Trenton.

Aside from these overall patterns, the different dynamics that occur in more concentrated markets versus more competitive markets are also of interest. Consider Bend, which is initially populated by twelve fringe stores but only six chain stores, operated by three different firms. In the counterfactual simulations, the number of fringe firms nearly doubles after 20 years, but in both the actual and expectations-only simulations, two fringe firms actually exit on average over the same period. It appears that the mere threat of Walmart's entry is enough to hinder fringe expansion in this market.

On the other hand, Boulder is initially served primarily by several moderately large chain firms, which collectively have twice as many stores as fringe firms. In this case Walmart's entry primarily limits expansion by chain firms, but also results in exit of about four fringe firms on average after 20 years. In the counterfactual simulations where Walmart does not enter at all, fringe firms actually appear to be edged out to a larger extent, due to more dominant chains, while under both the actual policies and the policies where Walmart is expected to enter, the more limited expansion by chain firms allows more fringe firms to remain in the market.

Thus, there appears to be much heterogeneity in the effect of Walmart's entry on the fringe. Both the initial market structure and other market-specific factors are important determinants of the nature of Walmart's competitive effect. The most interesting finding, however, is that the mere threat of Walmart has such a substantial effect on industry dynamics.

7 Conclusion

While recently developed two-step estimation methods have made it possible to estimate large-scale dynamic games, performing simulations for counterfactual work or generating data remains severely constrained by the curse of dimensionality that arises from simultaneous moves. We recast the standard discrete-time, simultaneous-move game as a sequential-move game in continuous time. This breaks the curse of dimensionality, greatly expanding the breadth and applicability of these structural methods and making even full-solution

estimation feasible for very large games.

Furthermore, by building on an underlying discrete-choice random utility model, our model preserves many of the desirable features of discrete-time models. In particular, we show that the insights from two-step estimation methods can be applied directly in our framework, resulting in another order of magnitude reduction in computational costs during estimation. We also show how to extend the model to accommodate incomplete sampling schemes, including time-aggregated data. Both are likely to be relevant for real-world datasets.

Our empirical application shows that our model is both feasible to estimate and can be used to run counterfactuals on problems with extremely large state spaces. Our empirical results are of interest in their own right, establishing substantial heterogeneity across markets in the behavior of supermarkets as well as establishing the importance of accounting for dynamics. Our counterfactuals reveal that the mere threat of Walmart entry has a substantial effect on industry dynamics, with the dominant firm type (either chains or fringe) holding back on expansion when Walmart may enter. This allows the weaker firm type to prosper more than it would if Walmart was banned.

A Proofs

A.1 Proof of Proposition 1

Let $v_{jk} = \psi_{jk} + V_{l(j,k)}$ denote the choice-specific value function, net of the choice-specific shock, for choice j in state k . Given the additively separable structure of instantaneous payoffs, we can express the conditional choice probabilities in terms of the social surplus function of McFadden (1981):

$$S(v_{0k}, \dots, v_{J-1,k} \mid k) \equiv \mathbb{E} \max_{j \in \mathcal{A}} \{v_{jk} + \varepsilon_j\}.$$

This function depends implicitly on the joint distribution of ε . Under [Assumption 2](#), by the Williams-Daly-Zachary theorem ([Rust, 1994](#), Theorem 3.1), the function S exists, S is additive in the sense that for any $\alpha \in \mathbb{R}$,

$$S(v_{0k} + \alpha, \dots, v_{J-1,k} + \alpha \mid k) = S(v_{0k}, \dots, v_{J-1,k} \mid k) + \alpha,$$

and the vector of CCPs equals the gradient of S : $\sigma_k = \nabla S(v_{0k}, \dots, v_{J-1,k} \mid k)$. Let $\tilde{v}_k = (v_{0k} - v_{j',k}, \dots, v_{J-1,k} - v_{j',k})$ denote the $J - 1$ vector of differenced choice-specific value functions, relative to choice j' , with the j' -th component omitted.

It follows from Proposition 1 of [Hotz and Miller \(1993\)](#) (also see Lemma 3.1 of [Rust](#)

(1994)) that there is a one-to-one function H which maps the differenced choice-specific value functions \tilde{v}_k in \mathbb{R}^{J-1} to choice probabilities σ_k in Δ^J , the J -dimensional unit simplex in \mathbb{R}^J . The result follows by noting that the j -th component of the inverse mapping yields $\tilde{v}_{jk} = \psi_{jk} - \psi_{j'k} + V_{l(j,k)} - V_{l(j',k)}$ as a function of j , j' , and σ_k .

