Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition*

Peter Arcidiacono[1], Patrick Bayer[1],
Jason R. Blevins[2], and Paul B. Ellickson[3]

[1] *Duke University and NBER*
[2] *Ohio State University*
[3] *Simon Graduate School of Business, University of Rochester*

December 17, 2015

Abstract. This paper develops a dynamic model of retail competition and uses it to study the impact of the expansion of a new national competitor on the structure of urban markets. In order to accommodate substantial heterogeneity (both observed and unobserved) across agents and markets, the paper first develops a general framework for estimating and solving dynamic discrete choice models in continuous time that is computationally light and readily applicable to dynamic games. In the proposed framework, players face a standard dynamic discrete choice problem at decision times that occur stochastically. The resulting stochastic-sequential structure naturally admits the use of CCP methods for estimation and makes it possible to compute counterfactual simulations for relatively high-dimensional games. The model and method are applied to the retail grocery industry, into which Wal-Mart began rapidly expanding in the early 1990s, eventually attaining a dominant position. We find that Wal-Mart's expansion into groceries came mostly at the expense of the large incumbent supermarket chains, rather than the single-store outlets that bore the brunt of its earlier conquest of the broader general merchandise sector. Instead, we find that independent grocers actually thrive when Wal-Mart enters, leading to an overall reduction in market concentration. These competitive effects are strongest in larger markets and those into which Wal-Mart expanded most rapidly, suggesting a diminishing role of scale and a greater emphasis on differentiation in this previously mature industry.

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

JEL Classification: C57, C35, L11, L81.

---

*We thank the attendees of the 2009 Cowles Foundation Conference on Structural Microeconomics, the 2010 cemmap Conference on Matching and Sorting, the 2012 NBER/NSF/CEME Conference on the Econometrics of Dynamic Games, the 2013 Meeting of the Midwest Econometrics Group, the 2014 Meeting of the Midwest Economics Association, the 2014 International Industrial Organization Conference, the 2014 University of Calgary Empirical

1 Introduction

Beginning with the static equilibrium entry models of Bresnahan and Reiss (1991) and Berry (1992), the modern empirical literature in industrial organization has sought to understand the determinants of market structure and its impact on market power and the nature of competition in oligopolistic industries. Recent papers have extended these models to permit forward-looking behavior on the part of firms, as well as more complex forms of investment and post-entry competition.[1] Adding such dynamic considerations broadens the nature of strategic interactions among firms and permits the study of a range of fundamentally dynamic phenomena such as preemption, predation, and limit pricing.

Using existing methods, incorporating forward-looking behavior in models of strategic interaction has been computationally costly, making it infeasible to compute the dynamic equilibrium unless the state space is sufficiently small. As a result, empirical researchers have often had to sacrifice much of the rich firm and market heterogeneity that can be incorporated in static models in order to study the kinds of interesting strategic behavior that can result when firms are forward-looking. With this trade-off in mind, the central goal of this paper is to provide a new approach for computing dynamic equilibrium models of market competition that is much lighter computationally and, therefore, permits the study of dynamic behavior without preventing researchers from incorporating dimensions of heterogeneity (both observed and unobserved) that may be critical for understanding key aspects of the nature of market competition.

At the outset, it is important to clarify that the key computational challenge in the dynamic games literature is related to the computation of the dynamic equilibrium (e.g., for use in counterfactuals) rather than estimation per se. In particular, since the seminal work of Hotz and Miller (1993) and Hotz, Miller, Sanders, and Smith (1994), conditional choice probability (CCP) estimators have been applied to a wide range of dynamic discrete choice problems including, more recently, simultaneous-move dynamic games (Aguirregabiria and Mira, 2007, Bajari et al., 2007, Pesendorfer and Schmidt-Dengler, 2008, Pakes et al., 2007). The key advantage of CCP estimation is that it eliminates the need to compute the full

Microeconomics Workshop, and the 2015 Econometric Society World Congress as well as seminar participants at Chicago (Economics and Booth), Columbia, Duke, Harvard, Iowa, Johns Hopkins, Kentucky, London School of Economics, Michigan, Northwestern, Ohio State (Economics and Fisher), Penn State, Rochester, Toronto, UBC (Economics and Sauder), UC Davis, UCLA, Virginia, Western Ontario, Washington University (Olin), Wisconsin, and Yale for useful comments. Timothy Schwuchow provided excellent research assistance.

[1] Building on earlier methodological contributions pioneered by Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pesendorfer and Schmidt-Dengler (2008), and Pakes, Ostrovsky, and Berry (2007), empirical researchers have recently examined the impact of environmental regulations on entry, investment and market power in the cement industry (Ryan, 2012), the effect of demand fluctuations in the concrete industry (Collard-Wexler, 2013), and the impact of increased royalty fees on the variety of products offered by commercial radio stations (Sweeting, 2013).

solution of the dynamic game, allowing empirical researchers to estimate relatively high-dimensional games, including many that can not be solved directly even once.

While knowledge of the parameters of many dynamic games can be informative, a key limitation of the application of CCP estimation to high-dimensional problems is that it is often impossible to compute counterfactual simulations of the estimated model. In the context of games, the calculation of players' expectations over all combinations of actions of their rivals grows exponentially in the number of players, making it computationally challenging to compute the equilibrium even in some relatively simple economic contexts.[2]

In this paper, we develop a characterization of a dynamic game in continuous time that not only alleviates some of the computational difficulties associated with simultaneous move games, but also links naturally with the existing literature on dynamic discrete choice models and dynamic discrete games. The key feature of our approach is that players face a standard dynamic discrete choice problem at decision times that occur stochastically. The resulting stochastic-sequential structure naturally admits the use of CCP methods for estimation and makes it possible to compute counterfactual simulations for relatively high-dimensional games.

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 continues to eliminate the need to compute the full solution of the model for estimation. Using our framework, the two-step estimators of Aguirregabiria and Mira (2002, 2007), Bajari et al. (2007), Hotz et al. (1994), Pakes et al. (2007), and Pesendorfer and Schmidt-Dengler (2008) can be applied in continuous time. 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.

We take advantage of this new formulation of a dynamic game in continuous time to study the impact of a new national competitor on the structure of urban markets across the United States. Specifically, we examine the impact of Wal-Mart's rapid expansion into the retail grocery industry from 1994–2006.[3] In particular, we model the decisions of Wal-Mart and its rivals over whether to operate grocery stores in a market and at what scale (i.e.,

[2] These 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).

[3] Our methods have also been applied by Nevskaya and Albuquerque (2012) to online games, by Schiraldi, Smith, and Takahasi (2012) to supermarkets, by Mazur (2014) to airlines, and by Cosman (2014) to bars in Chicago.

number of stores). We include the choices of Wal-Mart and up to seven competing chains, as well as the single-store entry and exit decisions of several dozen fringe players. 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 Wal-Mart, 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. We estimate the model using CCP methods and solve counterfactually for the equilibrium under several scenarios designed to measure how Wal-Mart's entry into the retail grocery industry affects the profitability and decision-making of rival chain and fringe firms.[4]

The estimates imply that Wal-Mart's entry has a substantial impact on market structure that is heterogeneous both across markets and firm types. In particular, we find that Wal-Mart's expansion came mostly at the expense of the large, incumbent grocery chains, leading some to exit the market completely and sharply diminishing the scale of others. In contrast, the small fringe firms actually thrive in the new market structure, suggesting a greater ability to differentiate themselves (e.g., in terms of product offerings or geography) from both Wal-Mart and the remaining chains. Taken as a whole, market concentration is sharply reduced by Wal-Mart's entry in these markets.

Notably, this new, entrepreneurial activity is strongest in the larger markets and those into which Wal-Mart expanded the fastest. In contrast, in another set of (primarily smaller, Western) markets, Wal-Mart's entry greatly concentrates the market in a way that closely resembles the impact of its initial entry into discount retailing two decades earlier (Jia, 2008). Wal-Mart's entry is felt most directly by the fringe firms in these markets because, much like the rural markets Wal-Mart focused on originally, chain stores were less established in these markets to begin with. However, in the vast majority of grocery markets, the fringe actually benefits from Wal-Mart's presence.

A comparison of the results for specifications with and without unobserved heterogeneity reveals that the inclusion of unobserved heterogeneity is essential for uncovering these qualitatively distinct economic implications of Wal-Mart's entry across markets. While there is still some variation across markets, the specification without unobserved heterogeneity implies that Wal-Mart's entry decreases market concentration in every market in the sample and sharply understates the positive impact on small, independent stores, especially in the larger markets. Taken as a whole, the results of our analysis demonstrate the importance of incorporating substantial heterogeneity both across markets and firm types in estimating dynamic games of retail entry and competition, thereby highlighting the advantage of com-

[4]Our counterfactual simulations involve calculating value functions at up to 157 million states in each of 205 markets, yet by taking advantage of the continuous time formulation this is computationally feasible.

putationally light approaches for estimating and solving dynamic models with large state spaces.

Our paper relates to the literature on both estimation of dynamic models in continuous time as well as the empirical literature on entry in retail markets.[5] Heckman and Singer (1986) advocated using continuous time models for economic processes to avoid specifying unnatural, fixed decision intervals for agents and because, unlike discrete time models, these models are functionally invariant when applied to data recorded at different time intervals. In effect, agents need not move simultaneously and move times need not match the intervals at which observations are recorded (e.g., annual or quarterly).

Continuous-time methods also have a long history in the search literature. Our approach is most closely connected to the literature in which job offers arrive at an exogenous rate. Another vein of the search literature, however, assumes that individuals receive offers according to their search intensity, effectively choosing a hazard rate for job offers. This is the spirit of Doraszelski and Judd (2012) who first illustrated the computational advantages of casting dynamic games into continuous time. 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). In our paper, we take advantage of the insight of Doraszelski and Judd (2012) regarding how continuous time, and more specifically sequential state-to-state transitions, can reduce the computational burden of dynamic games but reframe the problem in a way that (i) naturally parallels the discrete time discrete choice literature, (ii) retains the computational advantages of CCP methods in estimation, and (iii) circumvents issues of multiple equilibria in estimation.

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 establishes conditions for identification of the model primitives and then Section 6 develops our estimators and discusses issues associated with time aggregation. Section 7 introduces and describes the results of our empirical analysis of the market structure of grocery store chains in geographically separate U.S. markets. Section 8 concludes.

[5] On the empirical side, our paper is the first to estimate structurally the impact of Wal-Mart on both chain and single-store firms. Ellickson and Grieco (2013) examine the impact of Wal-Mart on the structure of the supermarket industry using descriptive methods from the treatment effects literature, while Basker and Noel (2009) and Matsa (2011) look at its impact on prices and quality. Wal-Mart's effect on discount retail has been analyzed by Jia (2008), Holmes (2011), and Ellickson, Houghton, and Timmins (2013).

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 a multi-agent context in the following section. We begin this section by laying out the notation and structure of the model in a general manner. 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 are not a result of actions by the agent. The elements of Q_0, denoted by q_{kl}, are the rates at which particular state transitions occur and are nonnegative and bounded. Second, a Poisson arrival process with rate parameter λ governs when the agent can move. When a move arrival occurs, the agent chooses an action j from the discrete choice set $\mathcal{A} = \{0, \dots, J-1\}$.

A finite-state Markov jump process 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 $l \neq k$

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

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

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

is the negative of the overall rate at which the process leaves state k. 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).

We assume that the agent is forward-looking and discounts future payoffs at rate $\rho \in (0, \infty)$. While in state k, the agent receives flow utility u_k with $|u_k| < \infty$. At rate λ the agent makes a decision, choosing an action $j \in \mathcal{A}$ and receiving an additively separable instantaneous payoff $\psi_{jk} + \varepsilon_j$, where ψ_{jk} is the mean payoff (or cost) associated with making choice j in state k, with $|\psi_{jk}| < \infty$, and $\varepsilon_j \in \mathbb{R}$ is an instantaneous choice-specific payoff shock which is observed by the agent but not by the econometrician.[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. Let $l(j,k)$ denote the state that results upon making choice j in state k. We use the convention that $j = 0$ is a costless continuation choice (i.e., $\psi_{0k} = 0$ and $l(0,k) = k$ for all k) and assume that the remaining actions are meaningfully distinct (i.e., $l(j,k) \neq l(j',k)$ for $j' \neq j$ for all k).

We can now specify 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:[7]

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

The denominator contains the discount factor plus the sum of 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 instantaneous and future value obtained when a move arrival occurs in state k. The expectation is with respect to the joint distribution of $\varepsilon = (\varepsilon_0, \dots, \varepsilon_{J-1})^\top$.

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

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

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

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

[6] Although the choice-specific shocks ε_j vary over time, we omit the t subscript for simplicity. We also assume the distribution of ε_j is identical across states, to avoid conditioning on k throughout, but allowing it to depend on k does not present additional difficulties. Finally, as formalized in Section 4, we assume the errors are i.i.d. with joint density f, finite first moments, and support $\mathbb{R}^J$.

[7] To derive the Bellman equation, note that event probabilities over a small time increment h under the Poisson assumption are proportional to h and the discount factor is $1/(1 + \rho h)$, then take the limit as $h \to 0$.

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 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 accommodating 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 single state variable captures 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 the initial 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

$$V_k = \frac{u_k + q_{k1}V_{k+1} + q_{k2}V_{k+2} + \lambda \operatorname{E} \max\left\{V_k + \varepsilon_0, V_0 + c + \varepsilon_1\right\}}{\rho + q_{k1} + q_{k2} + \lambda}, \tag{2}$$

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 CCP representation of the value function.

3 CCP Representation

In traditional discrete time dynamic discrete choice models, agents typically make decisions simultaneously at pre-determined intervals. In contrast, in our continuous time framework, only a single decision or state change occurs at any given instant (almost surely) and moves occur at random time intervals. Despite these key differences, our framework preserves the basic stochastic structure for decision-making from the discrete time literature and, as a result, many of the insights of the literature following Hotz and Miller (1993) on expressing value functions in terms of conditional choice probabilities (CCPs) apply here as well. In particular, as the following proposition shows, it is possible to eliminate the value functions on the right-hand side of (1), implying that no fixed point problem needs to be solved in estimation. All proofs are given in the Appendix.

Proposition 1. *The value function can be written as*

$$V(\sigma) = [(\rho+\lambda)I - \lambda\Sigma(\sigma) - Q_0]^{-1}\,[u + \lambda E(\sigma)],$$

where I is the $K \times K$ identity matrix, $\Sigma(\sigma)$ is the $K \times K$ state transition matrix induced by the choice probabilities σ due to actions by the agent, u is the $K \times 1$ vector of flow payoffs, and $E(\sigma)$ is the $K \times 1$ vector containing the ex-ante expected values of the instantaneous payoffs in each state, $\sum_j \sigma_{jk}\,[\psi_{jk} + e_{jk}(\sigma)]$ where $e_{jk}(\sigma)$ is the expected value of ε_{jk} given that choice j is optimal,

$$\frac{1}{\sigma_{jk}} \int \varepsilon_{jk} \cdot 1\left\{\varepsilon_{j'k} - \varepsilon_{jk} \le \psi_{jk} - \psi_{j'k} + V_{l(j,k)}(\sigma) - V_{l(j',k)}(\sigma)\ \forall j'\right\} f(\varepsilon_k)\, d\varepsilon_k.$$

The result above shows that the value function can be written as a linear function of the payoffs, CCPs, and hazards for nature. This is analogous to the results of Aguirregabiria and Mira (2002) for discrete time models.

