

The Dynamics of Retail Oligopoly*

Arie Beresteanu^a, Paul B. Ellickson^b, Sanjog Misra^c, and James C. Reeder,
III^d

^aUniversity of Pittsburgh

^bSimon Business School, University of Rochester

^cBooth School of Business, University of Chicago

^dKU School of Business, University of Kansas

February 23, 2024

Abstract

This paper empirically examines competition between supermarkets as a dynamic discrete game between heterogeneous players. We focus on the overall impact of Wal-Mart's entry on incumbent supermarket firms, quantifying the effects on prices, producer surplus, consumer welfare and overall competitive structure. Employing a thirteen-year panel dataset of store level observations that includes every supermarket firm operating in the United States alongside the rapid proliferation of Wal-Mart Supercenters, we propose and estimate a dynamic structural model of chain level competition in which incumbent firms decide each period whether to add or subtract stores or exit the market entirely, and potential entrants choose whether or not to enter. Product market competition is modeled as a discrete-choice demand system, incorporating detailed information on prices and characteristics of chains, as well as unobserved heterogeneity in chain-level quality. Our estimation approach combines two-step estimation techniques with a novel random forest based value function approximation algorithm to accommodate the high-dimensional structure of the underlying state space.

Keywords: Retail Grocery, Dynamic Oligopoly, Value Function Approximation, Machine Learning, Random Forest

*The authors would like to thank seminar participants at Cornell, Duke, Harvard, Northwestern, Penn State, the Federal Reserve, the Federal Trade Commission, the Department of Justice, and the AEA winter meetings. Any remaining errors and omissions are our own.

1 Introduction

This paper examines competition between supermarkets as a dynamic discrete game between heterogeneous players. We focus on the overall impact of Wal-Mart’s entry on incumbent supermarket firms, quantifying the effects on prices, producer surplus, consumer welfare and overall competitive structure. We propose and estimate a dynamic structural model of chain level competition using a thirteen-year panel dataset of store-level observations, which incorporates all supermarket firms operating in the US, alongside the rapid expansion of Wal-Mart Supercenters. In this model, existing firms make decisions each period to either add or remove stores, or completely exit the market, while potential entrants decide on whether to enter the market or not. Product market competition is modeled as a discrete-choice demand system, incorporating detailed information on prices and characteristics of chains, as well as unobserved heterogeneity in chain-level quality.

The theoretical framework proposed in this paper is based on the Markov perfect equilibrium (MPE) framework of Ericson and Pakes (1995), in which firms make discrete strategic investments that increase the quality of their products. In the context of retail competition, in which firms operate a chain of individual stores, quality is a function of the total number of stores operated by each firm, the individual characteristics of their stores, and their overall retail format (conventional supermarket or supercenter). Allowing firms to adjust all of these features independently would yield an intractably complex dynamic control problem. Instead, our strategy is to focus on a single dimension of quality (store density) and allow firms to differ by format (supermarket or supercenter). Product market competition is modeled using a discrete choice model of demand, in which firms may also differ by an additional dimension of overall firm quality that is fixed over time. We assume that the economically relevant features of the industry can be encoded into a state vector that includes each firm’s store density, its overall format, and its quality level. Firms receive state dependent payoffs in the product market and influence the evolution of the state vector through their entry, exit, and investment decisions. In particular, incumbent firms can adjust their chain size

each period by either opening new stores, closing existing ones, or even exiting the market entirely. One potential entrant can enter into each format *de novo* every period. Equilibrium obtains when all firms choose strategies that maximize their expected discounted profits, given the expected actions of their rivals.

We estimate this model of competition using a detailed panel dataset that follows the entire supermarket industry over thirteen consecutive periods (years) spanning 1994 to 2006. Our empirical framework extends the two-step estimation approach proposed by Aguirregabiria and Mira (2007) by incorporating a novel value function approximation algorithm to tackle the large scale of the underlying state space. In the first step, we recover the firm’s policy functions governing entry, exit, and investment. These functions characterize firms beliefs regarding the evolution of the common state variables and the actions of their rivals. We also estimate the per-period payoff that each firm receives as a function of the current state. In the second step, we use the structure of the dynamic optimization problem to recover the parameters that make those beliefs consistent with an MPE. Following Hotz and Miller (1993), this is accomplished by replacing the continuation values in the best response probability functions that characterize the MPE with inverted conditional choice probabilities (CCPs) that can be recovered non-parametrically from the data.