A.2 Proof of Proposition 2

Recall from the proof of Proposition 1 that, by the Williams-Daly-Zachary theorem (Rust, 1994, Theorem 3.1) and Proposition 1 of Hotz and Miller (1993), we have $\tilde{v}_k = H^{-1}(\sigma_k)$. Then, by the aforementioned additivity property of S ,

$$\begin{aligned} \mathbb{E} \max_j \{v_{jk} + \varepsilon_j\} &= \mathbb{E} \max_j \{v_{jk} - v_{j'k} + \varepsilon_j - \varepsilon_{j'}\} + v_{j'k} + \mathbb{E} \varepsilon_{j'} \\ &= S(v_{0k} - v_{j'k}, \dots, v_{J-1,k} - v_{j'k} \mid k) + v_{j'k} + \mathbb{E} \varepsilon_{j'} \\ &= V_{l(j',k)} + \psi_{j'k} + \Gamma^2(j', \sigma_k), \end{aligned}$$

where $\Gamma^2(j', \sigma_k) = S(H^{-1}(\sigma_k) \mid k) + \mathbb{E} \varepsilon_{j'}$.

A.3 Proof of Proposition 3

For simplicity, suppose that $j = 0$ is a continuation action such that $l(0, k) = k$.²⁰ Let $(j_k^1, \dots, j_k^{D_k})$ denote a generic sequence of D_k decisions by which state k^* is attainable from state k . Similarly, let l_k^d denote the intermediate state in which the d -th decision is made, where $l_k^1 = k$ and $l_k^d = l(j_k^{d-1}, l_k^{d-1})$. Then, by recursively applying Proposition 1 for the continuation choice $j = 0$, we can write

$$V_k = V_{k^*} + \sum_{d=1}^{D_k} \left(\psi_{j_k^d, l_k^d} - \psi_{0, l_k^d} \right) + \sum_{d=1}^{D_k} \Gamma^1(0, j_k^d, \sigma_{l_k^d}). \quad (20)$$

Applying a similar procedure for each $l \neq k$ for which $q_{kl} > 0$ implies that we can write the differences $V_l - V_k$ on the right-hand side of (2) in terms of a difference of terms of the form in (20), where the V_{k^*} term cancels leaving only sums of instantaneous payoffs ψ_{jk} and functions of the CCPs σ_k . Finally, using Proposition 2 and additivity, we can express the remaining term $\lambda \mathbb{E} \max_j \{\psi_{jk} + \varepsilon_j + V_{l(j,k)} - V_k\}$ as $\lambda \Gamma^2(0, \sigma_k) + \lambda \psi_{0k}$.

A.4 Proof of Proposition 4

Given a collection of equilibrium best response probabilities $\{\sigma_i\}_{i=1}^N$, we can obtain a matrix expression for the value function $V_i(\sigma_i)$ by rewriting (7). Let $\Sigma_m(\sigma_m)$ denote the $K \times K$

²⁰Otherwise, we have to begin from some suitable state k' and choice j' such that $l(j', k') = k$ so that V_k is on the left-hand side of (4).

state transition matrix induced by the choice probabilities σ_m and the continuation state function $l(m, \cdot, \cdot)$. Let \tilde{Q}_0 denote the matrix formed by replacing the diagonal elements of Q_0 with zeros. Finally, let $E_i(\sigma_i)$ be the $K \times 1$ matrix containing the ex-ante expected value of the immediate payoff (both the instantaneous payoff and the choice-specific shock) for player i . That is, the k -th element of $E_i(\sigma_i)$ is $\sum_j \sigma_{ijk} [\psi_{ijk} + e_{ijk}(\sigma_i)]$ where $e_{ijk}(\sigma_i)$ is the expected value of ε_{ijk} given that choice j is optimal:

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

Given [Proposition 1](#), the difference $V_{i,l(i,j,k)}(\sigma_i) - V_{i,l(i,j',k)}(\sigma_i)$ can be expressed as a function of payoffs and choice probabilities σ_i .