While Proposition 1 can be applied to any setting that follows our stochastic-sequential structure, computation is further simplified in some cases. Namely, when a certain finite dependence condition holds, we can avoid calculating the $K \times K$ matrix inverse in Proposition 1 and instead express the value function in terms of CCPs for a limited number of states.

To do so, 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). Using the CCPs 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 opportunity to move. The second result establishes a similar CCP representation for the final term in the Bellman equation.

Proposition 2. *There exists a function $\Gamma^1(j, j', \sigma_k)$ such that for all $j, j' \in \mathcal{A}$,*

$$V_{l(j,k)} = V_{l(j',k)} + \psi_{j'k} - \psi_{jk} + \Gamma^1(j, j', \sigma_k). \tag{3}$$

Proposition 3. *There exists a function* $\Gamma^2(j', \sigma_k)$ *such that for all* $j' \in \mathcal{A}$,

$$\mathrm{E}\max_j \left\{ \psi_{jk} + \varepsilon_j + V_{l(j,k)} \right\} = V_{l(j',k)} + \psi_{j'k} + \Gamma^2(j', \sigma_k). \tag{4}$$

Proposition 2 states that the valuations can be linked across states. The intuition for Proposition 3 is that we can express the left hand side of (4) relative to $V_{l(j',k)} + \psi_{j'k}$ for an action j' of our choosing. By doing so, the terms inside the $\mathrm{E}\max$ term will consist of differences in value functions and instantaneous payoffs. These differences, as established by Proposition 2, can be expressed as functions of conditional choice probabilities.

For a concrete example, consider the case where the ε's follow the type I extreme value distribution. In this case, closed form expressions exist for both Γ^1 and Γ^2:[8]

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

where γ is Euler's constant.

Importantly, Proposition 2 allows us to link value functions across many states. 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, there are classes of models where the remaining value functions on the right hand side of (1) 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 4. *For a given state* k*, suppose that for any state* $l \neq k$ *with* $q_{kl} > 0$ *there exists a state* k^* *that is attainable from both* k *and* l*. Then, there exists a function* $\Gamma_k(\psi, Q_0, \sigma)$

[8] These expressions have closed forms in specific cases (e.g., multinomial logit or nested logit error structures) (Arcidiacono and Miller, 2011). This will be more difficult in other settings.

such that

$$\rho V_k = u_k + \Gamma_k(\psi, Q_0, \sigma). \tag{5}$$

The function Γ_k for each state may depend on the model primitives ψ and Q_0 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 and no $K \times K$ matrix need be stored or inverted. This can often lead to large computational gains. We now provide an example showing how to apply these propositions.

3.1 Example: A Single-Agent Renewal Model

Recall the bus engine replacement example of Section 2.1, in which the value function was characterized by (2). Applying Proposition 3 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:

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

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 4:

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

This example illustrates one of the benefits of continuous time over discrete time when using conditional choice probabilities. In particular, it is possible to write the value function directly in terms of CCPs in the continuous time framework without differencing with respect to a particular choice. As a result, as illustrated above, we only need to estimate replacement probabilities for a series of states k, $k+1$, and $k+2$ for any given k. In contrast, discrete time renewal problems require differencing, as illustrated by Arcidiacono and Miller (2011). This means that accurate estimates of the conditional probability of replacing the engine at very low mileages are needed for CCP estimation. Because these are low probability events, such estimates will likely depend heavily on the smoothing parameters or functional forms used to mitigate the associated small sample problems.

4 Dynamic Discrete Games

The potential advantages of modeling decisions using a continuous time framework are particularly salient in 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 settings.

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 transitions 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 N players indexed by $i = 1, \dots, N$. 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: $\mathcal{A}_{ik} = \{0, \dots, J-1\}$ for all i and k. 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 N independent, competing Poisson processes with rate λ generating move arrivals for each of the N players.[9] The next event to occur is determined by the earliest arrival of one of these $N + 1$ processes.

Let ς_i denote player i's beliefs regarding the actions of rival players, given by a collection of $(N-1) \times J \times K$ 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

[9]For simplicity, we assume the move arrival rates are equal for each firm.

player i in state k can be defined recursively as:

$$V_{ik}(\varsigma_i) = \frac{u_{ik} + \sum_{l \neq k} q_{kl} V_{il}(\varsigma_i) + \sum_{m \neq i} \lambda \sum_j \varsigma_{imjk} V_{i,l(m,j,k)}(\varsigma_i) + \lambda \operatorname{E} \max_j \left\{ \psi_{ijk} + \varepsilon_{ij} + V_{i,l(i,j,k)}(\varsigma_i) \right\}}{\rho + \sum_{l \neq k} q_{kl} + N\lambda}. \tag{6}$$

The denominator of this expression the sum of the discount factor and the rates of exogenous state changes and moves by players. The numerator is best understood by looking at each term separately. The first term is the flow payoff that accrues to firm i each instant the model remains in state k. Next, we have a sum over possible exogenous state changes, weighted by the rates at which those changes occur. The third term is a sum over the rate-weighted state changes that could result due to the action of a rival firm. The final term is the rate-weighted continuation value that occurs when the agent moves and optimally makes a discrete choice.

For dynamic games, we work under assumptions that generalize those used above for single agent models and largely mirror the standard assumptions used in discrete time.

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

Assumption 2 (Bounded Rates and Payoffs)**.** The discount rate ρ, move arrival rate, rates of state changes due to nature, and payoffs are all bounded for all $i = 1, \ldots N$, $j = 0, \ldots, J-1$, $k = 1, \ldots, K$, $l = 1, \ldots, K$ with $l \neq k$: (a) $0 < \rho < \infty$, (b) $0 < \lambda < \infty$, (c) $0 \leq q_{kl} < \infty$, (d) $|u_{ik}| < \infty$, and (e) $|\psi_{ijk}| < \infty$.

Assumption 3 (Additive Separability)**.** For each player i and in each state k the instantaneous payoff associated with choice j is additively separable as $\psi_{ijk} + \varepsilon_{ijk}$.

Assumption 4 (Distinct Actions)**.** For all $i = 1, \ldots, N$ and $k = 1, \ldots, K$, the continuation state function $l(i, j, k)$ and the choice-specific payoffs ψ_{ijk} satisfy the following two properties: (a) choice $j = 0$ is a costless continuation choice with $l(i, j, k) = k$ and $\psi_{ijk} = 0$, and (b) all choices j are meaningfully distinct in the sense that the continuation states differ: $l(i, j, k) \neq l(i, j', k)$ for all $j = 0, \ldots, J-1$ and $j' \neq j$.

Assumption 5 (Private Information)**.** The errors ε_{ik} are i.i.d. over time and across players with joint distribution F which is absolutely continuous with respect to Lebesgue measure (with joint density f), has finite first moments, and has support equal to $\mathbb{R}^J$.

Following Maskin and Tirole (2001), we focus on Markov perfect equilibria in pure strategies, as is standard in the discrete-time games 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, $\{\varsigma_1, \ldots, \varsigma_N\}$, and a collection of model primitives, a Markov strategy for player i is a *best response* if

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

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

$$\sigma_{ijk} = \Pr\left[\delta_i(k, \varepsilon_i; \varsigma_i) = j \mid k\right]. \tag{7}$$

Definition. A collection of Markov strategies $\{\delta_1, \ldots, \delta_N\}$ and beliefs $\{\varsigma_1, \ldots, \varsigma_N\}$ 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 $\{\sigma_1, \ldots, \sigma_N\}$ where each probability in σ_i is a best response given beliefs σ_{-i}.

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

Proposition 5. *If Assumptions 1–5 hold, then a Markov perfect equilibrium exists.*

4.1 CCP Representation

As in single-agent models, the value function for any player i in the multi-player case can be expressed in terms of reduced-form CCPs and hazards. The following proposition formalizes this and generalizes Proposition 1 to the multi-player setting.

Proposition 6. *If Assumptions 1–5 hold, then for each player i*

$$V_i(\sigma) = \left[(\rho + N\lambda) I - \sum_{m=1}^{N} \lambda \Sigma_m(\sigma_m) - Q_0\right]^{-1} \left[u_i + \lambda E_i(\sigma)\right] \tag{8}$$

where $\Sigma_m(\sigma_m)$ is the $K \times K$ state transition matrix induced by the actions of player m given the choice probabilities σ_m and where $E_i(\sigma)$ is a $K \times 1$ vector where each element k is the *ex-ante expected value of the choice-specific payoff in state k, $\sum_j \sigma_{ijk} [\psi_{ijk} + e_{ijk}(\sigma)]$ where $e_{ijk}(\sigma)$ is the expected value of ε_{ijk} given that choice j is optimal,*

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

This representation mirrors those of Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008) in discrete time and forms the basis of a two-step estimator discussed in the next section.

We further note that the propositions in Section 3 apply to games as well, though the attainability condition may be more difficult to satisfy. Namely, it is possible to eliminate the value functions in the fourth term of the numerator of (6) using Proposition 3:

$$V_{ik} = \frac{u_{ik} + \sum_{l \neq k} q_{kl} V_{il} + \sum_{m \neq i} \lambda \sum_j \sigma_{mjk} V_{i,l(m,j,k)} + \lambda \Gamma^2(0, \sigma_{ik})}{\rho + \sum_{l \neq k} q_{kl} + (N-1)\lambda}. \tag{9}$$

Eliminating the other value functions, however, is problematic as each player may only be able to move the process to some subset of the state space via a unilateral action, since they only have direct control over their own state.

There are important classes of models, however, for which the remaining value functions can be eliminated. In models with a terminal choice, such as a firm permanently exiting a market, the value function for the terminal choice does not include other value functions. Similarly, in models in which a player's action can reset the game for all players, the value function for reset can be expressed in terms of appropriate CCPs.

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. To see this, 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 the economy-wide state vector associated with state k is $x^k = (x_1^k, \ldots, x_N^k)$, which contains the store counts of all firms in the market (including potential entrants, namely firms with zero stores). Let $l^*(i, k, x_i')$ denote the index of the state that is equal to the initial state x^k, but where firm i has x_i' stores instead of x_i^k. Applying Proposition 2 and normalizing the value of zero stores to zero, we can express V_{ik} as:

$$V_{ik} = \sum_{x_i'=1}^{x_i^k} \Gamma^1 \left(0, -1, \sigma_{i,l^*(i,k,x_i')} \right) + x_i^k c. \tag{10}$$

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:

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

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 Identification

The structural model primitives are $(\psi_1, \ldots, \psi_N, u_1, \ldots, u_N, F, \rho, Q_0)$. Following Magnac and Thesmar (2002), Pesendorfer and Schmidt-Dengler (2008), and several others, we assume that the distribution of idiosyncratic errors, F, and the discount rate, ρ, are known and establish conditions for the identification of the remaining primitives. This proceeds in two steps. First, we establish that the aggregate intensity matrix Q for the continuous time model, and hence nature's intensity matrix Q_0, are identified even with only discrete time data. To do so, we exploit a priori restrictions on Q that arise from the construction of the model itself. Second, we show that knowledge of Q allows us to identify the remaining structural primitives—the flow payoffs $(u_1, \ldots, u_N)$ and the instantaneous payoffs $(\psi_1, \ldots, \psi_N)$—under conditions that are similar to those used for identification of discrete time models (e.g., payoff exclusion restrictions).

We follow most papers in the literature on two-step estimation (e.g., Bajari et al., 2007, Aguirregabiria and Mira, 2007) and assume that a single Markov perfect equilibrium is played in each state k and that all players expect the same equilibrium to be played at all times both in and out of sample. This assumption is analogous to similar assumptions commonly used in two-step estimation of discrete time models (see Aguirregabiria and Mira (2010) for a survey).

Assumption 6 (Multiple Equilibria)**.** (a) In each state $k = 1, \ldots, K$, a single Markov perfect equilibrium is played which results in an intensity matrix Q. (b) The distribution of state transitions at any point t is consistent with the intensity matrix Q.

The purpose of this assumption is to guarantee that we can consistently estimate either the state-to-state transition probability matrix based on discrete time data observed at fixed intervals of length Δ,[10] or the aggregate intensity matrix itself, Q, based on continuous time observations. We first discuss the choice of the rate of move arrivals before turning to identification of the intensity matrix and then the structural primitives.

5.1 Uniformization and the Rate of Move Arrivals

In discrete time models, the lengths of decision periods in the structural model are not known but are chosen by the researcher (usually to match the data sampling period). Because agents can choose to do nothing each period, the choice probabilities adjust to reflect the chosen time period. For example, the probability that a firm does not enter a market using quarterly data would be larger than with annual data (given the same intrinsic rates).

In a similar way, the rate of move arrivals in our model, λ, is not a parameter of interest but is chosen by the researcher. Introducing λ while giving players a continuation choice ($j = 0$) is related to a technique known as *uniformization* in the stochastic process literature (see e.g., Puterman, 2005, Ch. 11). In our case, it also has a convenient behavioral interpretation and allows us to maintain a multinomial choice structure that is analogous to discrete time models.

In a generic, stationary Markov jump process the rate at which the process leaves state k may in general differ across states. However, the same process has an equivalent representation in terms of a state-independent Poisson "clock process" with a sufficiently large common rate γ that governs *potential* transitions out of each state and an embedded Markov chain governing the actual state transitions. In this representation, with each arrival of the Poisson process the probability of remaining in the same state can be nonzero.

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 γ be the fastest transition rate

$$\gamma = N\lambda + \max_k \sum_{l \neq k} q_{kl}.$$

If we consider γ to be the overall move arrival rate for players, including nature, and if $Q = \sum_{i=0}^{N} Q_i$ is the aggregate intensity matrix with elements ω_{kl}, then at each arrival of

[10] Extending this to cases with irregular time intervals is straightforward and also helps with identification. See Blevins (2016) for a summary of results on identification of models with irregularly spaced observations including Cuthbert (1973), Singer and Spilerman (1976), and Hansen and Sargent (1983).

the Poisson process with rate γ the probability of the state transitioning from k to l is

$$p_{kl} = \begin{cases} \omega_{kl}/\gamma & \text{if } l \neq k, \\ 1 - \sum_{k \neq l} \omega_{kl}/\gamma & \text{if } l = k. \end{cases}$$

In other words, given Q we can construct an embedded Markov chain with transition matrix $I + \frac{1}{\gamma}Q$ that characterizes the state transitions for each arrival of the Poisson process.[11]

In the context of our dynamic discrete choice model, uniformization with a fixed rate γ is analogous to fixing the overall rate of decisions, λ, while allowing the choice $j = 0$ to a continuation decision. Choice probabilities and other model implications should then be interpreted relative to the choice of λ and the unit of time. In our application we set $\lambda = 1$ which implies that agents make decisions *on average* once per unit of time (e.g., one year) without restricting the actual realized move times to a fixed interval.