A challenge to estimation stems from our large state space and rich choice set, which makes generating the continuation values computationally burdensome. With the array of continuous variables that describe the current state facing each player in the market, which is further complicated by a varying number of players each round, current methods proposed in the dynamic games literature proved insufficient. To this end, we extend the dynamic games literature by developing a new estimation strategy. Our proposed method leverages advances in machine learning to create a data-adaptive kernel weighed projection of the continuation value.¹ By using the adaptive, nearest-neighbor design of the regression forest to pre-specify

¹The use of data-adaptive kernels is not new in the machine learning literature, as they have been employed in survival analysis (Hothorn et al. (2004)), quantile regressions (Meinshausen and Ridgeway (2006)), and, most recently, the causal forest approach (Wager and Athey (2018)).

the relationship of observed, or calculated, state variables, we are able to tie the current period's guess of the continuation value to any combination of observed state variables that appear one period ahead. In essence, our estimation approach allows us to generate and then freeze a random forest for updating the continuation value in a similar manner as Crawford and Shum (2005) or Sweeting (2013), but avoiding the use of a restrictive linear form for the mapping from basis functions to outcomes.

We build upon a rich literature on the estimation of dynamic discrete choice and dynamic games that includes the seminal methodological contributions of Aguirregabiria and Mira (2007), Bajari et al. (2007), Pakes et al. (2007) and Pesendorfer and Schmidt-Dengler (2007), as well as the more application oriented papers of Ryan (2012), Collard-Wexler (2013), Dunne et al. (2013), Barwick and Pathak (2015), Fowlie et al. (2016), Kalouptsi (2018) and Hollenbeck (2017).²

We also connect to the literature that considers supermarket competition, and the impact of Walmart on retailing more specifically. Important contributions include Foster et al. (2006); Basker and Noel (2009); Hausman and Leibtag (2007); Holmes (2011); Ellickson and Grieco (2013); Arcidiacono et al. (2016, 2020); Thomassen et al. (2017); Lagakos (2016); Atkin et al. (2018); Handbury (2021); Houde et al. (2023).

We find that Supercenters enjoy a cost advantage in building new stores within an established location. However, the Supercenter pays a sizeable up-front fixed cost to enter a new market. This is not surprising, as Supercenters often rely on well-organized distribution networks to enhance their cost positions. Through our counterfactual analysis, we find that markets do concentrate in a scenario where Supercenters are allowed free rein to grow. In the scenario where Supercenters are banned, we see fiercer competition among supermarkets, where smaller supermarkets are allowed to survive. That said, even with the greater amount of competition in the scenario with only supermarkets, the estimated consumer surplus is marketed lower compared to the one where Supercenters are allowed. This occurs for two

²For an overview of the rich empirical literature that has developed over the past two decades, see Aguirregabiria et al. (2021).

primary reasons. First, the technological superiority of the Supercenter format is to the benefit of the consumer. Second, the investment in gaining a lower marginal cost position suppressed prices in the market containing Supercenters, and these lower prices paid by the consumers are of great benefit. As we show, the customers value the trade-off of fewer options as being less important than having access to the dominant technology and lower prices in this industry.

The paper is organized as follows. Section 2 describes the construction of the dataset. Section 3 describes the theoretical framework. The empirical framework is described in Section 4. The results of the first and second steps of the estimation are presented in Section 5, while the results of the policy experiments are contained in Section 6. Section 7 concludes.

2 Data and Sample Construction

The data for the supermarket industry are constructed from yearly snapshots of the Trade Dimension's Retail Tenant Database (RTD) spanning the years 1994 to 2006, while market specific population levels and growth rates are drawn from the US Census. The RTD includes information on every supermarket and supercenter operating in the US during this period. The (establishment level) definition of a supermarket employed by Trade Dimensions is the government and industry standard: a store selling a full line of food products and generating at least \$2 million in yearly revenues.³ Every outlet of all the major US supermarket chains is well above this threshold, as are all Walmart supercenters.