Then, in matrix form,

$$\begin{aligned} V_i(\sigma_i) & \left[(\rho + \mathcal{N}\lambda) I - (Q_0 - \tilde{Q}_0) \right] \\ & = u_i + \tilde{Q}_0 V_i(\sigma_i) + \sum_{m \neq i} \lambda \Sigma_m(\sigma_m) V_i(\sigma_i) + \lambda [\Sigma_i(\sigma_i) V_i(\sigma_i) + E_i(\sigma_i)]. \end{aligned}$$

Collecting terms involving $V_i(\sigma_i)$ yields

$$V_i(\sigma_i) \left[(\rho + \mathcal{N}\lambda) I - \sum_{m=1}^{\mathcal{N}} \lambda \Sigma_m(\sigma_m) - Q_0 \right] = u_i + \lambda E_i(\sigma_i).$$

The matrix on the left hand side is strictly diagonally dominant since the diagonal of Q equals the off-diagonal row sums, the elements of each matrix $\Sigma_m(\sigma_m)$ are in $[0, 1]$ for all m , and $\rho > 0$ by [Assumption 1](#). Therefore, by the Levy-Desplanques theorem, this matrix is nonsingular ([Horn and Johnson, 1985](#), Theorem 6.1.10). Hence,

$$V_i(\sigma_i) = \left[(\rho + \mathcal{N}\lambda) I - \sum_{m=1}^{\mathcal{N}} \lambda \Sigma_m(\sigma_m) - Q_0 \right]^{-1} [u_i + \lambda E_i(\sigma_i)].$$

This representation is similar to the linear expression for the *ex-ante* value function in the discrete time model of [Pesendorfer and Schmidt-Dengler \(2007\)](#).

Now, define the mapping $\Upsilon : [0, 1]^{\mathcal{N} \times J \times K} \rightarrow [0, 1]^{\mathcal{N} \times J \times K}$ by stacking the best response probabilities. This mapping defines a fixed point problem for the equilibrium choice probabilities σ_{ijk} as follows:

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

The mapping Υ is a continuous function from a compact space onto itself. By Brouwer's

theorem, it has a fixed point. The fixed point probabilities imply Markov strategies that constitute a Markov perfect equilibrium.

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.
- Altuğ, S. and R. A. Miller (1998). The effect of work experience on female wages and labour supply. *Review of Economic Studies* 65, 45–85.
- Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79, 1823–1867.
- Arcidiacono, P., H. Sieg, and F. Sloan (2007). Living rationally under the volcano? An empirical analysis of heavy drinking and smoking. *International Economic Review* 48, 37–65.
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370.
- Bajari, P., H. Hong, and S. P. Ryan (2007). Identification and estimation of discrete games of complete information. *Econometrica* 78, 1529–1568.
- Bellman, R. E. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University Press.
- Beresteanu, A., P. B. Ellickson, and S. Misra (2007). The dynamics of retail oligopoly. Working paper, University of Rochester.
- Billingsley, P. (1961). *Statistical Inference for Markov Processes*. Chicago: University of Chicago Press.
- Doraszelski, U. and K. L. Judd (2007). Dynamic stochastic games with sequential state-to-state transitions. Working paper, Harvard University.
- Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics* 3, 53–93.

- Doraszelski, U. and M. Satterthwaite (2010). Computable Markov-perfect industry dynamics. *RAND Journal of Economics* 41, 215–243.
- Einav, L. (2010). Not all rivals look alike: Estimating an equilibrium model of the release date timing game. *Economic Inquiry* 48, 369–390.
- Elfving, G. (1937). Zur theorie der Markoffschen ketten. *Acta Societatis Scientiarum Fennicæ Nova Series A* 2, 1–17.
- Ellickson, P. B. and P. L. Grieco (2013). Wal-Mart and the geography of grocery retailing. *Journal of Urban Economics*.
- Ellickson, P. B., S. Houghton, and C. Timmins (2010). Estimating network economies in retail chains: A revealed preference approach. Working Paper 15832, NBER.
- Geweke, J. (1978). Temporal aggregation in the multiple regression model. *Econometrica* 46, 643–661.
- Geweke, J., R. C. Marshall, and G. A. Zarkin (1986). Exact inference for continuous time Markov chain models. *Review of Economic Studies* 53, 653–69.
- Gotz, G. A. and J. J. McCall (1980). Estimation in sequential decision-making models: A methodological note. *Economics Letters* 6, 131–136.
- Hansen, L. P. and T. J. Sargent (1983). The dimensionality of the aliasing problem in models with rational spectral densities. *Econometrica* 51, 377–387.
- Holmes, T. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica* 79, 253–302.
- Horn, R. A. and C. R. Johnson (1985). *Matrix Analysis*. New York: Cambridge University Press.
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–529.
- Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61, 265–289.
- Jia, P. (2008). What happens when Wal-Mart comes to town: An empirical analysis of the discount retail industry. *Econometrica* 76, 1263–1316.
- Johansen, S. (1974). Some results on the imbedding problem for finite Markov chains. *Journal of the London Mathematical Society* s2-8, 345–351.