5.2 Identification of Q

With continuous-time data, identification and estimation of the intensity matrix for finite-state Markov jump processes is straightforward and well-established (Billingsley, 1961). On the other hand, when a continuous-time process is only sampled at discrete points in time, the parameters of the underlying continuous-time model may not be uniquely identified.[12] This is known as the aliasing problem and has been studied by many authors in the context of continuous-time vector autoregression models (Phillips, 1973, Hansen and Sargent, 1983, Geweke, 1978, Kessler and Rahbek, 2004, McCrorie, 2003, Blevins, 2016). In the present model, the concern is that there may be multiple Q matrices which give rise to the same transition probability matrix $P(\Delta)$.

Formally, 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

[11] To see the equivalence, let $\tilde{Q} = (\tilde{\omega}_{kl})$ denote the intensity matrix after uniformization. For transitions from k to $l \neq k$ we have $\tilde{\omega}_{kl} = \gamma(\frac{\omega_{kl}}{\gamma}) = \omega_{kl}$ and for transitions back to k we have $\tilde{\omega}_{kk} = -\gamma + \gamma p_{kk} = -\sum_{l \neq k} \omega_{kl} = \omega_{kk}$.

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

the *transition matrix*, which is the matrix exponential of ΔQ:

$$P(\Delta) = \exp(\Delta Q) = \sum_{j=0}^{\infty} \frac{(\Delta Q)^j}{j!}. \tag{11}$$

These transition probabilities account for all paths of intermediate jumps to other states between observed states (including possibly no jumps at all). We establish conditions under which there is a unique Q matrix that is consistent both with the model and (11).[13]

The theoretical model restricts Q to a lower-dimensional subspace since it is sparse and must satisfy certain within-row and across-row restrictions. We show that this sparse structure leads to unique identification of Q by establishing identification conditions based on linear restrictions on the Q matrix of the form $R\,\text{vec}(Q) = r$.[14] The following proposition establishes that there are sufficiently many restrictions of full rank to identify Q in a broad class of continuous time discrete choice games.

Proposition 7. *Suppose that Assumptions 1–6 hold and let $x^k = (x_1^k, \dots, x_N^k)$ be the vector of player-specific states corresponding to state k, where each component $x_i^k \in \mathcal{X}_i$ can only be affected by the action of player i. Suppose that there are $|\mathcal{X}_i| = \kappa$ possible states per player, J actions per player, and $N \geq 2$ players. Suppose that Q has distinct eigenvalues that do not differ by an integer multiple of $2\pi i/\Delta$. Then Q is generically identified[15] from $P(\Delta)$ if $J \leq \frac{\kappa^N + 2N - 1}{2N}$.*

To better understand the conclusion, we consider some common cases. First, in any non-trivial binary choice game Q is identified. That is, if $J = 2$ and $\kappa \geq 2$ then Q is identified for any value of N. Similarly, for $J = 3$ choices and $\kappa \geq 3$, then Q is identified for any N. Finally, for any $N \geq 1$ and $\kappa > 1$, the model is identified as long as J is not too large, where the upper bound on J is increasing exponentially in N. The result can be extended to cases with exogenous state variables and where the number of endogenous states per player, κ, may differ across players and states.

For identification purposes, we assume that given the aggregate intensity matrix Q we can determine the player-specific intensity matrices Q_i for $i = 0, \dots, N$.

Assumption 7. The mapping $Q \to \{Q_0, Q_1, \dots, Q_N\}$ is known.

[13] A related issue is the embeddability problem: could $P(\Delta)$ have been generated by a Markov jump process with intensity matrix Q? 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 $K = 3$.

[14] For the $K \times K$ matrix $Q = (q_{kl})$, $\text{vec}(Q)$ is the vector obtained by stacking the columns of Q.

[15] We say Q is *generically identified* if it is identified except possibly for a measure zero set of population Q matrices. See Phillips (1973, p. 357) for a detailed discussion of generic identification in a 3×3 model.

This assumption is satisfied in applications where firms cannot change each other's state variables and where actions by nature can be distinguished from the actions of firms. It holds trivially in the single-agent example above and in our empirical application, where the Q matrix is sparse and each non-zero off-diagonal element corresponds to a single matrix Q_i. A sufficient condition in general is that the continuation states resulting from the actions of different players are distinct: for all players i and $m \neq i$ and all states k, $\{l(i,j,k) : j = 1,\dots,J-1\} \cap \{l(m,j,k) : j = 1,\dots,J-1\} = \varnothing$.

5.3 Identification of the Instantaneous and Flow Payoffs

After identifying Q, and hence the CCPs σ, identification of the instantaneous and flow payoffs can proceed in a similar manner as in discrete time models. Let $V_i = (V_{i1},\dots,V_{iK})^\top$ and $u_i = (u_{i1},\dots,u_{iK})^\top$ denote the K-vectors of valuations and flow payoffs for player i in each state. Let $\psi_{ij} = (\psi_{ij1},\dots,\psi_{ijK})^\top$ denote the K-vector of instantaneous payoffs for player i making choice j in each state and let $\psi_i = (\psi_{i,1}^\top,\dots,\psi_{i,J-1}^\top)^\top$.

Using Proposition 2 and Assumption 4, for the baseline choice $j' = 0$ the differences $v_{ijk} \equiv \psi_{ijk} + V_{i,l(i,j,k)} - V_{ik}$ are identified from the CCPs and Q_0 for all players i, choices j, and states k. Letting $v_{ij} = (v_{ij1},\dots,v_{ijK})^\top$, we can stack equations across states k to write $v_{ij} = \psi_{ij} + (S_{ij} - I_K)V_i$, where S_{ij} is the $K \times K$ permutation matrix induced by the continuation state function $l(i,j,\cdot)$ and I_K is the $K \times K$ identity matrix. Finally, we can use the linear representation established in Proposition 6 to write V_i in terms of u_i and identified quantities as $V_i = \Xi_i^{-1}(u_i + \lambda E_i)$, where Ξ_i is the first matrix in square brackets in (8) and we have dropped the explicit dependence on σ and Q_0. The resulting system of K equations for player i and choice j is $v_{ij} = \psi_{ij} + (S_{ij} - I_K)\Xi_i^{-1}(u_i + \lambda E_i)$.

Define $\tilde{S}_{ij} = (S_{ij} - I_K)\Xi_i^{-1}$ and stack the equations for all choices $j = 1,\dots,J-1$ to obtain a system of equations for the payoffs ψ_i and u_i:

$$\begin{bmatrix} I_K & 0 & \dots & 0 & \tilde{S}_{i,1} \\ 0 & I_K & \dots & 0 & \tilde{S}_{i,2} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I_K & \tilde{S}_{i,J-1} \end{bmatrix} \begin{bmatrix} \psi_{i,1} \\ \vdots \\ \psi_{i,J-1} \\ u_i \end{bmatrix} = \begin{bmatrix} v_{i,1} - \lambda \tilde{S}_{i,1} E_i \\ v_{i,1} - \lambda \tilde{S}_{i,2} E_i \\ \vdots \\ v_{i,J-1} - \lambda \tilde{S}_{i,J-1} E_i \end{bmatrix},$$

or more simply $X_i \begin{bmatrix} \psi_i \\ u_i \end{bmatrix} = y_i$, where the matrix X_i and vector y_i are defined accordingly.

There are JK unknown payoffs but X_i only has $(J-1)K$ rows. We can complete the system with at least K additional linear restrictions on ψ_i and u_i of the form

$$R_i \begin{bmatrix} \psi_i \\ u_i \end{bmatrix} = r_i, \tag{12}$$

where R_i and r_i have at least K rows, each representing a restriction, and where R_i has JK columns, each corresponding to an element of ψ_i or u_i. Then the full system is

$$\begin{bmatrix} X_i \\ R_i \end{bmatrix} \begin{bmatrix} \psi_i \\ u_i \end{bmatrix} = \begin{bmatrix} y_i \\ r_i \end{bmatrix}$$

and if $\left[\begin{smallmatrix} X_i \\ R_i \end{smallmatrix}\right]$ has full column rank then ψ_i and u_i are identified.[16]

Proposition 8. *If Assumptions 1–7 hold, then for each player i the $(J-1)K \times JK$ matrix X_i has full rank. Furthermore, if for player i there exist restrictions on ψ_i and u_i as in (12) and if the matrix $\left[\begin{smallmatrix} X_i \\ R_i \end{smallmatrix}\right]$ has rank JK, then ψ_i and u_i are identified.*

Examples of appropriate restrictions include states where the continuation values are known, for example, if $u_{ik} = 0$ when a firm has permanently exited. This can be represented by the matrix $R_i = \left[\begin{smallmatrix} 0 & \dots & 0 & 1 & 0 & \dots & 0 \end{smallmatrix}\right]$ with $r_i = 0$, where the single nonzero element corresponds to u_{ik}. Payoff exclusion or exchangeability restrictions of the form $u_{ik} = u_{ik'}$ for $k' \neq k$ may also be used, for example, where k and k' are two states that differ only by a rival-specific state and are payoff-equivalent to firm i. This can be represented by a matrix $R_i = \left[\begin{smallmatrix} 0 & \dots & 0 & 1 & 0 & \dots & 0 & -1 & 0 & \dots & 0 \end{smallmatrix}\right]$ with $r_i = 0$, where the 1 and -1 elements correspond to u_{ik} and $u_{ik'}$ respectively. Finally, states where the instantaneous payoffs are the same can provide restrictions, for example, if entry costs or scrap values are constant across states implying $\psi_{ijk} - \psi_{ijk'} = 0$ for all i, some choice j, all states k and k'. Recall that in the single-agent renewal example, the replacement cost did not depend on the mileage state. These restrictions could be represented by the $(K-1) \times JK$ matrix

$$R_i = \begin{bmatrix} 1 & \cdots & 0 & -1 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & -1 & \vdots & \ddots & \vdots \\ 0 & \cdots & 1 & -1 & 0 & \cdots & 0 \end{bmatrix}$$

with $r_i = \begin{bmatrix} 0 & \cdots & 0 \end{bmatrix}^\top$. This alone yields $K-1$ linearly independent restrictions. In our empirical application in Section 7, we make use of restrictions of all three types mentioned

[16]Because payoffs in our model depend on rival actions only through the state, the number of required restrictions for a game with N players is only linear in N. In discrete time, simultaneous-move models the payoffs depend separately on the state and the actions of all N players, so the number of required restrictions is exponential in N (Pesendorfer and Schmidt-Dengler, 2008).

above: permanent exit, exchangeability, and constant entry costs and scrap values.

6 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.[17] CCP-based estimation methods, on the other hand, use two steps: CCPs are estimated in a first step and used to approximate the value function in a closed-form inversion or simulation step.[18] The approximate value function is then used in the likelihood function or the GMM criterion function to estimate the structural parameters.

Single-agent applications of our model may be estimated using full-solution methods, but in the rest of this section we focus on describing the two-step estimators. We begin with the simplest case in which continuous-time data is available. Next, 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. We conclude the section with a comparison of continuous time and discrete time methods.

6.1 Two-Step Estimation with Continuous Time Data

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.

Here we consider a dataset of observations $\{k_{mn}, t_{mn} : m = 1, \dots, M, n = 1, \dots, T_m\}$ sampled in continuous time over the interval $[0, \overline{T}]$ where k_{mn} is the state immediately prior to the n-th state change in market m and t_{mn} is the time of the state change.

Step 1: Estimating the Reduced-Form Hazards In Step 1, one estimates the hazards of actions and state changes nonparametrically. For example, these hazards can be estimated by maximum likelihood. Let $h_{ijk} = \lambda\sigma_{ijk}$ denote the hazard for an active player i choosing action j in state k and let

$$h = (q_{12}, q_{13}, \dots, q_{K-1,K}, \lambda\sigma_{111}, \dots, \lambda\sigma_{1Jk}, \dots \lambda\sigma_{N11}, \dots, \lambda\sigma_{NJK}) \tag{13}$$

[17]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.

[18]These methods were 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 (2008), Pakes et al. (2007), and Arcidiacono and Miller (2011).

denote the vector of distinct hazards for nature and all state-specific, non-continuation hazards of players. Let $\mathcal{H} \subset \mathbb{R}^{K(K-1)+N(J-1)K}$ denote the space of admissible vectors h.

In state k, the probability of the next state change occurring within τ units of time is

$$1 - \exp\left[-\tau\left(\sum_{l\neq k} q_{kl} + \sum_i \lambda \sum_{j\neq 0} \sigma_{ijk}\right)\right]. \tag{14}$$

This is the cumulative distribution function (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.

Differentiating with respect to τ yields the density for the time of the next state change, which is the exponential probability density function with the same rate parameter as before:

$$\left(\sum_{l\neq k} q_{kl} + \sum_i \lambda \sum_{j\neq 0} \sigma_{ijk}\right) \exp\left[-\tau\left(\sum_{l\neq k} q_{kl} + \sum_i \lambda \sum_{j\neq 0} \sigma_{ijk}\right)\right]. \tag{15}$$

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

$$\frac{\lambda\sigma_{ijk}}{\sum_{l\neq k} q_{kl} + \sum_i \lambda \sum_{j\neq 0} \sigma_{ijk}}. \tag{16}$$

Now, define g to be the exponential term from (14) and (15) restated as a function of h:

$$g(\tau, k; h) = \exp\left[-\tau\left(\sum_{l\neq k} q_{kl} + \sum_i \sum_{j\neq 0} h_{ijk}\right)\right]. \tag{17}$$

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 (15) and (16),

$$\lambda\sigma_{ijk}\, g(\tau, k; h),$$

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

$$q_{kl}\, g(\tau, k; h).$$

Now consider a continuous time sample of M markets. Define $\tau_{mn} \equiv t_{mn} - t_{m,n-1}$ to be the holding time between events. For the interval between the last event and the end of the sampling period, we define the final state and interval length as $k_{m,T+1} \equiv k_{m,T}$ and $\tau_{m,T+1} \equiv \overline{T} - t_{mT}$. Let $I_{mn}(i, j)$ be the indicator for whether the n-th move in market m

was a move by player i and the choice was j and let $I_{mn}(0,l)$ be the indicator for whether the n-th move in market m was a move by nature to state l. Then, the maximum likelihood estimates of the hazards h are

$$\hat{h} = \arg\max_{h \in \mathcal{H}} \left\{ \sum_{m=1}^{M} \sum_{n=1}^{T} \left[\ln g(\tau_{mn}, k_{mn}; h) + \sum_{l \neq k_{mn}} I_{mn}(0,l) \ln q_{kl} \right.\right.$$
$$\left.\left. + \sum_{i} \sum_{j \neq 0} I_{mn}(i,j) \ln h_{ijk} \right] + \ln g(\tau_{m,T+1}, k_{m,T+1}; h) \right\}.$$

The last term is the natural log 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.

Step 2: Estimating the Structural Payoff Parameters In Step 2, we use the estimated hazards from Step 1 to estimate the structural parameters $\theta \in \Theta$, which determine the flow payoffs u and instantaneous payoffs ψ. The main idea is to express the structural conditional choice-specific hazards as functions of θ and the estimated hazards $\hat{h}$ so that no fixed-point problem needs to be solved. Recall that, given the estimated hazards, we can estimate the choice probability for $j \neq 0$ as $\hat{h}_{ijk}/\lambda$ and $1 - \sum_{j \neq 0} \hat{h}_{ijk}/\lambda$ for $j = 0$.