Information on average weekly volume, store size, number of checkouts, number of employees (full time equivalents), and the overall format of the store (e.g. Supercenter or conventional supermarket) is gathered through quarterly surveys sent to store managers. These surveys are then compared with similar surveys given to the principal food broker assigned to each store and further verified via repeated phone calls. Each store is assigned a

³Foodstores with less than \$2 million in revenues are classified as convenience stores and are not included in the dataset. Firms in this segment operate very small stores and compete with only the smallest grocery stores.

unique identifier code that remains with the store regardless of ownership, which we used to construct the overall store panel. In addition, each store has a unique *firm* code, which we used to identify the ultimate owner. The availability of reliable firm identifiers is critical in the supermarket industry, since parent firms often operate stores under several “flag names,” especially when the stores have been acquired by merger. To avoid problems of false exits and entries, we treat stores acquired in a merger as having always belonged to their final owner. Also, when a firm is taken private or bought out by a public holding company, we do not treat the event as an entry (or exit).

Previous empirical studies of the supermarket industry (Ellickson (2007), Smith (2004)) suggest dividing the retail grocery market into two distinct submarkets: supermarkets and grocery stores. Supermarkets compete in a tight regional oligopoly that is not a significant strategic rival of the much smaller and highly fragmented grocery segment. Furthermore, the number of firms in these oligopolies do not increase with market size, yielding an equilibrium firm count that is apparently not impacted by population size or growth. The RTD includes information on both types of firms. Since we are primarily concerned with competition between retail oligopolists, we focus only on the “top” firms in each market. Following standard practice, we define a local market here to be a US Metropolitan Statistical Area (MSA). We then include in our panel only those firms that served at least 5% of the market in which they operated in at least one period. Because the top supermarket firms do not compete significantly with the grocery firms in the fringe, this should not introduce any selection problems.⁴

The discrete choice model we use to characterize product market competition requires us to specify and collect data on the sales of the outside good. Obvious consumer alternatives to supermarkets include grocery stores, convenience stores, liquor stores, restaurants, and cafeterias. Therefore, we assume that total sales of the outside good are equal to the combined sales of all food and beverage stores (NAICS 445 - of which supermarkets are a subset)

⁴Ellickson (2006) and Smith (2004) both present empirical results that support the separate submarkets claim.

and all food service and drinking establishments (NAICS 772) less the sales accounted for by supermarkets alone. Data on total sales is taken from the 1997 Census of Retail Trade. To construct the share of the outside good, we use the Census dataset to construct an MSA specific multiplier characterizing the ratio of total sales in both categories (445 and 772) to total sales in supermarkets alone (NAICS 44511). We then use this multiplier to impute the total sales in both categories for each MSA in our dataset, using the observed revenue of the supermarkets as our baseline measure of sales. We are thereby implicitly assuming that the ratio is constant over time.

Estimating this demand system also requires data on firm level prices, which we acquired from the American Chamber of Commerce Researchers Association (ACCRA). The ACCRA collects data from over 250 U.S. towns and cities on the prices of various retail products (26 of which are grocery items) for use in the construction of their *Cost of Living Index*. The ACCRA sends representatives to several supermarkets in each geographic market with the goal of collecting a representative sample of prices at the major chains. They are given a specific list of products for which to collect individual prices (e.g. 50 oz. Cascade dishwashing powder). We obtained their disaggregated dataset, so we observe the store name and individual prices for each product. We then use these individual prices to construct a price index (using the same weights employed by ACCRA) for each store in their dataset that is inflated to match average weekly grocery expenditures, as reported by the BLS. Since we are modeling competition at the firm level (and assuming prices are set at that level as well), we then aggregated these indices up to the level of the firm (in each market) and matched them to the corresponding firms in our panel, yielding a total of 649 MSA/firm level observations on price. Since ACCRA only began recording the names of the individual stores in 2004, we have prices for only a single period. Summary statistics are provided in Table 1.

Note that Walmart supercenters are almost twice as large as the stores operated by supermarket firms, with three times as many checkouts. Store size here is grocery space, but checkouts are the total for the store. Supermarket chains operate almost three times

as many stores per market as Walmart supercenters, but capture roughly the same share of the market. Walmart supercenters are about 14% cheaper than conventional supermarkets, which is consistent with prior studies (Hausman and Leibtag, 2007; Basker and Noel, 2009; Arcidiacono et al., 2020).