- Karlin, S. and H. M. Taylor (1975). *A First Course in Stochastic Processes* (Second ed.). San Diego, CA: Academic Press.
- Keane, M. and K. I. Wolpin (1997). The career decisions of young men. *Journal of Political Economy* 105, 473–522.
- Kingman, J. F. C. (1962). The imbedding problem for finite Markov chains. *Probability Theory and Related Fields* 1, 14–24.
- Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory* 100, 191–219.
- McFadden, D. (1981). Econometric models of probabilistic choice. In C. F. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data*. Cambridge, MA: MIT Press.
- Milgrom, P. R. and R. J. Weber (1985). Distributional strategies for games with incomplete information. *Mathematics of Operations Research* 10, 619–632.
- Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political Economy* 92, 1086–1120.
- Moler, C. and C. V. Loan (1978). Nineteen dubious ways to compute the exponential of a matrix. *SIAM Review* 20, 801–836.
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* 54, 755–784.
- Pakes, A. and P. McGuire (1994). Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *The RAND Journal of Economics* 25, 555–589.
- Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric Markov perfect equilibrium, and the ‘curse’ of dimensionality. *Econometrica* 69, 1261–1281.
- Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics* 38, 373–399.
- Pesendorfer, M. and P. Schmidt-Dengler (2007). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75, 901–928.
- Phillips, P. C. (1973). The problem of identification in finite parameter continuous time models. *Journal of Econometrics* 1, 351–362.

- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55, 999–1013.
- Rust, J. (1994). Estimation of dynamic structural models, problems and prospects: Discrete decision processes. In C. Sims (Ed.), *Advances in Econometrics: Sixth World Congress*, Volume 2, Chapter 4, pp. 119–170. Cambridge: Cambridge University Press.
- Schmidt-Dengler, P. (2006). The timing of new technology adoption: The case of MRI. Working paper, London School of Economics.
- Sidje, R. B. (1998). Expokit: A software package for computing matrix exponentials. *ACM Transactions on Mathematical Software* 24, 130–156.
- Sims, C. A. (1971). Discrete approximations to continuous time distributed lags in econometrics. *Econometrica* 39, 545–563.
- Singer, B. and S. Spilerman (1976). The representation of social processes by Markov models. *The American Journal of Sociology* 82, 1–54.
- Su, C.-L. and K. L. Judd (2012). Constrained optimization approaches to estimation of structural models. *Econometrica* 80, 2213–2230.
- Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica* 76, 1375–1411.
- Wolpin, K. I. (1984). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy* 92, 852–874.

Table 1: Summary Statistics

	Mean	S.D.	Max.
Number of Chains Present ^a	2.559	0.024	7
Average No. of Stores per Chain ^b	3.727	0.040	32
Number of Walmarts Present ^a	1.004	0.142	12
Number of Fringe Firms Present ^a	12.997	0.823	47
Number of New Chain Stores ^c	0.277	0.012	5
Number of Exiting Chain Stores	0.224	0.011	7
Number of New Fringe Stores	0.825	0.021	10
Number of Exiting Fringe Stores	0.908	0.023	11
Number of New Walmarts	0.177	0.008	3
Number of Exiting Walmarts	0.002	0.001	1
Population Increase	0.042	0.004	1
Population Decrease	0.004	0.001	1

^a Sample size is 2910, ^b Sample size is 7446 and removes all market-period combinations where the chain operates no stores, ^c Sample size in this and all remaining rows is 2686.

Table 2: Response to Initial Walmart Entry

	Year Before	Year During	Year After
Number of New Chain Stores	0.311 (0.064)	0.211 (0.054)	0.189 (0.041)
Number of Exiting Chain Stores	0.122 (0.038)	0.156 (0.044)	0.189 (0.050)
Number of New Fringe Stores	0.867 (0.117)	0.711 (0.105)	0.767 (0.102)
Number of Exiting Fringe Stores	0.789 (0.114)	0.844 (0.118)	0.833 (0.132)

Standard errors in parentheses. Based on 90 markets where Walmart is first observed to enter.