Then, let $\Lambda : \Theta \times \mathcal{H} \to \mathcal{H} : (\theta, h) \mapsto \Lambda(\theta, h)$ denote the mapping by which, given vectors θ and h, a new vector $h' = \Lambda(\theta, h)$ is determined using the inverse CCP mapping. This proceeds in two steps. First, given θ and h, we construct Q and σ and obtain new value functions. In the second step, given the new value functions, we determine the new choice probabilities σ' and then form the new vector h'. Note that equilibrium conditional choice probabilities must satisfy $h = \Lambda(\theta, h)$.

Here, Λ can be either of two mappings. If the finite dependence property of Proposition 4 holds, Λ can be the mapping defined by the representation in (5). More generally, under Proposition 6, Λ can be the mapping defined by the linear representation of (8) for the infinite-horizon case. Although we only consider these two cases here, it is also possible to use other two-step approaches such as the forward-simulation-based estimators of Hotz et al. (1994) and Bajari et al. (2007).

In state k, 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

$$\Lambda_{ijk}(\theta, \hat{h})\, g(\tau, k; \Lambda(\theta, \hat{h})),$$

where $\Lambda_{ijk}(\theta, h)$ denotes the element of $\Lambda(\theta, h)$ corresponding to the hazard of player i

playing action j in state k. The second stage estimates are then

$$\hat{\theta} = \arg\max_{\theta} \left\{ \sum_{m=1}^{M} \sum_{n=1}^{T} \left[\ln g(\tau_{mn}, k_{mn}; \Lambda(\theta, \hat{h})) + \sum_{i} \sum_{j \neq 0} I_{mn}(i,j) \ln \Lambda_{ijk_{mn}}(\theta, \hat{h}) \right] \right.$$
$$\left. + \ln g(\tau_{m,T+1}, k_{m,T+1}; \Lambda(\theta, \hat{h})) \right\}.$$

This is a fairly standard two-step estimator and it is consistent and asymptotically normal under suitable regularity conditions. In the next section we consider two-step estimation in the leading case of discrete time data. There, we formally state sufficient conditions for consistency and asymptotic normality. The properties of the estimator above, with continuous time data, are largely similar.

6.2 Two-Step Estimation with Discrete Time Data

Often the exact sequence of events and event times are not observed, but rather the state is only observed at discrete points in time. Here, we consider estimation with a dataset of observations $\{k_{mn} : m = 1, \ldots, M, n = 0, \ldots, T\}$ which are sampled at times on the lattice $\{n\Delta : n = 0, \ldots, T\}$. We first show that it is still easy to carry out two-step estimation. We then briefly discuss some relevant computational details.

We now formalize our assumptions in order to define the two-step estimator with discrete time data and establish its large sample properties. These are standard regularity conditions requiring markets to be independent, the parameter space to be compact, the population parameter vector to be identified, and the hazard mapping Λ to be sufficiently smooth.

Assumption 8. Θ is compact and the true parameters θ^0 lie in the interior of Θ.

Assumption 9. For any $\theta \in \Theta$ with $\theta \neq \theta^0$ and any h such that $h = \Lambda(\theta, h)$, we have $h \neq h^0$ for the hazards h^0 implied by θ^0.

Assumption 10. $\Lambda : \Theta \times \mathcal{H} \to \mathcal{H} : (\theta, h) \mapsto \Lambda(\theta, h)$ is twice continuously differentiable.

Define the pseudo likelihood function

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

where $P_{k,l}(\Delta; h)$ denotes the (k, l) element of the transition matrix induced by h. Recall that for any $h \in \mathcal{H}$ a corresponding matrix $Q = \sum_{i=0}^{N} Q_i$ can be constructed.

Suppose we have a $\sqrt{M}$-consistent first stage M-estimator $\hat{h}$ for h^0. For example, the nonparametric maximum likelihood estimator of h based on $P(\Delta; h)$ is an estimator of this kind. We define the pseudo maximum likelihood estimator $\hat{\theta}$ of θ^0 as

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

Under the maintained assumptions, we show that $\hat{\theta}$ is consistent and asymptotically normal.

Proposition 9. *Suppose that Assumptions 1–10 hold and that $\hat{h}$ is an M-estimator of h^0 such that $\sqrt{M}(\hat{h} - h^0) \xrightarrow{\text{d}} \mathrm{N}(0, \Sigma)$. Then*

$$\sqrt{M}(\hat{\theta} - \theta^0) \xrightarrow{\text{d}} \mathrm{N}\left(0, \Omega_{\theta\theta^\top}^{-1} + \Omega_{\theta\theta^\top}^{-1} \Omega_{\theta h^\top} \Sigma \Omega_{\theta h^\top}^\top \Omega_{\theta\theta^\top}^{-1}\right)$$

where $\Omega_{\theta\theta^\top} = \mathrm{E}[\nabla_\theta\, s_m\, \nabla_{\theta^\top} s_m]$ and $\Omega_{\theta h^\top} = \mathrm{E}[\nabla_\theta\, s_m\, \nabla_{h^\top} s_m]$ and s_m is the pseudo-score

$$s_m \equiv \sum_{n=1}^{T} \ln P_{k_{m,n-1},k_{mn}} \left(\Delta; \Lambda(\theta^0, h^0)\right).$$

6.2.1 Computational Considerations

The matrix exponential can be computed using one of many known algorithms (cf. Moler and Loan, 1978, Sidje, 1998). When the Q matrix is large, this may seem to introduce a dimensionality problem 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 MT rows of $P(\Delta)$ need be calculated to estimate the model, corresponding to the number of observations.[19]

We now provide some intuition for why discrete-time data will not substantially complicate the problem, utilizing the uniformization procedure discussed in Section 5.1 under which we factor the overall jump process into a Poisson process with rate γ and an embedded Markov chain $\mathcal{Z}(q, \theta)$. The Markov chain depends on the rates of state changes for nature, q, and the structural parameters θ through the CCPs. 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

[19] 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 $P(\Delta)$.

length ($\Delta = 1$) satisfy

$$\left(\hat{q}, \hat{\theta}\right) = \arg\max_{(q,\theta)} \sum_{n=1}^{T} \ln \left[\sum_{r=0}^{\infty} \frac{\gamma^r \ \exp(-\gamma)}{r!} a_n^\top \mathcal{Z}(q,\theta)^r a_{n+1} \right]. \tag{18}$$

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 in (18) above also suggests a natural simulation-based estimator for $P(\Delta)$. Namely, we can use the expression inside the sum for the first $R < \infty$ terms and then 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.

6.3 Unobserved Heterogeneity

Our methods can also be extended to accommodate permanent unobserved heterogeneity using finite mixture distributions. In particular, suppose that T observations are sampled at intervals of length Δ for each of M markets, where each market is one of Z types. Let $\pi(z, k_{m1})$ denote the population probability of being type z conditional on the initial state.[20] We can then integrate with respect to the distribution of the unobserved state, so that the first-step maximum likelihood problem becomes

$$\left(\hat{h}, \hat{\pi}\right) = \arg\max_{(h,\pi)} \sum_{m=1}^{M} \ln \left[\sum_{z=1}^{Z} \pi(z, k_{m1}) \prod_{n=1}^{N} \ln P_{k_{m,n-1},k_{mn}}(\Delta; h, z) \right], \tag{19}$$

where $P(\Delta; h, z)$ is the transition matrix for type z as a function of the hazards, conditional on the observed and unobserved states. Both the hazards and type probabilities are estimated in the first stage using the EM algorithm as in Arcidiacono and Miller (2011).

Bayes' rule gives the probability of market m being in unobserved state z given the data. Denoting $\pi_m(z)$ as this conditional probability, it is defined as:

$$\pi_m(z) = \frac{\pi(z, k_{m1}) \prod_{n=1}^{T} P_{k_{m,n-1},k_{mn}}(\Delta; \hat{h}, z)}{\sum_{z'} \pi(z', k_{m1}) \prod_{n=1}^{T} P_{k_{m,n-1},k_{mn}}(\Delta; \hat{h}, z')}.$$

These probabilities are then used as weights in the second-step pseudo likelihood function

[20] By letting $\pi(z, k_{m1})$ depend on k_{m1}, we allow for an initial conditions problem.

to estimate the structural parameters:

$$\hat{\theta} = \arg\max_{\theta} \sum_{m=1}^{M} \sum_{z} \pi_m(z) \sum_{n=1}^{T} \ln P_{k_{m,n-1},k_{mn}}(\Delta; \Lambda(\theta, \hat{h}), z).$$

6.4 Comparison to Discrete Time Methods

Characterizing dynamic problems in continuous time has both advantages and disadvantages when compared to discrete time formulations.

From a computational perspective, the key benefit to working in continuous time concerns the treatment of counterfactuals. Here there are two advantages over discrete time. First, even if the data is time aggregated, the counterfactual analysis can leverage the full computational benefit of the underlying continuous time process. In particular, instantaneous representations of the value functions can be used to solve for counterfactual hazards, greatly reducing computational times because only one event occurs in any given instant. In the typical discrete time set up, all players move simultaneously, sharply increasing the computational burden. Second, one potential source of multiple equilibria—simultaneous moves—is eliminated in the continuous time context. Multiple equilibria make interpreting counterfactuals difficult due to ambiguity regarding which equilibrium would be played in the counterfactual environment.

On the other hand, estimation is more complicated in continuous time when data are time aggregated, as it requires integrating out over the possible paths between the observed state changes. It is important to keep in mind, however, that estimation is not the main bottleneck for research, as consistent estimates can be obtained with two-step estimation in either discrete or continuous time with low computational burden.

Continuous and discrete time representations of dynamic games also have different economic implications due primarily to the contrast between simultaneous and (stochastic) sequential moves. The applicability of each approach depends on the salient features of the economic setting.[21]

It is well-known, for example, that in a discrete time setting with asymmetric information, agents can sometimes make "mistakes" that arise from the simultaneity of choice (for example, simultaneously opening in the same location as a rival with a low ex-ante probability of entering because the rival happened to receive a "high" idiosyncratic shock). This is obviously most damaging in a static model (where firms cannot correct their mistakes) but may be unattractive in a dynamic setting as well if, for example, the choice requires

[21] The methods developed here apply to stationary environments. While dynamic games are typically estimated assuming stationarity, there are exceptions. See Beauchamp (2015) and Igami (2014) for examples of empirical games in non-stationary environments. Further work is necessary to understand whether the tools developed here can be extended to non-stationary, continuous time environments.

substantial investments to un-wind. That said, if such "mistakes" are a salient feature of the economic environment because, for example, firms only observe their rival's actions with a lag, a discrete time model might be able to capture this behavior more simply than a continuous time model.

More broadly, the regularity of decision timing implied by the discrete time model may be preferred in some settings. For example, there may be institutional features that restrict the timing of when decisions can be made. Board meetings may occur at pre-specified (e.g. quarterly) times during the year or long-term contracts (or implicit agreements) might prevent an agent from changing its actions for a pre-specified period of time. In such cases, a discrete time model may be more appropriate. On the other hand, a discrete time model may force many outcomes to occur at the same instant that are more naturally viewed as occurring sequentially (e.g. the opening or closing of stores, changes in installed capacity, or the release of new products). Generally, the appropriate choice of modeling assumption will depend on the institutional setting being considered.

7 Wal-Mart's Entry into the Supermarket Industry

Our empirical application considers the impact of Wal-Mart's entry into the supermarket industry. In 1994, the first year for which we have data, Wal-Mart owned 97 supercenter outlets. However, by 2006, the last year of our data, they operated 2225 such outlets and ranked first among all grocery firms in terms of overall sales. Much of this expansion came at the expense of incumbent grocers. The question is exactly which types of firms were most impacted and how the competitive landscape evolved in response.

Wal-Mart first gained national prominence through its gradual rise to an ultimately dominant position in the discount store (general merchandise) industry. Due to its large role in the U.S. economy (Wal-Mart accounted for 8.8 percent of (non-automobile) retail sales in 2004), Wal-Mart has attracted significant attention from both the popular press and academic researchers. Much of the debate centers on Wal-Mart's overall impact on consumer surplus, labor market outcomes, and local competitors. Our focus here will be on its impact on competition and market structure.

There is no doubt that Wal-Mart has had a significant impact on retail competition. Due to its scale and operational efficiency, Wal-Mart is often able to undercut the prices of its rivals, both in general merchandise and groceries. For example, in the context of groceries, Basker and Noel (2009) find that Wal-Mart is able to set prices that are on average 10 percent lower than their competitors and that this differential appears to be increasing over time. They also find evidence of a competitive response: in the short run, competing grocery stores reduce prices by 1-2% when a Wal-Mart enters. They note that the response is mostly due to smaller scale competitors—the reaction by the top 3 national chains is only half as large. Matsa (2011) looks at Wal-Mart's impact on supermarket stock-outs (a measure of quality) and finds that entry by Wal-Mart decreases stock-outs at competing supermarkets, but that this impact is instead centered on large-scale competitors. The smaller rivals either cut prices or exit. Since much of Wal-Mart's advantage appears to derive from its enormous scale and intensive investment in information technology (Basker, 2007), there is particular concern that Wal-Mart stifles small scale entrepreneurial activity. This was particularly salient in the context of the discount store industry (Wal-Mart's original line of businesss) which was, prior to entry by Wal-Mart, served by a collection of single-store outlets that typically focused on a more narrow line of goods. By offering greater breadth and depth of assortment, Wal-Mart consistently leveraged its greater scale to undercut prices and consolidate purchases (by offering one stop shopping). The impact on small scale firms was unambiguously negative—Wal-Mart displaced the "mom and pop" stores. Jia (2008) concludes that Wal-Mart's expansion alone drove 50-70 percent of the net exit of small discount retailers from the late 1980s to mid 1990s. Focusing on the big box industry as a whole, Haltiwanger, Jarmin, and Krizan (2010) find a large negative impact of big box chains (including Wal-Mart) on single store retailers and small chains, continuing a long term trend toward larger chains throughout much of retail.

However, Wal-Mart's impact on the structure of the grocery industry is less clear. While Matsa (2011) finds a significant negative impact on the survival probabilities of small-scale rivals consistent with earlier experience in the discount industry (and a shift up market by the larger chains), Ellickson and Grieco (2013) find a large, geographically localized negative impact of Wal-Mart centered on the large grocery chains, with no measurable impact on the small firms at all. In the analysis presented below, we find that Wal-Mart's negative impact falls almost entirely on the large chains and is actually associated with an *expansion* by the single store segment. This sharp contrast with the earlier case of the discount segment is striking, and illustrates the benefit of a dynamic structural model as the overall shift is an equilibrium result that evolves slowly over time.

7.1 Data

Our data for the supermarket industry are drawn 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 and a variety of in-store features are linked to the firm level through a firm identity code, which can also be used to identify the location of the nearest distribution facility. In addition to the Trade Dimensions data, which consists of yearly snapshots of the entire industry, we also have information on the exact opening dates of the Wal-Mart supercenters that were gathered from a variety of online sources.[22]

For market definition, we focus on Metropolitan Statistical Areas (MSAs) with population under 500,000, yielding a total of 309 markets. For our purposes, 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. We allow for up to seven chain players in each MSA who may or may not be active in the market at any given time. 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, we allow for ten potential fringe firm entrants in each MSA, so the number of fringe firms is the number of incumbent fringe firms plus ten.