Table 1: Summary Statistics

	Format	
	Supercenter	Supermarket
Store Size	65.2 (22.8)	36.3 (15.2)
Checkouts	29.4 (6.36)	10.1 (3.94)
Stores per Market	3.64 (5.48)	10.4 (22.6)
Market Share	16.6 (13.9)	15.1 (10.2)
Basket Price	82.08 (6.31)	95.66 (10.28)
Firms per MSA	.70 (.64)	4.38 (1.42)

Store size is in 1000s of square feet.

Table 2: Action Frequencies (per period)

	Potential Entrants			Incumbents	
	Supercenter	Supermarket		Supercenter	Supermarket
Don't Enter	93.5%	92.9%	Exit	1%	2.7%
Build 1	5.3%	6.1%	Close 2+		2.5%
Build 2	1.2%	.6%	Close 1		6.3%
Build 3		.4%	Do Nothing	71%	74%
			Open 1	18%	9%
			Open 2+	10%	5.2%

3 A Model of Retail Chain Competition

The discrete game structure and notation employed here closely follows Sweeting (2013), with differences noted as they arise. We characterize the overall game structure at a fairly

high level here, providing the more specific details that map to estimation in subsequent sections.

In each of $m = 1, \dots, M$ markets (MSAs), grocery chains (both supermarkets and supercenters) indexed $f = 1, \dots, \mathcal{F}_m$ play a discrete time game over periods $t = 1, \dots, T$. Although almost all chains complete in more than one market, we assume that these markets are independent (i.e. there are no benefits from common ownership or multi-market contact). Markets are differentiated by two common state variables characterizing their population levels (pop_m) and growth rates (g_m). These are taken to be exogenous, deterministic and commonly observed. For convenience, we omit the market subscript m .

Firms compete by offering a basket of groceries to each consumer in the market. These firm-level baskets are differentiated by price and by the characteristics of the chains that provide them. The characteristics that define each chain are its format (either supercenter or supermarket), the number of stores it operates per capita (a form of capacity, as well as a proxy for the travel distance to its customer base), and its overall quality level (inferred from the demand system). Chains are assumed to set market-level prices to maximize per-period profit.⁵ Prices are therefore determined by the static Nash equilibrium in the product market. The collection of state variables, commonly observed by all market participants, is denoted \mathcal{M}_{jft} , where j distinguishes a firm f 's state in period t .

The per-period profits that firms earn are denoted $\pi^v(\mathcal{M}_{jft})$. The full per-period flow profit function includes the variable profits from product market competition, as well as the cost of adjusting the state (e.g., by opening or closing stores) or exiting. These costs are denoted $C(a)$, as they depend on the firm's choice of action a . Each firm also privately observes state variables (cost shocks) associated with opening and closing stores, which is explained below. For each firm in each year in each market, the researcher directly observes the firm's format ($o \in SM, SC$), the number of stores it operates, the price it charges in the product market and its total revenue. Each firm's overall quality level is inferred from the

⁵Hitsch et al. (2021) and DellaVigna and Gentzkow (2019) both present empirical results consistent with the market-level pricing assumption.

estimated demand system and treated as observed (by rival firms and the researcher) in the dynamic game. Marginal costs are also recovered from the demand system estimates, via the first order conditions of the static pricing game.

The timing of the game is as follows. Each period, all currently active (incumbent) firms observe the current state (the market population and growth rate, the firm’s own store count, quality level, and format, as well as those of its rivals). Each incumbent then decides whether to open additional stores, close existing stores, maintain the current portfolio or exit the market entirely. The number of new store openings is capped at two, as is the total number of closures (larger changes are rarely observed). The choice set $\mathcal{A}_f(\mathcal{M}_{jft})$ depends on the current state, as a firm with one store cannot choose to close two. Coupled with each action $a_{ft} \in \mathcal{A}_f(\mathcal{M}_{jft})$ is a cost/payoff shock $\varepsilon_{ft}(a_{ft})$, which is privately observed by the focal firm and assumed to be independent and identically distributed (iid) Type I Extreme Value, times a scale parameter⁶ θ^ε that accommodates heteroskedasticity by population and firm type via the following relation:

$$\theta^\varepsilon = \exp(\beta_1 * \log(pop) + \beta_2 * SC_f).$$

The flow profit of firm f in period t can then be written as

$$\begin{aligned} \pi_{ft}(a_{ft}, \mathcal{M}_{jft}, \theta, \gamma) + \theta^\varepsilon \varepsilon_{ft}(a_{ft}) \\ = \pi^v(\mathcal{M}_{jft}, \gamma) - C(a_{ft})\theta^C + \theta^\varepsilon \varepsilon_{ft}(a_{ft}) \end{aligned} \tag{1}$$

where π_{ft} is the net profit, π_{ft}^v is variable profit and $C()$ is the total cost function. Note that we are making use of the typical “time to build assumption” whereby variable profits depend on the current state, adjustment costs are incurred in the current period t , and changes to the state occur in the subsequent period $t + 1$.