Table 3: Chain Store Parameters

	Coefficient	S.E.
Constant	4.7687	0.1424
Number of chain stores	-0.0611	0.0056
Number of Walmarts	-0.4248	0.0237
Number of Fringe stores	-0.0665	0.0054
Population	0.2281	0.0377
Number of Own Stores	-0.0439	0.0215
Number of Own Stores Squared / 100	-0.3132	0.1057
Store building cost	-5.3886	0.0886
Chain entry cost	-17.9249	0.1267
Exit value	15.8385	0.1285
Unobserved State	-1.0256	0.2308
Store building cost \times unobserved state	3.5965	0.2225
Exit value \times unobserved state	4.0021	0.3259
Number of own stores \times unobserved state	0.3114	0.0383
Chain entry cost \times unobserved state	-4.5012	0.2698

Table 4: Fringe Store Parameters

	Coefficient	S.E.
Constant	-12.8481	0.0803
Number of Chain stores	-0.0236	0.0033
Number of Walmarts	-0.0324	0.0123
Number of Fringe Stores	0.1736	0.008
Number of Fringe Stores Squared /100	-0.3233	0.0186
Population	0.2176	0.0211
Entry Cost	-5.0703	0.0329
Entry Cost \times unobserved state	1.2347	0.0829
Unobserved State	-2.4846	0.1125
Unobserved State \times number of fringe stores	0.0437	0.0059

Table 5: Counterfactual simulations for selected cities starting with 1994 population and market structure

City	Type	Period	Actual CCPs				Walmart Absent			Counterfactual CCPs		
			Chain Firms	Chains	Fringe	Walmart	Chain Firms	Chains	Fringe	Chain Firms	Chains	Fringe
Bend 87,688	H,-	Year 0	3.0	6.0	12.0	0.0	3.0	6.0	12.0	3.0	6.0	12.0
		Year 5	2.1	5.3	11.1	0.5	2.1	5.4	11.1	2.8	5.7	15.6
		Year 10	1.9	5.2	10.5	0.9	2.0	5.6	10.8	2.7	5.3	19.1
		Year 20	1.7	5.0	10.0	1.4	1.8	6.0	11.0	2.4	4.8	22.4
Merced 191,195	H,-	Year 0	1.0	6.0	15.0	0.0	1.0	6.0	15.0	1.0	6.0	15.0
		Year 5	1.1	6.0	13.0	0.5	1.1	6.1	13.0	1.1	5.6	18.5
		Year 10	1.1	6.0	11.9	0.9	1.1	6.1	12.0	1.1	5.1	22.0
		Year 20	1.1	5.8	10.7	1.4	1.1	6.2	11.8	1.1	4.4	26.7
Santa Cruz 255,602	M,0	Year 0	3.0	10.0	20.0	0.0	3.0	10.0	20.0	3.0	10.0	20.0
		Year 5	2.6	9.2	21.3	0.5	2.6	9.2	21.6	3.1	10.7	19.3
		Year 10	2.2	8.4	22.5	1.1	2.2	8.6	23.3	3.2	11.4	18.6
		Year 20	1.7	7.1	24.3	2.3	1.9	8.0	26.4	3.3	12.8	17.4
Trenton 328,368	M,0	Year 0	3.0	6.0	24.0	0.0	3.0	6.0	24.0	3.0	6.0	24.0
		Year 5	2.0	4.9	25.7	0.7	2.0	5.0	25.9	2.9	6.6	23.2
		Year 10	1.7	4.4	27.0	1.5	1.7	4.7	27.8	2.9	7.3	22.5
		Year 20	1.2	3.4	28.3	3.0	1.5	4.7	29.8	2.9	8.6	21.1
Erie 279,117	L,+	Year 0	2.0	13.0	23.0	0.0	2.0	13.0	23.0	2.0	13.0	23.0
		Year 5	2.2	14.0	21.8	1.2	2.3	14.4	21.1	4.0	19.0	14.6
		Year 10	2.2	13.9	21.7	2.6	2.3	15.6	19.6	4.5	24.9	9.0
		Year 20	2.0	12.0	23.3	5.0	2.3	17.5	17.8	4.6	29.5	4.0
Boulder 243,377	H,+	Year 0	4.0	21.0	10.0	0.0	4.0	21.0	10.0	4.0	21.0	10.0
		Year 5	4.2	23.7	7.8	1.7	4.2	24.8	7.3	4.8	26.2	6.1
		Year 10	4.1	23.7	7.6	3.6	4.2	27.7	5.9	5.1	30.8	4.0
		Year 20	3.8	20.5	10.0	7.1	4.2	30.6	5.2	5.2	34.2	2.5

Note: H, M, and L refer to population growth types, -, 0, + refer to the value of the unobserved state. Walmart absent refers to the case where Walmart may enter, but has not done so. In the counterfactual equilibrium, it is common knowledge that Walmart is prohibited from entering.

Differences across these two columns are then driven by differences in expectations affecting current choices.