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.[23] 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. In particular, the growth category of a city is fast if the annual growth rate is greater than 2% (74 cities), moderate if the annual growth rate is between 1% and 2% (106 cities), and slow if the annual growth rate is less than 1% (129 cities). The parameters governing population transitions are indexed by these growth categories.

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 Wal-Marts 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, Wal-Mart, and fringe firms, respectively. The corresponding figures for store closings are 0.224, 0.002, and 0.908, revealing that Wal-Mart virtually never exits during our sample period.

[22] We use data on the exact date of Wal-Mart entry for a robustness specification below to examine how well our estimator does when it is based only on time aggregated data. Data on the exact date of entry comes from two sources. The first is trendresults.com, which provides the opening date for all Wal-Mart stores. Some supercenters, however, entered as a result of a conversion from a discount store to a supercenter and this dataset misses these conversion dates. We supplement these data with data from Emek Basker, who collected entry dates directly from Wal-Mart's website. Unfortunately, some of the dates in this dataset were clustered around particular days, suggesting the possibility that the exact entry date was approximated in some cases. For the robustness analysis below, we only used exact Wal-Mart entry dates if we were confident in the timing. This occurred for 256 of the 475 Wal-Mart entries.

[23] The discretization was such that differences in log population between adjacent categories was equal.

Table 2 looks at entry and exit decisions for chain firms and fringe firms one year before, the year of, and the year after initial entry by Wal-Mart. Here, we see that chain and fringe firms both respond negatively to Wal-Mart. The number of new chain stores falls from 0.311 in the period before Wal-Mart enters to 0.189 in the period after—a 40% drop. Similarly, the number of stores that close 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 Wal-Mart's presence is more detrimental to chain firms than fringe firms.[24]

7.2 Model

To quantify Wal-Mart's impact on large versus small rivals, and allow for heterogeneous competitive effects across firm types, there are three types of firms in our model: chain firms (who can operate many stores), Wal-Mart (who can also operate many stores), and fringe firms (who can operate at most one store each). We assume that the chain firms (including Wal-Mart) make strategic decisions within each local market (MSA), but independent decisions across markets (i.e. they do not choose their entire spatial layout jointly, but rather make optimal decisions on a market by market basis). 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, normalized at $\lambda = 1$, and q_1 and q_{-1} are the rates of moving up and down in population, respectively.

Our model is a continuous-time, discrete action version of the dynamic oligopoly model of Ericson and Pakes (1995) and Pakes and McGuire (1994), in which heterogeneous firms make entry, exit, and investment decisions. Firms in our model are differentiated by type (Wal-Mart, chain, or fringe) and by the number of stores they operate. Firms invest by building new stores and disinvest by closing stores.

Since the state variables are discrete, we enumerate all possible states by an integer scalar index $k = 1, \ldots, K$. The state of the market at each instant can be summarized by

[24] This table also highlights an advantage of using a model where the frequency of moves can differ from the sampling frequency of the data. Note that the numbers of entering and exiting chain stores in the year of Wal-Mart's initial entry are bracketed by the corresponding values the year before and the year after Wal-Mart's entry. In markets where chain and fringe stores saw little change in their building patterns, this suggests that Wal-Mart entered later in the period. In contrast, when Wal-Mart enters early in the period, exit by chain and fringe stores is more likely to occur within the period.

a vector x_k containing the number of stores operated by each chain firm, $s^{\mathrm{c}}_{1k}, s^{\mathrm{c}}_{2k}, \ldots$, each fringe firm, $s^{\mathrm{f}}_{1k}, s^{\mathrm{f}}_{2k}, \ldots$, and Wal-Mart, s^{w}_{k}, along with the current population, d_k:

$$x_k = (s^{\mathrm{c}}_{1k}, s^{\mathrm{c}}_{2k}, \ldots, s^{\mathrm{f}}_{1k}, s^{\mathrm{f}}_{2k}, \ldots, s^{\mathrm{w}}_{k}, d_k).$$

Therefore, each value of k represents an encoded state vector and the function $l(i, j, k)$ gives the state conditional on firm i taking action j in state k. Additionally, each market is characterized by a time-invariant unobserved type z, which is observed by the firms in the market but not by the econometrician.[25] Hence, the full state vector at any instant can be written as (x_k, z).

7.2.1 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 given by:

$$V_{ik} = \frac{u_{ik} + \sum_{j\in\{-1,1\}} q_j V_{i,l(0,j,k)} + \sum_{m\neq i} \lambda \sum_{j\in\{-1,1\}} \sigma_{mjk} V_{i,l(m,j,k)} + \lambda \,\mathrm{E} \max_j \left\{ V_{i,l(i,j,k)} + \psi_{ijk} + \varepsilon_{ijk} \right\}}{\rho + \sum_{j\in\{-1,1\}} q_j + \sum_{m\neq i} \lambda \sum_{j\in\{-1,1\}} \sigma_{mjk} + \lambda}. \tag{20}$$

In this expression, nature is indexed by $i = 0$, the choices are $j = 1$ and $j = -1$, and the costs ψ_{ijk} reflect the costs of initial entry, building new stores, or closing stores depending on the values of the player identity i, choice j, and state k. We specify these costs for each firm type below.

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 (in effect, the continuation value for exit is identically zero). 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 (20) using Propositions 2 and 3. As a result (and exploiting Proposition 4), the value function on the left-hand side of (20) can be expressed as a function of the flow payoffs, the move arrival parameters, and the probabilities of making particular decisions. Because Wal-Mart essentially never closes stores or exits markets, applying our finite dependence representation to recover to their parameters would be inappropriate.[26] Nonetheless, we are able to fully account for its strategic impact on rivals' actions by using our first stage estimates to capture its rivals' beliefs about Wal-Mart's equilibrium policy functions. This flexibility is an additional advantage of two-step

[25] We use $Z = 5$ points of support, $z \in \{-1.3998, -0.5319, 0.0, 0.5319, 1.3998\}$, based on a discrete approximation to a standard normal random variable.

[26] We could, however, have used the representation in Proposition 6 to recover these parameters. We do not do so because they are not relevant to the policy simulations we consider.

estimation.

7.2.2 Flow Profits and Choice-Specific Payoffs for Chain Firms

We specify the flow payoff u_{ik} for chain firms in terms of per-store latent revenue and total cost. These are linear functions of market population, d_k, the number of own stores, s^{c}_{ik}, the number of competing chain stores, $\tilde{s}^{\mathrm{c}}_{ik}$, the number of Wal-Mart stores, s^{w}_k (Wal-Mart), and the number of fringe stores, s^{f}_k. Revenues also depend on an unobserved (to the econometrician) characteristic of the market, z, which reflects the tastes of consumers in a given market for particular types of products. Flow profits for a chain firm i in state k are

$$u^{\mathrm{c}}_{ik} = s^{\mathrm{c}}_{ik}\left(\beta^{\mathrm{c}}_0 + \beta^{\mathrm{c}}_1 \tilde{s}^{\mathrm{c}}_{ik} + \beta^{\mathrm{c}}_2 s^{\mathrm{w}}_k + \beta^{\mathrm{c}}_3 s^{\mathrm{f}}_k + \beta^{\mathrm{c}}_4 s^{\mathrm{c}}_{ik} + \beta^{\mathrm{c}}_5 d_k + \beta^{\mathrm{c}}_6 z + \beta^{\mathrm{c}}_7 z s^{\mathrm{c}}_{ik}\right) + e^{\mathrm{c}}_{ik},$$

where $e_{\mathrm{c}ik}$ is (the negative of) the flow cost of operating a set of stores:

$$e^{\mathrm{c}}_{ik} = \mu^{\mathrm{c}}_1 s^{\mathrm{c}}_{ik} + \mu^{\mathrm{c}}_2 (s^{\mathrm{c}}_{ik})^2 + \mu^{\mathrm{c}}_3 (s^{\mathrm{c}}_{ik})^3.$$

A cubic cost function allows there to be regions of increasing and then decreasing returns to scale, ensuring that for each state the optimal value of s^{c}_{ik} is finite. Collecting terms yields

$$\begin{aligned} u^{\mathrm{c}}_{ik} &= s^{\mathrm{c}}_{ik}\left((\beta^{\mathrm{c}}_0 + \mu^{\mathrm{c}}_1) + \beta^{\mathrm{c}}_1 \tilde{s}^{\mathrm{c}}_{ik} + \beta^{\mathrm{c}}_2 s^{\mathrm{w}}_k + \beta^{\mathrm{c}}_3 s^{\mathrm{f}}_k + (\beta^{\mathrm{c}}_4 + \mu^{\mathrm{c}}_2) s^{\mathrm{c}}_{ik} + \mu^{\mathrm{c}}_3 (s^{\mathrm{c}}_{ik})^2 + \beta^{\mathrm{c}}_5 d_k + \beta^{\mathrm{c}}_6 z + \beta^{\mathrm{c}}_7 z s^{\mathrm{c}}_{ik}\right) \\ &= s^{\mathrm{c}}_{ik}\left(\theta^{\mathrm{c}}_0 + \theta^{\mathrm{c}}_1 \tilde{s}^{\mathrm{c}}_{ik} + \theta^{\mathrm{c}}_2 s^{\mathrm{w}}_k + \theta^{\mathrm{c}}_3 s^{\mathrm{f}}_k + \theta^{\mathrm{c}}_4 s^{\mathrm{c}}_{ik} + \theta^{\mathrm{c}}_5 (s^{\mathrm{c}}_{ik})^2 + \theta^{\mathrm{c}}_6 d_k + \theta^{\mathrm{c}}_7 z + \theta^{\mathrm{c}}_8 z s^{\mathrm{c}}_{ik}\right). \end{aligned}$$

The choice-specific instantaneous payoffs ψ_{ijk} depend on the unobserved state z and differ according to whether firm i is an incumbent ($s^{\mathrm{c}}_{ik} > 0$) or new entrant ($s^{\mathrm{c}}_{ik} = 0$) and whether the choice is building a new store ($j = 1$) or closing an existing store ($j = -1$):

$$\psi_{ijk} = \begin{cases} \eta^{\mathrm{c}}_0 + \eta^{\mathrm{c}}_1 z + \kappa^{\mathrm{c}}_0 + \kappa^{\mathrm{c}}_1 z & \text{if } s^{\mathrm{c}}_{ik} = 0 \text{ and } j = 1, \\ \kappa^{\mathrm{c}}_0 + \kappa^{\mathrm{c}}_1 z & \text{if } s^{\mathrm{c}}_{ik} > 0 \text{ and } j = 1, \\ \phi^{\mathrm{c}}_0 + \phi^{\mathrm{c}}_1 z & \text{if } s^{\mathrm{c}}_{ik} > 0 \text{ and } j = -1, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, the structural parameters of interest for chain firms are the coefficients of the per-store payoff function and the parameters of the instantaneous payoffs:

$$\theta^{\mathrm{c}} = (\theta^{\mathrm{c}}_0, \ldots, \theta^{\mathrm{c}}_8, \eta^{\mathrm{c}}_0, \eta^{\mathrm{c}}_1, \kappa^{\mathrm{c}}_0, \kappa^{\mathrm{c}}_1, \phi^{\mathrm{c}}_0, \phi^{\mathrm{c}}_1).$$

We assume that there are seven chain firms in all markets with the number of potential entrants in a market equal to seven minus the observed number of chain firms in the market.

The two-step estimation we employ means we do not need to place any restrictions on the number of chain stores a firm operates, though we will need to do so when we solve for counterfactual choice probabilities.

7.2.3 Flow Profits and Choice-Specific Payoffs for Fringe Firms

Flow profits for fringe stores have a similar linear form to that of chain firms, though with different coefficients and a different flow cost function, e^{f}_{ik}. Namely, an operating fringe store has profits given by:[27]

$$u^{\mathrm{f}}_{ik} = \beta^{\mathrm{f}}_0 + \beta^{\mathrm{f}}_1 s^{\mathrm{c}}_k + \beta^{\mathrm{f}}_2 s^{\mathrm{w}}_k + \beta^{\mathrm{f}}_3 s^{\mathrm{f}}_k + \beta^{\mathrm{f}}_4 d_k + \beta^{\mathrm{f}}_5 z + \beta^{\mathrm{f}}_6 z s^{\mathrm{f}}_k + e^{\mathrm{f}}_{ik},$$

Fringe competitors often rely on the same suppliers to deliver goods to their stores. Hence, there may be some density economies 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:

$$e^{\mathrm{f}}_{ik} = \mu^{\mathrm{f}}_0 + \mu^{\mathrm{f}}_1 s^{\mathrm{f}}_k + \mu^{\mathrm{f}}_2 (s^{\mathrm{f}}_k)^2.$$

Collecting terms yields the flow profit function

$$\begin{aligned} u^{\mathrm{f}}_{ik} &= (\beta^{\mathrm{f}}_0 + \mu^{\mathrm{f}}_0) + \beta^{\mathrm{f}}_1 s^{\mathrm{c}}_k + \beta^{\mathrm{f}}_2 s^{\mathrm{w}}_k + (\beta^{\mathrm{f}}_3 + \mu^{\mathrm{f}}_1) s^{\mathrm{f}}_k + \mu^{\mathrm{f}}_2 (s^{\mathrm{f}}_k)^2 + \beta^{\mathrm{f}}_4 d_k + \beta^{\mathrm{f}}_5 z + \beta^{\mathrm{f}}_6 z s^{\mathrm{f}}_k \\ &= \theta^{\mathrm{f}}_0 + \theta^{\mathrm{f}}_1 s^{\mathrm{c}}_k + \theta^{\mathrm{f}}_2 s^{\mathrm{w}}_k + \theta^{\mathrm{f}}_3 s^{\mathrm{f}}_k + \theta^{\mathrm{f}}_4 (s^{\mathrm{f}}_k)^2 + \theta^{\mathrm{f}}_5 d_k + \theta^{\mathrm{f}}_6 z + \theta^{\mathrm{f}}_7 z s^{\mathrm{f}}_k. \end{aligned}$$

Recall that fringe firms can operate at most one store. Therefore, potential fringe entrants only choose whether to enter (and build a single store) or not and incumbent fringe firms only decide whether to close their store or not. We assume there are always ten potential fringe entrants. Therefore, the choice-specific instantaneous payoffs for fringe firms, ψ_{ijk}, represent entry costs for new entrants (for which $s^{\mathrm{f}}_{ik} = 0$) and exit values for incumbents (for which $s^{\mathrm{f}}_{ik} = 1$). Since fringe firms can operate at most one store, we cannot distinguish the entry cost from the building cost. Because we estimate a fixed flow cost parameter for fringe firms, θ^{f}_0, we normalize the exit value to zero. As with chain firms, we allow the entry cost to depend on the unobserved state z:

$$\psi_{ijk} = \begin{cases} \eta^{\mathrm{f}}_0 + \eta^{\mathrm{f}}_1 z & \text{if } s^{\mathrm{f}}_{ik} = 0 \text{ and } j = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, the structural parameters of interest for fringe firms are the coefficients of the

[27] Note that there is no term for own number of stores here as fringe stores can only operate one store.

payoff function and the parameters of the instantaneous payoffs:

$$\theta^{\mathrm{f}} = (\theta_0^{\mathrm{f}}, \ldots, \theta_7^{\mathrm{f}}, \eta_0^{\mathrm{c}}, \eta_1^{\mathrm{c}}).$$

7.3 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 the change in population over time, and then estimating the structural parameters of the profit functions, taking the reduced form hazards as given.