Two potential entrants (one of each format type) may also choose to enter the market

⁶The scale parameter is identified here since we observe (and condition upon) revenues when estimating the parameters governing dynamic investments.

in each period. Entrants may enter with either one or two stores and are assigned a quality draw from the (recovered and known) distribution of firm-level qualities (inferred from the estimated demand system). Because stores take a period to build, flow profits in the initial period include only the entry costs (and no current payoff). The entry costs are parameterized as a linear function of population and the number of stores initially opened. After making these choices (but before they are realized), each incumbent firm competes in the product market, earning variable profits given by $\pi_{ft}(\mathcal{M}_{jft}, \gamma)$, where γ is a parameter vector indexing the consumer utility function (to be specified below). Thus, each firm's flow profits include the costs/payoffs of opening or closing stores (or choosing to exit), as well as their associated shocks, plus the variable profit from product market competition.⁷

Assuming that firms play stationary Markov Perfect Nash Equilibria (MPNE), let the mapping from states to actions be denoted $\Gamma_f : (\mathcal{M}_{jft}, \varepsilon_{ft}) \rightarrow a_{ft}$. Firm f 's value in state $(\mathcal{M}_{jft}, \varepsilon_{ft})$ when it employs an optimal strategy (and all rivals follow Γ) is given by:

$$\begin{aligned} V_f^\Gamma(\mathcal{M}_{jft}, \varepsilon_{ft}) &= \max_{a \in A_f(\mathcal{M}_{jft})} [\pi(a, \mathcal{M}_{jft}) + \theta^\varepsilon \varepsilon_{ft}(a)] \\ &+ \beta \int \bar{V}_f^\Gamma(\mathcal{M}_{jft+1}) g(\mathcal{M}_{jft+1} | a, \Gamma_{-f}, \mathcal{M}_{jft}) d\mathcal{M}_{jft+1} \end{aligned} \quad (2)$$

where $g(\cdot)$ is the transition kernel given choice a and rival strategies Γ_{-f} , and $\bar{V}_f^\Gamma(\mathcal{M}_{jft+1})$ is the ex ante value function, obtained by integrating over the unobserved state variables

$$\bar{V}_f^\Gamma(\mathcal{M}_{jft}) = \int V_f^\Gamma(\mathcal{M}_{jft}, \varepsilon_{ft}) f(\varepsilon_{ft}) d\varepsilon_{ft} \quad (3)$$

The distributional assumption on ε yields closed form solutions for the optimal strategy

⁷Per-period fixed costs are not separately identified from entry costs and salvage values, and are therefore normalized to zero. For more detail on the normalizations required for estimation, as well as subsequent counterfactual simulation and interpretation in the context of dynamic discrete games, see Aguirregabiria and Suzuki (2014) and Kalouptsi et al. (2021).

profile of firm f :

$$P^{\Gamma_f}(a, \mathcal{M}_{jft}, \Gamma_{-f}) = \frac{\exp\left(\frac{v_f^{\Gamma}(a, \mathcal{M}_{jft}, \Gamma_{-f})}{\theta^\varepsilon}\right)}{\sum_{a' \in \mathcal{A}_f(\mathcal{M}_{jft})} \exp\left(\frac{v_f^{\Gamma}(a', \mathcal{M}_{jft}, \Gamma_{-f})}{\theta^\varepsilon}\right)} \quad (4)$$

where $v_f^{\Gamma}(a, \mathcal{M}_{jft}, \Gamma_{-f})$ is the following choice specific value function (net of shock)

$$\begin{aligned} v_f^{\Gamma}(a, \mathcal{M}_{jft}, \Gamma_{-f}) &= \pi(a, \mathcal{M}_{jft}) \\ &+ \beta \int \bar{V}_f^{\Gamma}(\mathcal{M}_{jft+1}) g(\mathcal{M}_{jft+1} | a, \Gamma_{-f}, \mathcal{M}_{jft}) d\mathcal{M}_{jft+1} \end{aligned} \quad (5)$$

Note that, given the Type 1 Extreme Value assumption for the private shocks, the ex ante value function can now be expressed via the well-known log sum formula

$$\bar{V}_f^{\Gamma}(\mathcal{M}_{jft}) = \theta^\varepsilon \log \left[\sum_{a'_t \in A} \exp\left(\frac{v_f^{\Gamma}(a', \mathcal{M}_{jft}, \Gamma_{-f})}{\theta^\varepsilon}\right) \right] + \theta^\varepsilon \lambda^{Euler} \quad (6)$$

where λ^{Euler} is Euler's constant.

A Markov Perfect Equilibrium of the discrete game is a fixed point of the collection of best response probability functions (4) of all four player types: incumbent supermarkets and supercenters, as well as one potential entrant each period of each type. While existence of equilibria follows from Brouwer's fixed point theorem, uniqueness is not guaranteed. For the purposes of estimation, we follow typical practice in the empirical games literature in assuming that a unique equilibrium is played in the data and can then be conditioned upon directly as part of our two-step estimation procedure.

4 Estimation Overview

The estimation of single agent dynamic discrete choice problems typically proceeds by matching the model implied conditional choice probabilities (CCPs), represented above by equation

(4), to the discrete actions (choices) observed in the data, using either a GMM or (pseudo) MLE criterion. Because the choice specific value functions that appear on the right-hand side of (4) are not economic primitives, they must either be solved for using a nested fixed point routine (as in Rust (1987)) or replaced by a sample analog as part of a two-step estimation procedure (as in Hotz and Miller (1993) or Hotz et al. (1994)).

Dynamic discrete games raise additional complications due to the doubly nested structure induced by the future value and rival action components of the agent’s decision problem. A powerful solution to this computational problem involves replacing the CCPs of rival firms with sample analogs, effectively turning the estimation problem into a collection of single-agent games against nature.⁸ This two-step approach has been implemented successfully in many contexts. However, the high-dimensional nature of our underlying state space, arising from the large number of actual and potential players, as well as the continuous nature of some of the state variables (e.g. firm quality and store density), makes a direct application of these existing methods intractable here. A scalable approach is also needed to perform the back-end counterfactuals that constitute the key substantive contribution of the paper.

To proceed, we propose an alternative estimation approach that involves the following (CCP-based) two steps. First, rivals actions are found by integrating the right-hand side of equation (5)). Then, to solve the individual dynamic programming problem, we employ value function iterations via approximation. The integrated value function is then interpolated on points outside of the finite grid on which the full problem is solved.⁹ Exploiting recent

⁸This method was pioneered in a series of contemporaneous papers (Aguirregabiria and Mira, 2007; Bajari et al., 2007; Pakes et al., 2007; Pesendorfer and Schmidt-Dengler, 2007) following a strategy originally proposed by Rust (1994). Arcidiacono and Ellickson (2011) provide a description of these approaches, highlighting the various benefits and trade-offs of the alternative estimation approaches.

⁹Value function approximation was originally proposed in Bellman et al. (1963) and subsequently built upon and extended by many authors (see Judd (1998) and Powell (2007)). Here, we use the value function iteration approach developed for a single-agent problem in Crawford and Shum (2005) in a game context similar to Sweeting (2013). One key difference is that Sweeting (2013) employed a parametric policy iteration approach based on Benitez-Silva et al. (2000), while we use value function interpolation and iteration instead. Another distinction is that we use a random forest for function approximation rather than a linear estimator. Other value function approximation approaches have been employed successfully by many authors, including Keane and Wolpin (1994), Barwick and Pathak (2015), Fowlie et al. (2016) and Kalouptside (2018), often using a Lasso approach to select over a set of basis functions defined by polynomial, sieve or series approximation.

advances in machine learning, we build our approximation to the value function in two steps. First, we use forward simulation of actions and payoffs to generate an analog of the “ground truth” upon which to anchor the structure of a random forest. We then fit a regression forest to these outcomes to identify the “nearest neighbors” to each point in the finite grid using the data-adaptive kernel induced by the forest routine. This structure is then frozen, and the resulting partitions and partition weights are used in the value function iteration procedure. The key innovation here lies in using the forward simulated payoff functions to anchor the overall “shape” of the resulting forest, and then freezing that structure in place for the iterative updating process (the nested fixed point routine). In other words, we update the weighted average, but not the weights or groupings. Avoiding the high degree of smoothing typically associated with series approximation should allow the approximated surface to better reflect the nonlinearities and discontinuous “jumps” in the value function that are often observed in discrete games (Doraszelski and Pakes, 2007).