7.3.1 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, Wal-Mart, and fringe).[28] In particular, let $h(\alpha)$ denote the vector of reduced form hazards as in (13) but written as a function of a parameter vector α. Let $\tilde{\sigma}_{ij}(k, z, \alpha)$ denote the reduced form probability of firm i making choice j in state (k, z), which has the form

$$\tilde{\sigma}_{ij}(k, z, \alpha) = \frac{\exp(\phi_j(k, z, \alpha))}{\sum_{j' \in \mathcal{A}_{ik}} \exp(\phi_{j'}(k, z, \alpha))},$$

where $\phi_j(k, z, \alpha)$ is a flexible function of the state variables. The hazards are

$$h(\alpha) = (q_{-1}, q_1, \lambda\tilde{\sigma}_{11}(1, 1, \alpha), \ldots, \lambda\tilde{\sigma}_{NJ}(K, Z, \alpha))$$

and the likelihood of firm i making choice j in state k in a market with unobserved state z with an interval of length τ since the previous event, is

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

Since we have annual data, we simulate potential 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. For each period we draw R simulated paths, randomly assigning each

[28] The 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 Wal-Marts and its square, 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.

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.

For simplicity, we focus on a single observation n in market m where, over an interval of length $\Delta = 1$, the starting and ending states are $\underline{k}$ and $\overline{k}$ and W events were known to occur in between. Let $k_w^{(r)}$ denote the state immediately preceding event w in simulation r, with $w = 1, \ldots, W+1$. Since we observe the states at the beginning and end, all simulated paths $r = 1, \ldots, R$ are such that $k_1^{(r)} = \underline{k}$ and $k_{W+1}^{(r)} = \overline{k}$. 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 and let $t_w^{(r)}$ and $\tau_w^{(r)}$ be the absolute time and holding time of simulated event w. Conditional on knowing the unobserved state z, the simulated likelihood for the single observation n in market m is

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

Since z is unobserved, we estimate the reduced form hazards using mixture distributions. Higher values of the unobserved state may make it easier or harder to operate as a chain, fringe, or Wal-Mart store respectively. We discretize the standard normal distribution into five points and then estimate the population probabilities of being at each of these points.[29]

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 T periods each, integrating with respect to the distribution of the unobserved market states yields

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

We obtain $(\tilde{\alpha}, \tilde{\pi})$ using the EM algorithm following Arcidiacono and Miller (2011). These 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 market being in each of the unobserved states.

[29]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 Wal-Marts, and the number of fringe stores, all interacted with population, and the city growth type.

7.3.2 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 firms and fringe stores. As noted above, there is not enough observed variation in Wal-Mart's actions to feasibly estimate their structural parameters using our finite dependence representation. However, their strategic impact is captured by the first stage estimates, which reflect their rivals' beliefs over Wal-Mart's expected actions. Our counterfactual experiments will then include only scenarios in which Wal-Mart's structural parameters do not play a role (e.g., equilibria in which Wal-Mart no longer exists). Let $\pi_m(z)$ denote the probability of MSA m being in unobserved state z given the data. Using Bayes' rule, we have

$$\pi_m(z) = \frac{\pi(z, k_{m1}) \prod_{n=1}^{T} \tilde{L}_{mn}(h(\tilde{\alpha}); z)}{\sum_{z'} \pi(z', k_{m1}) \prod_{n=1}^{T} \tilde{L}_{mn}(h(\tilde{\alpha}); z')}. \tag{21}$$

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

Next, using Proposition 4 we express the value function in (20) as a function of the structural parameters, θ, and the reduced form hazards from the first stage, $h(\tilde{\alpha})$. Let Λ denote the mapping which yields the new implied hazards as a function of θ and h. In practice we will use $h = h(\tilde{\alpha})$, based on the first stage estimates, and then estimate θ. Let $\tilde{\sigma}_{ij}(k, z, \theta) = \Lambda_{ijkz}(\theta, h(\tilde{\alpha}))/\lambda$ denote the rescaled element of $\Lambda(\theta, h(\tilde{\alpha}))$ that corresponds to the implied probability of firm i taking action j in state (k, z). The second-step pseudo-likelihood function used to estimate θ can then be written as $\check{L}_{mn}(\theta; \tilde{\alpha}, z) \equiv \tilde{L}_{mn}(\Lambda(\theta, h(\tilde{\alpha})); z)$. We use the same simulation draws as in Step 1 but replace each $\tilde{\sigma}_{ij}$ with the corresponding structural probability which is a function of θ. The second stage estimates are then given by

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

7.4 Results

The structural parameter estimates for chain stores are presented in Table 3. In the first set of columns we present results in which all entry and exit decisions are aggregated over the year. In the second set we use the information on the exact date of Wal-Mart entry, while the third set removes the controls for unobserved heterogeneity. In all cases, we calculate standard errors using the approach of Ackerberg, Chen, and Hahn (2012).

All three sets of estimates show Wal-Mart having substantial effects on chain flow profits that dwarf the effects of other chain and fringe stores. Despite these large effects, the estimates aggregating over Wal-Mart entry times (column 1) and using the exact date of Wal-Mart entry (column 2) show virtually identical parameter estimates, suggesting that, in this case, integrating out over entry times does not contaminate the estimates.

The effects of competition on chain profits, however, are substantially lower in the third set of columns that do not control for unobserved heterogeneity. The coefficients on the number of Wal-Marts and number of fringe stores are over thirty percent higher in columns 1 and 2 than in column 3, with the effect of the other chain stores almost four times as high. This is to be expected since, all else equal, higher unobserved demand will be correlated with more entry leading to estimates of competition that are biased downward. Controlling for such bias is important for our counterfactual analysis since the degree to which different firm types face differential competitive pressures from each type of rival will determine who thrives and who fails as the market evolves.

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 the three specifications for fringe firms are presented in Table 4. We again see negative effects of Wal-Mart on fringe profits in all three specifications with the time-aggregated results (column 1) very close to those that condition on Wal-Mart's exact entry times (column 2). In contrast to the chain stores, Wal-Mart's effects on fringe stores are smaller when we account for unobserved heterogeneity. An additional chain store negatively affects profits of fringe stores, with the effect being about half that of an additional Wal-Mart. By comparison, the effect of an additional chain store from a competing chain is a little over one-sixth of the effect of an additional Wal-Mart for chain profits. Hence, Wal-Mart appears to have a greater relative effect on chain stores than fringe stores. Moreover, the impact of competition from both types of rivals is smaller for the fringe than for the chains, suggesting that fringe stores are indeed more differentiated in product space than their chain rivals. Since differentiation in this industry mainly involves focusing on more narrow segments of the consumer base (e.g. ethnic foods, organic meats and produce) it makes sense that the benefits of scale would be more muted here, creating an opening for smaller scale firms to counter the cost advantage enjoyed by the larger chains.

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 then dominating as the number of fringe competitors increases. This positive spillover is robust to the inclusion of correlated unobservables. 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.

7.5 Counterfactuals

The focus of our empirical analysis is on the differential impact of Wal-Mart on the chain and fringe segments. While Wal-Mart's earlier impact on the discount sector was unambiguously detrimental to small-scale rivals, the impact on the grocery industry is much less clear. To evaluate the impact of Wal-Mart's entry on the market structure of the grocery industry, we conducted counterfactual experiments for each of the 205 markets that did not have a Wal-Mart outlet in the beginning of our sample. In particular, we computed equilibrium policy functions for a counterfactual scenario in which Wal-Mart does not exist and compared the temporal evolution of these markets under this counterfactual to the evolution implied by the estimated first stage policy functions recovered from the true data.[30] Using these two sets of policy functions (true and counterfactual) we then simulated 10,000 future histories from the first period in the data (1994), and averaged over them (at the market level) to characterize differences in long-run equilibrium outcomes.

Table 5 illustrates the long-run impact of Wal-Mart's entry. The two panels of Table 5 contain the simulated equilibrium outcomes at year 20 (which corresponds to 2014 in calendar time) and includes several measures of market structure, including the average number of chain firms, chain stores, fringe stores, Wal-Mart stores, the average market shares of each of these three player types, the share (by store count) of the largest (C1) and three largest (C3) firms, and the Herfindahl-Hirschman Index (HHI) computed using the share of total square feet of selling space.[31] Column 1 contains the average initial market population, while column 2 contains the number of markets in a given category. The first row of each panel averages across all 205 markets, while the next four rows in each panel break this average out by census region.

Several clear patterns emerge. First, Wal-Mart has a sizable negative impact on both the number of chain firms and number of chain stores that are active in a given market. Looking across regions, we find that the impact is fairly consistent across the Midwest, Northeast and South, but only about half as large in the West. This is consistent with both the higher level of political resistance that Wal-Mart has faced in these markets, and the fact that the more tightly clustered population centers in these markets are less suited to Wal-Mart's more diffuse outlet structure.

Second, Wal-Mart's presence actually leads to an expansion of the fringe. Note that

[30] There is no guarantee that our counterfactual policy functions are unique. The lower computational burden of continuous time, however, makes it easier to search for multiple equilibria, while also eliminating simultaneity as a likely cause of multiplicity. Using 10,000 highly-varied draws on starting values for the value functions and iterating to a fixed point showed no evidence of multiple equilibria in our counterfactual simulations.

[31] Our square footage calculation assumes that all Wal-Mart stores are 62,200 square feet, all chain stores are 35,500 square feet and all fringe stores are 13,500 feet, which correspond to the empirical averages from the data.

this is in sharp contrast to what occurred in the discount store industry, where the small "mom and pop" stores retreated in the face of Wal-Mart's expansion. There are at least two likely reasons for this contrast. First, unlike the earlier experience of rural discount (general merchandise) stores, Wal-Mart faced a large number of well-established chain stores in the supermarket industry that had already made similar investments in scale and IT and were providing a range of products that overlapped very closely with Wal-Mart's offerings. Second, and consistent with this overlap, as captured in the first stage policy function estimates, Wal-Mart competes more directly with chain firms than those in the fringe. As noted above, firms in the fringe are much more likely to be horizontally differentiated into a distinct local niche (such as providing an ethnic or gourmet focus) than the full-service chains. As Wal-Mart displaces the chains, this likely provides an even greater opportunity for differentiation by the fringe segment, while Wal-Mart's huge cost advantage (reflected in its large competitive impact on chain flow profits) represents a direct challenge to the survival of competing chains.

The overall (and regional) impact of Wal-Mart actually leads to a decrease in market concentration along all three measures (C1, C3 and HHI). This is driven by the displacement of the chains by firms from the fringe. Even though Wal-Mart eventually becomes a large player in many of these markets, the decrease in the number of chain firms is significantly smaller than the decrease in the number of chain stores, leading to a more uniform market structure when Wal-Mart is present. While we do not have information on prices, under most models of retail competition a more uniform structure would yield tougher price competition amongst the remaining firms.

Table 6 cuts the counterfactual results along several additional dimensions. The first two columns contain the number of markets in each category and the average number of Wal-Mart's predicted to enter under the "actual" scenario. The remaining columns report the percentage change in the various market structure measures employed in Table 5. The first panel breaks the results out by census region (as in Table 5), illustrating the importance of accounting for unobserved heterogeneity in capturing regional variation (we will return to this point shortly). The next panel breaks the results out by market size. We label markets as either "small" or "large" based on whether their initial population is below or above the median value. Here we see a sharp contrast: while the negative impact on chains is the same in both small and large markets, the positive impact on the fringe is much stronger in the larger markets. This is consistent with the notion that the fringe is exploiting a greater opportunity for horizontal differentiation, as these opportunities would naturally increase with market size (i.e., a larger population yields more sub-markets large enough to support one or more stores tailored to more unique tastes). Similarly, we divide markets by "Slow", "Moderate", or "Fast" growth rates based on whether they are in the lower, middle, or upper tercile of the population growth distribution.

The last panel breaks markets out by their "unobserved type" as assigned by the finite mixture model employed above.[32] While there were five points of support in the estimation, no markets were pre-dominantly assigned to the highest type. Higher values of the unobserved type are clearly associated with greater Wal-Mart entry and higher initial population levels. Viewed geographically (results not shown), markets with low values of the unobserved state tend to be in the western and mountain states as well as New England, while markets with high values of the unobserved state tend to be in the south and southern Atlantic states. Again, this is consistent with Wal-Mart's center of regional strength and areas where opposition to its expansion is weakest. The highest values of the unobserved state are also associated with a greater initial fringe presence.

Strikingly, in markets assigned the lowest value for the unobserved type, Wal-Mart has a strong negative impact on both chain and fringe stores. As a result, concentration actually increases with Wal-Mart's presence, as it moves towards being the dominant firm in the market—an outcome that closely matches what happened in the discount industry. In the case of the supermarket industry, however, this outcome represents only a very small fraction of the overall set of markets (9 out of 205). Geographically, these are small, rural markets that initially had few chains and were essentially dominated by fringe players. Wal-Mart pushes both types of firms out and actually increases the equilibrium level of concentration. It seems likely that these markets are too small to support a diverse range of offerings and Wal-Mart's scale advantage dominates.

However, this is far from the modal outcome across the whole industry. For all other values of the unobserved state, the average impact on the fringe is positive and sharply increasing with the value of the unobserved state. The overall structure becomes less concentrated with Wal-Mart's presence. This is due to both the contraction of the chain segment, which is hurt more the greater the value of the unobserved state, and the fact that the fringe competes less directly with (i.e. is more differentiated from) Wal-Mart.

Table 7 considers the impact of ignoring heterogeneity, and presents the same sets of conditional means employed in Table 6, but using parameters (and counterfactual computations) estimated without accounting for unobserved heterogeneity. Perhaps not surprisingly, these results miss the rich heterogeneity apparent in Table 6: the regional variation is muted and the asymmetric impact in large and small markets vanishes. The last panel breaks things out using the values of the unobserved state recovered earlier (which are ignored in this specification). Note that we now lose the traditional story altogether, as the small set of rural "mom and pop" markets are now averaged in with the rest. Most importantly, the

[32] The probability $\pi_m(z)$ that market m is of unobserved type z is defined in (21). We take the mean state for a market to be $\overline{z}_m = \sum_z z\pi_m(z)$ and round $\overline{z}_m$ to the nearest mass point to label cities, respectively, as "More Negative", "Negative", "Zero", "Positive", or "More Positive".

offsetting impact of Wal-Mart's expansion, namely its complementary relationship with the fringe, is sharply reduced as we no longer capture its competitive separation from Wal-Mart and the interplay with market size.

Table 8 explores the temporal evolution of market structure, providing some insight into why the expansion of the fringe did not show up in earlier studies and illustrating the importance of a dynamic structural model in this context. Focusing on the small market versus large market comparison, the table shows the equilibrium market structure in years 5, 10, 15 and 20. As in year 20, the expansion of the fringe is most pronounced in earlier years in the larger markets where the scope for differentiation is greatest. However, across both sizes, the expansion of the fringe evolves more slowly than the contraction of the chains, reflecting the subtle impact of dynamics. Recall that Wal-Mart's direct impact on the flow profits of the fringe is negative, but the overall impact is positive since Wal-Mart pushes out the chains, who compete more directly with them, shifting the competitive landscape toward the smaller scale competitors. Interestingly, Wal-Mart has recently started to shift its focus toward offering much smaller stores (e.g. Wal-Mart Neighborhood Markets and Wal-Mart Express, which are closer in size to corner grocers and convenience stores), perhaps acknowledging the importance of these more localized offerings and the decreasing role of local scale.

8 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 computational burden that arises due to simultaneous moves. We recast the standard discrete-time, simultaneous-move game as a sequential-move game in continuous time. This significantly reduces the computational cost, greatly expanding the breadth and applicability of these structural methods.

By building on an underlying discrete-choice random utility framework, 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 an 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.

Using this formulation of a dynamic game in continuous time, we develop a dynamic model of retail competition that allows for substantial heterogeneity (both observed and unobserved) across agents and markets. We use the model to study the impact of Wal-Mart's entry into the retail grocery industry on market structure. The results imply that

Wal-Mart's entry on market structure varies greatly across markets, leading to an increase in market concentration in some markets which were initially served primarily by smaller fringe stores and to a sharp decrease in concentration in the majority of markets that were characterized by the presence of a number of large, dominant chains.

The inclusion of unobserved heterogeneity in the model is essential for uncovering these qualitatively distinct economic implications of Wal-Mart's entry across markets. Taken as a whole, the results of our analysis demonstrate the importance of incorporating substantial heterogeneity both across markets and firm types in estimating dynamic games of retail entry and competition, thereby highlighting the advantage of computationally light approaches for estimating and solving dynamic models with large state spaces.

A Proofs

A.1 Proof of Proposition 1

Let $\Sigma(\sigma)$ denote the $K \times K$ state transition matrix induced by the choice probabilities σ and the continuation state function $l(\cdot,\cdot)$. Let $\tilde{Q}_0$ denote the matrix formed by replacing the diagonal elements of Q_0 with zeros. Finally, let $E(\sigma)$ 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) as defined in Proposition 1.

We can rewrite the value function in (1) in matrix form as

$$\left[(\rho+\lambda)I - (Q_0 - \tilde{Q}_0)\right] V(\sigma) = u + \tilde{Q}_0 V(\sigma) + \lambda \left[\Sigma(\sigma)V(\sigma) + E(\sigma)\right].$$

Collecting terms involving $V(\sigma)$ yields

$$\left[(\rho+\lambda)I - \lambda\Sigma(\sigma) - Q_0\right] V(\sigma) = u + \lambda E(\sigma).$$

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 $\Sigma(\sigma)$ are in $[0,1]$, and $\rho > 0$ by assumption. Therefore, by the Levy-Desplanques theorem, this matrix is nonsingular (Horn and Johnson, 1985, Theorem 6.1.10). The result follows by multiplying both sides of the above equation by $\left[(\rho+\lambda)I - \lambda\Sigma(\sigma) - Q_0\right]^{-1}$.

A.2 Proof of Proposition 2

For choice j in state k let $v_{jk} = \psi_{jk} + V_{l(j,k)}$ denote an arbitrary choice-specific valuation and let $v_k = (v_{0k}, \dots, v_{J-1,k}) \in \mathbb{R}^J$ denote a J-vector of valuations in state k. Let σ_{jk} denote a choice probability and let $\sigma_k = (\sigma_{0k}, \dots, \sigma_{J-1,k}) \in \Delta^{J-1}$ denote an arbitrary J-vector of

CCPs where $\Delta^{J-1} = \{(p_0, \ldots, p_{J-1}) \in [0,1]^J : \sum_j p_j = 1\}$ is the unit J-simplex.

There is a mapping $H_k : \mathbb{R}^J \to \Delta^{J-1}$ from valuations to choice probabilities in state k where the j-th component is defined as $\sigma_{jk} = H_{jk}(v_k) = \Pr\left(j = \arg\max_{j' \in \mathcal{A}}\{v_{j',k} + \varepsilon_{j'}\}\right)$. Let $\mathcal{P}_k \subset \Delta^{J-1}$ be the space of all CCP vectors σ_k in state k. In other words, $\mathcal{P}_k$ is the codomain of H_k given by $\mathcal{P}_k = \{\sigma_k \in \Delta^{J-1} : \sigma_k = H_k(v_k), v_k \in \mathbb{R}^J\}$. For the normalizing choice j', let $\mathcal{V}$ denote the $(J-1)$-dimensional space of normalized valuations $\mathcal{V} = \{\tilde{v} \in \mathbb{R}^J : \tilde{v}_{j'} = 0\}$.

Now consider the mapping $\tilde{H}_k : \mathcal{V} \to \mathcal{P}_k$ defined by restricting H_k to the domain $\mathcal{V} \subset \mathbb{R}^J$ and the inverse correspondence $\tilde{H}_k^{-1}$. To see that the inverse $\tilde{H}_k^{-1}$ in $\mathcal{V}$ is nonempty, note that for any $\sigma_k \in \mathcal{P}_k$ there is a $v_k \in \mathbb{R}^J$ with $\tilde{H}_k(v_k) = \sigma_k$ by definition. Because the choice probabilities are invariant to the normalization of the valuations, there is a corresponding $\tilde{v}_k \in \mathcal{V}$ with $\tilde{H}_k(\tilde{v}_k) = \tilde{H}_k(v_k) = \sigma_k$. By Proposition 1 of Hotz and Miller (1993), also restated as Lemma 3.1 by Rust (1994), $\tilde{H}_k$ is one-to-one and therefore the inverse $\tilde{H}_k^{-1}(\sigma_k)$ is unique. The result follows by noting that the j-th component of the inverse, $\tilde{H}_{jk}^{-1}(\sigma_k)$, 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.3 Proof of Proposition 3

Furthermore, since the payoffs are additively separable we can relate the CCPs to the social surplus function of McFadden (1981): $S_k(v_k) \equiv \mathrm{E}\left[\max_{j \in \mathcal{A}}\{v_{jk} + \varepsilon_j\} \mid k\right]$. By the Williams-Daly-Zachary theorem (Rust, 1994, Theorem 3.1), S_k exists and for any $\alpha \in \mathbb{R}$, $S_k(v_k + \alpha) = S_k(v_k) + \alpha$. Using this additivity property for $\alpha = v_{j',k}$,

$$\begin{aligned}
\mathrm{E}\left[\max_j \{v_{jk} + \varepsilon_j\} \mid k\right] &= \mathrm{E}\left[\max_j \{v_{jk} - v_{j'k} + \varepsilon_j - \varepsilon_{j'}\} \mid k\right] + v_{j'k} + \mathrm{E}[\varepsilon_{j'} \mid k] \\
&= S_k(\tilde{v}_k) + v_{j'k} + \mathrm{E}[\varepsilon_{j'} \mid k] \\
&= V_{l(j',k)} + \psi_{j'k} + \Gamma^2(j', \sigma_k),
\end{aligned}$$

where, recalling from the proof of Proposition 2 that $v_{j',k} = V_{l(j',k)} + \psi_{j'k}$ and $\tilde{H}_k^{-1}(\sigma_k) = \tilde{v}_k$, we define $\Gamma^2(j', \sigma_k) = S_k(\tilde{H}_k^{-1}(\sigma_k)) + \mathrm{E}[\varepsilon_{j'} \mid k]$.

A.4 Proof of Proposition 4

Let $(j_k^1, \ldots, j_k^{D_k})$ 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 2 for the continuation choice $j = 0$, we can write

$$V_k = V_{k^*} + \sum_{d=1}^{D_k} \left(\psi_{j_k^d, l_k^d} - \psi_{0, l_k^d} \right) + \sum_{d=1}^{D_k} \Gamma^1(0, j_k^d, \sigma_{l_k^d}). \tag{22}$$

Recalling the Bellman equation from (1) and rearranging terms we can restate V_k as

$$\rho V_k = u_k + \sum_{l \neq k} q_{kl}(V_l - V_k) + \lambda \operatorname{E} \max_j \{\psi_{jk} + \varepsilon_j + V_{l(j,k)} - V_k\}. \tag{23}$$

Applying a similar procedure as in (22) 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 (23) in terms of a difference of terms of the form in (22), where the V_{k^*} term cancels leaving only sums of instantaneous payoffs ψ_{jk} and functions of the CCPs σ_k. Finally, using Proposition 3 and additivity, we can express the remaining term $\lambda \operatorname{E} \max_j \{\psi_{jk} + \varepsilon_j + V_{l(j,k)} - V_k\}$ as $\lambda \Gamma^2(0, \sigma_k) + \lambda \psi_{0k}$.

A.5 Proof of Propositions 5 and 6

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 (6). Let $\Sigma_m(\sigma_m)$ denote the $K \times K$ 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 $E_i(\sigma)$ 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 as defined in Proposition 6.

Then, we can rewrite (6) in matrix form as

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

Collecting terms involving $V_i(\sigma_i)$ yields

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

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 2. Therefore, by the Levy-Desplanques theorem, this matrix is nonsingular (Horn and Johnson, 1985, Theorem 6.1.10). Hence,

$$V_i(\sigma_i) = \left[(\rho + N\lambda)\, I - \sum_{m=1}^{N} \lambda \Sigma_m(\sigma_m) - Q_0\right]^{-1} \left[u_i + \lambda E_i(\sigma)\right]. \tag{24}$$

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

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

The mapping Υ is a continuous function from a compact, convex space into itself. By Brouwer's theorem, it has a fixed point. The fixed point probabilities imply Markov strategies that constitute a Markov perfect equilibrium.

A.6 Proof of Proposition 7

The proof proceeds along the lines of Theorem 1 of Blevins (2016), for identification of first-order systems of stochastic differential equations, while making use of properties that are specific to the case of Markov jump processes. In our model, Q can be expressed as a coefficient matrix in the (non-stochastic) system $P'(\Delta) = QP(\Delta)$ with $P(0) = I$. Recall from (11) that $P(\Delta)$ is the matrix exponential of ΔQ. There may be multiple solutions to (11), but we will show that Q is the only solution that satisfies the restrictions.

By Gantmacher (1959, VIII.8) and the distinct eigenvalue assumption on Q, all alternative solutions $\tilde{Q}$ to $\exp(\Delta\tilde{Q}) = P(\Delta)$ have the form $\tilde{Q} = Q + UDU^{-1}$ where U is the matrix of eigenvectors of Q and D is a diagonal matrix containing differences between the complex eigenvalues of Q and $\tilde{Q}$. This means that both the eigenvectors U and the real eigenvalues of Q are identified. But Q is an intensity matrix with zero row sums (i.e., $Qe = 0$, where e is a $K \times 1$ vector of ones) and so Q has a real eigenvalue equal to zero and hence the number of complex eigenvalues is at most $K - 1$.

For $\tilde{Q}$ to be admissible it must satisfy the prior restrictions $R\,\mathrm{vec}(\tilde{Q}) = r$. By the relationship between Q and $\tilde{Q}$ above, we have $R\,\mathrm{vec}(Q + UDU^{-1}) = r$. But $R\,\mathrm{vec}(Q) = r$, and by linearity of the vectorization operator, $R\,\mathrm{vec}(UDU^{-1}) = 0$. An equivalent representation is $R(U^{-\top} \otimes U)\,\mathrm{vec}(D) = 0$, where $U^{-\top} \otimes U$ is nonsingular. Following the proof of Theorem 1 of Blevins (2016), but with at most $K - 1$ complex eigenvalues, we can complete the system of equations to show that when there are at least $\left\lfloor \frac{K-1}{2} \right\rfloor$ linear restrictions and R has full rank, then D must be generically zero and therefore $\tilde{Q} = Q$.

Finally, we show that there are sufficiently many restrictions of full rank. The total number of distinct states is $K = \kappa^N$. Consider a single row of Q. The diagonal element is non-zero and for each player i, since there are J choices including the outside option, this leads to $J-1$ nonzero elements per row of the Q matrix associated with transitions due to player i. Therefore with N players there are a total of $K - N(J-1) - 1 = \kappa^N - NJ + N - 1$ zeros per row of Q. The order condition—that the number of restrictions is at least $\lfloor (K-1)/2 \rfloor$—is therefore satisfied when $\kappa^N - NJ + N - 1 \geq \left\lfloor (\kappa^N - 1)/2 \right\rfloor$ or when the number of choices J is such that $J \leq \frac{1}{N}\left[\kappa^N - \left\lfloor (\kappa^N - 1)/2 \right\rfloor + N - 1\right]$. The matrix R representing these restrictions has full rank because each restriction involves only a single element of Q and each row of R is a different row of the $K^2 \times K^2$ identity matrix. The sufficient condition stated in the text follows by noting that $\left\lfloor (\kappa^N - 1)/2 \right\rfloor \leq (\kappa^N - 1)/2$.

A.7 Proof of Proposition 8

Note that under Assumption 4, for any action $j > 0$ in any state k, the resulting state is always different from k. Therefore, the diagonal elements of S_{ij} are all zero and $S_{ij} - I_K$ has full rank for each $j > 0$. We established that Ξ_i is nonsingular in the proof of Proposition 5 above. It follows that X_i has full rank.

A.8 Proof of Proposition 9

Step 1: Uniform Convergence of L_M to L—We apply the uniform law of large numbers of Newey and McFadden (1994, Lemma 2.4) to establish uniform convergence. The data are independent and identically distributed, the parameter space is compact (Assumption 8), the observation likelihood is continuous at each (θ, h) with probability one, and the observation likelihood is strictly bounded between 0 and 1 under additive separability of the idiosyncratic shocks (Assumption 3) and Assumption 5, and since the rates λ and q_{kl} are bounded for all k and l (Assumption 2).

Step 2: Consistency of $\hat{\theta}$—By assumption, $\hat{h}$ is a $\sqrt{M}$-consistent M-estimator. Let $R_M(h) = \frac{1}{M}\sum_{m=1}^{M} r_m(h)$ denote the corresponding objective function and let $r_m \equiv r_m(h^0)$. (For example, for nonparametric MLE we would have $r_m(h) = \sum_{n=1}^{T} \ln P_{k_{m,n-1},k_{mn}}(\Delta; h)$). We have shown above that L_M converges uniformly in probability to L. L is also uniformly continuous in θ and h. Therefore, by Lemma 24.1 of Gourieroux and Monfort (1995), $L_M(\theta, \hat{h})$ converges in probability to $L(\theta, h^0)$ uniformly in θ. Then, Assumption 9 implies that θ^0 is the only element of Θ for which $h^0 = \Lambda(\theta, h^0)$. By the Kullback-Leibler information inequality, it follows that θ^0 is the unique maximizer of $L(\theta, h^0)$ in Θ. It follows by Theorem 2.1 of Newey and McFadden (1994) that $\hat{\theta} = \arg\max_{\theta \in \Theta} L_M(\theta, \hat{h}) \xrightarrow{\text{p}} \theta^0$.

Step 3: Asymptotic Normality of $\hat{\theta}$—The first order conditions for $\hat{\theta}$ are $\nabla_\theta L_M(\hat{\theta}, \hat{h}) = 0$.

By a mean value expansion between (θ^0, h^0) and $(\hat\theta, \hat h)$ and by consistency of the latter,

$$0 = \nabla_\theta L_M(\theta^0, h^0) + \nabla_{\theta\theta^\top} L_M(\theta^0, h^0)(\hat\theta - \theta^0) + \nabla_{\theta h^\top} L_M(\theta^0, h^0)(\hat h - h^0) + o_p(1).$$

By the central limit theorem and information matrix equality, $\nabla_{\theta\theta^\top} L_M(\theta^0, h^0) \xrightarrow{\text{p}} -\Omega_{\theta\theta^\top}$ and $\nabla_{\theta h^\top} L_M(\theta^0, h^0) \xrightarrow{\text{p}} -\Omega_{\theta h^\top}$. It follows that

$$\sqrt{M}(\hat\theta - \theta^0) = \Omega_{\theta\theta^\top}^{-1} \left\{ \Omega_{\theta h^\top} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^{M} \nabla_h r_m \right) + \left(\frac{1}{\sqrt{M}} \sum_{m=1}^{M} \nabla_\theta s_m \right) \right\} + o_p(1).$$

Finally, under the maintained assumptions Λ is continuously differentiable and the transition probabilities are bounded away from zero. The regularity conditions of Theorem 5.1 of Newey and McFadden (1994) are satisfied, so the generalized information matrix equality holds (Newey and McFadden, 1994, p. 2163), $\mathrm{E}[\nabla_h r_m \nabla_{\theta^\top} s_m] = 0$, and $\mathrm{E}[\nabla_h r_m \nabla_{h^\top} s_m] = I$. Thus,

$$\frac{1}{\sqrt{M}} \left(\sum_{m=1}^{M} \nabla_\theta s_m \right) - \Omega_{\theta h^\top} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^{M} \nabla_h r_m \right) \xrightarrow{\text{d}} \mathrm{N}\left(0, \Omega_{\theta\theta^\top} + \Omega_{\theta h^\top} \Sigma \Omega_{\theta h^\top}^\top \right).$$

The result holds by applying the continuous mapping theorem.

References

Ackerberg, D., X. Chen, and J. Hahn (2012). A practical asymptotic variance estimator for two-step semiparametric estimators. *Review of Economics and Statistics 94*, 481–498.

Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica 70*, 1519–1543.

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

Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics 156*, 38–67.

Arcidiacono, P. 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.

Basker, E. (2007). The causes and consequences of Wal-Mart's growth. *Journal of Economic Perspectives 21*, 177–198.

Basker, E. and M. Noel (2009). The evolving food chain: Competitive effects of Wal-Mart's entry into the supermarket industry. *Journal of Economics and Management Strategy 18*, 977–1009.

Beauchamp, A. (2015). Regulation, imperfect competition, and the U.S. abortion market. *International Economic Review 56*, 963–996.

Bellman, R. E. (1957). *Dynamic Programming.* Princeton, NJ: Princeton University Press.

Berry, S. T. (1992). Estimation of a model of entry in the airline industry. *Econometrica 60*, 889–917.

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

Blevins, J. R. (2016). Identifying restrictions for finite parameter continuous time models with discrete time data. *Econometric Theory*. Forthcoming.

Bresnahan, T. F. and P. C. Reiss (1991). Empirical models of discrete games. *Journal of Econometrics 48*, 57–81.

Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica 81*, 1003–1037.

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

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

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

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 75*, 1–14.

Ellickson, P. B., S. Houghton, and C. Timmins (2013). Estimating network economies in retail chains: A revealed preference approach. *The RAND Journal of Economics 44*(2), 169–193.

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

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

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

Gourieroux, C. and A. Monfort (1995). *Statistics and Econometric Models*, Volume 2. Cambridge University Press.

Haltiwanger, J., R. Jarmin, and C. Krizan (2010). Mom-and-pop meet big-box: Complements or substitutes? *Journal of Urban Economics 67*, 116–134.

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

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

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.

Igami, M. (2014). Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry. Working paper, Yale University.

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.

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

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

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica 70*, 801–816.

Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory 100*, 191–219.

Matsa, D. A. (2011). Competition and product quality in the supermarket industry. *Quarterly Journal of Economics 126*, 1539–1591.

Mazur, J. (2014). Can stricter bankruptcy laws discipline capital investment? Evidence from the U.S. airline industry. Working paper, Duke University.

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

McFadden, D. (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.

Moler, C. and C. V. Loan (1978). Nineteen dubious ways to compute the exponential of a matrix. *SIAM Review 20*, 801–836.

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

Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. In R. F. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, Amsterdam. North Holland.

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., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics 38*, 373–399.

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

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

Puterman, M. L. (2005). *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Second ed.). Wiley-Interscience.

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

Rust, J. (1994). Structural estimation of Markov decision processes. In R. F. Engle and D. L. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, Amsterdam, pp. 3081–3143. North Holland.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica 80*, 1019–1061.

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

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.

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

Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica 81*, 2013.

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica 76*, 1375–1411.

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 Wal-Marts 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 Wal-Marts	0.177	0.008	3
Number of Exiting Wal-Marts	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 Wal-Mart Entry

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

Standard errors in parentheses. Based on 90 markets where Wal-Mart is first observed to enter.

Table 3: Chain Firm Parameters

	Time-aggregated		With Wal-Mart Entry Times		No Unobserved Heterogeneity	
	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
Constant (θ_0^c)	4.470	(0.768)	4.403	(0.749)	2.561	(0.409)
Number of Chain Stores (θ_1^c)	-0.065	(0.024)	-0.067	(0.024)	-0.017	(0.014)
Number of Wal-Marts (θ_2^c)	-0.375	(0.148)	-0.383	(0.139)	-0.278	(0.108)
Number of Fringe Stores (θ_3^c)	-0.052	(0.017)	-0.053	(0.017)	-0.040	(0.012)
Number of Own Stores (θ_4^c)	-0.039	(0.081)	-0.044	(0.084)	0.104	(0.051)
Number of Own Stores Sq./100 ($100 \times \theta_5^c$)	-0.182	(0.432)	-0.165	(0.445)	-0.265	(0.166)
Population (θ_6^c)	0.176	(0.114)	0.213	(0.111)	0.267	(0.075)
Unobserved State (θ_7^c)	-0.956	(0.881)	-0.968	(0.806)		
Unobserved State $\times$ Number of Own Stores (θ_8^c)	0.245	(0.199)	0.249	(0.191)		
Entry Cost (η_0^c)	-18.377	(0.805)	-18.400	(0.807)	-17.643	(0.953)
Entry Cost $\times$ Unobserved State (η_1^c)	-5.151	(1.621)	-5.148	(1.676)		
Store Building Cost (κ_0^c)	-5.068	(0.876)	-5.073	(0.870)	-4.494	(0.782)
Store Building Cost $\times$ Unobserved State (κ_1^c)	3.513	(0.968)	3.508	(0.986)		
Exit Value (ϕ_0^c)	15.913	(0.888)	15.912	(0.896)	15.044	(0.633)
Exit Value $\times$ Unobserved State (ϕ_1^c)	4.166	(1.261)	4.126	(1.274)		

Table 4: Fringe Firm Parameters

	Time-aggregated		With Wal-Mart Entry Times		No Unobserved Heterogeneity	
	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
Constant (θ_1^{f})	-13.074	(0.080)	-13.092	(0.080)	-12.698	(0.067)
Number of Chain stores (θ_2^{f})	-0.021	(0.003)	-0.021	(0.003)	-0.018	(0.003)
Number of Wal-Marts (θ_2^{f})	-0.041	(0.012)	-0.042	(0.012)	-0.054	(0.012)
Number of Fringe Stores (θ_3^{f})	0.183	(0.008)	0.183	(0.008)	0.193	(0.008)
Number of Fringe Stores Squared / 100 ($100 \times \theta_4^{\mathrm{f}}$)	-0.349	(0.018)	-0.349	(0.019)	-0.369	(0.018)
Population (θ_5^{f})	0.240	(0.021)	0.248	(0.021)	0.170	(0.021)
Unobserved State (θ_6^{f})	-2.530	(0.107)	-2.544	(0.107)		
Unobserved State × Number of Fringe Stores (θ_7^{f})	0.050	(0.006)	0.051	(0.006)		
Entry Cost (η_0^{f})	-5.034	(0.033)	-5.034	(0.033)	-5.030	(0.033)
Entry Cost × Unobserved State (η_1^{f})	1.186	(0.079)	1.190	(0.079)		

Table 5: Counterfactual Simulations of Market Structure in Year 2014 With and Without Wal-Mart

	Markets	Initial Pop	Chain Firms	Chain Stores	Fringe Stores	Wal-Mart Stores	Chain Share	Wal-Mart Share	Fringe Share	C1	C3	HHI
						With Wal-Mart						
All Markets	205	176153	2.41	9.17	11.98	2.42	39.9%	10.8%	49.4%	25.4%	48.1%	0.22
Midwest	58	175371	1.75	5.88	14.36	2.07	27.3%	9.9%	62.7%	21.7%	39.1%	0.20
Northeast	22	205180	2.18	8.48	14.32	2.58	35.2%	10.7%	54.1%	24.0%	45.3%	0.21
South	83	170856	2.78	11.72	9.63	2.85	49.1%	12.1%	38.7%	29.1%	55.8%	0.24
West	42	172494	2.71	9.02	12.11	1.96	41.3%	9.3%	49.4%	23.8%	46.6%	0.20
						Absent Wal-Mart						
All Markets	205	176153	2.77	12.43	9.85	0.00	54.9%	0.0%	44.6%	29.9%	55.7%	0.26
Midwest	58	175371	2.13	8.41	11.81	0.00	42.0%	0.0%	58.0%	27.6%	47.5%	0.25
Northeast	22	205180	2.61	12.22	11.18	0.00	53.7%	0.0%	46.3%	30.6%	55.4%	0.27
South	83	170856	3.22	16.15	7.54	0.00	66.9%	0.0%	32.7%	33.2%	64.3%	0.28
West	42	172494	2.86	10.77	11.02	0.00	49.5%	0.0%	48.8%	25.9%	50.2%	0.23

Table 6: Counterfactual Simulations of Changes in Market Structure Due to Wal-Mart's Presence

	Markets	Initial Pop	Wal-Mart Stores	Chain Stores	Fringe Stores	Chain Share	Fringe Share	C1	C3	HHI
All Markets	205	176153	2.42	-26.3%	21.6%	-27.3%	10.6%	-15.0%	-13.7%	-16.6%
By Region										
Midwest	58	175371	2.07	-30.1%	21.6%	-34.8%	8.1%	-21.3%	-17.8%	-19.3%
Northeast	22	205180	2.58	-30.5%	28.1%	-34.5%	16.8%	-21.5%	-18.4%	-21.8%
South	83	170856	2.85	-27.4%	27.7%	-26.5%	18.7%	-12.4%	-13.2%	-15.6%
West	42	172494	1.96	-16.3%	9.9%	-16.6%	1.2%	-8.1%	-7.2%	-11.9%
By Market Size										
Small	104	117740	1.76	-24.3%	7.0%	-23.1%	5.0%	-12.1%	-9.9%	-14.1%
Large	101	236300	3.09	-27.4%	30.0%	-31.7%	16.2%	-18.1%	-17.9%	-19.4%
By Growth Type										
Slow	54	178252	2.26	-35.8%	40.6%	-36.5%	24.3%	-22.3%	-21.0%	-23.8%
Moderate	46	175444	2.17	-38.9%	18.5%	-38.3%	15.2%	-16.3%	-17.6%	-9.7%
Fast	105	175383	2.61	-17.4%	13.5%	-18.6%	1.0%	-10.1%	-8.4%	-14.9%
By Unobserved Type										
More Negative	9	106248	1.20	-18.5%	-17.2%	-6.7%	-6.2%	29.6%	27.3%	42.3%
Negative	68	127754	1.62	-15.9%	2.7%	-17.9%	-2.5%	-9.4%	-5.5%	-11.3%
Zero	96	184404	2.20	-27.8%	25.2%	-31.1%	18.2%	-20.0%	-19.2%	-22.0%
Positive	32	273906	5.08	-30.5%	59.1%	-33.7%	40.6%	-15.1%	-17.8%	-17.7%

Table 7: Counterfactual Simulations of Changes in Market Structure Absent Unobserved Heterogeneity

	Wal-Mart Stores	Chain Stores	Fringe Stores	Chain Share	Fringe Share	C1	C3	HHI
All Markets	3.15	-33.2%	7.3%	-33.1%	9.0%	-12.9%	-9.7%	-7.0%
				By Region				
Midwest	2.92	-38.4%	9.7%	-41.5%	8.3%	-17.2%	-12.0%	-6.6%
Northeast	3.24	-35.8%	7.0%	-36.0%	9.7%	-14.6%	-10.6%	-7.5%
South	3.28	-32.3%	7.4%	-29.8%	11.7%	-10.0%	-8.5%	-6.6%
West	3.15	-28.6%	3.4%	-29.4%	3.9%	-12.4%	-8.7%	-7.6%
				By Market Size				
Large	3.65	-30.5%	7.3%	-29.6%	7.2%	-12.2%	-9.5%	-7.8%
Small	2.66	-37.5%	7.1%	-36.2%	11.5%	-13.9%	-10.3%	-6.8%
				By Unobserved Type				
More Negative	2.22	-42.9%	15.6%	-43.9%	13.1%	-16.6%	-14.9%	-2.9%
Negative	2.71	-35.9%	6.0%	-36.9%	8.5%	-14.9%	-9.8%	-6.5%
Zero	3.24	-33.7%	6.7%	-32.8%	7.6%	-12.2%	-9.1%	-6.2%
Positive	4.07	-29.6%	9.3%	-24.4%	13.7%	-10.1%	-10.1%	-12.5%

Table 8: Temporal Evolution of Market Structure

Year	Market Size	WM Stores	Chain Stores	Fringe Stores	Chain Share	Fringe Share	C1	C3	HHI
5	Small	0.58	-6.8%	3.4%	-7.9%	-0.1%	-5.7%	-6.3%	-9.9%
5	Large	0.87	-10.0%	5.2%	-9.9%	2.9%	-5.2%	-7.5%	-10.1%
10	Small	1.06	-13.7%	5.4%	-14.3%	0.8%	-9.3%	-9.2%	-14.3%
10	Large	1.76	-16.8%	12.8%	-18.3%	6.6%	-10.2%	-12.0%	-15.6%
15	Small	1.45	-19.5%	6.5%	-19.4%	2.0%	-11.5%	-10.3%	-15.4%
15	Large	2.52	-22.3%	21.3%	-25.6%	11.1%	-14.6%	-15.2%	-18.2%
20	Small	1.76	-24.3%	7.0%	-23.1%	5.0%	-12.1%	-9.9%	-14.1%
20	Large	3.09	-27.4%	30.0%	-31.7%	16.2%	-18.1%	-17.9%	-19.4%