At a high level, our estimation approach proceeds as follows. We first recover estimates of the parameters characterizing the demand system using a standard discrete choice demand approach (Berry, 1994). Inverting the first order conditions of the static pricing problem provides estimates of the marginal costs of production, which are projected down onto store density and a dummy for format type (for prediction and extrapolation purposes). From these primitives, we can then construct estimates of variable profit for any point in the state space (including both points observed directly in the data and points outside of its support). These are used when constructing estimates of the flow profit term on the right-hand side of equation (5), as well as generating the outcomes needed to determine the structure of the regression forest.

We turn next to recovering an initial estimate of the partitions and partition weights that govern the random forest approximation to the value function surface. To do so, we first specify a finite grid of points on the natural state space upon which to solve the value function exactly. We then define a collection of aggregate states or “features” that summarize

the key aspects of the observed market structure at a given location in the state space from the perspective of a focal firm. Following the logic of Powell (2007), these are chosen to reflect moments or features of the underlying “natural” states then act as “basis functions” for the prediction problem. We then forward simulate payoffs for each focal firm on the grid to recover an initial empirical analog of the value function at each point on the grid. These simulated values are the outcomes that the random forest is tasked with predicting. The output of this prediction exercise is a collection of partitions of the aggregate state space, along with partition weights, than can then be used to predict values at points outside the fixed grid.

We next move to the estimation of the dynamic structural parameters, freezing the partitions and partition weights, but updating the values on the grid using the structure of the Bellman equation (the value function iteration process) and the CCPs of rival players. This procedure is nested inside a pseudo MLE routine to iteratively update the structural parameters, and repeated until the full process converges. Finally, the entire procedure is repeated once more to incorporate the resulting cost parameters into the initial forward simulation step to further refine the structure of the forest to better capture the full shape of the value function. Bootstrapped standard errors are computed by subsampling markets during this final step.

The estimation of the dynamic game proceeds in two main stages. In the first stage, or “set-up” step, we compile several different functions and approximations that characterize demand and static product market competition, as well as the initial partitions of the aggregate state space and partition weights. In this stage, the static parameters that define product market competition are recovered, a collection of aggregate states is defined, the reduced-form CCPs are estimated, three approximation grids are specified, and the distribution of possible one-period-ahead states is enumerated. Generating these functions in this initial stage dramatically reduces the computational burden of the second (main estimation) stage. In the second stage, the dynamic parameters are recovered using value function ap-

proximation coupled to a standard pseudo MLE routine. We enumerate, in greater detail, each step of the estimation process in what follows.

4.1 Estimation of the Static Parameters

The first stage of our estimation procedure involves recovering the parameters that govern the demand system (using the observed revenues and prices), inverting the first order conditions to infer the marginal costs of each firm, and then constructing a measure of variable profits for every market/firm combination observed in the full dataset.

We begin by assuming that an individual consumer i 's utility for a weekly basket of groceries purchased at chain f in week t is given by:

$$u_{ift} = \gamma^{pm} + \gamma_o^{pm} \cdot SC_f + \gamma_d^{pm} \cdot \frac{stores_{ft}}{pop_t} - \gamma_p^{pm} \cdot p_{ft} + \xi_f + \Delta\xi_{ft} + \varepsilon_{ift}^{pm} \quad (7)$$

where SC_f is a supercenter (firm) indicator, $\frac{stores_{ft}}{pop_t}$ is the “store density” of the firm (a proxy for travel distance), p_{ft} is the price of a basket of groceries, ξ_f is the latent quality of firm f , $\Delta\xi_{ft}$ is a transient shock to this quality, and ε_{ift}^{pm} is the usual Type I Extreme Value “demand shock”. Given this simple logit structure, the firm level market shares take on the familiar analytic forms, and the first order conditions of the static pricing problem can be inverted to obtain estimates of the marginal costs associated with each chain-level product, namely a weekly basket of groceries. The estimated costs are then projected down onto firm characteristics (state variables) as follows:

$$\ln(mc_{ft}) = \gamma^{mc} + \gamma_o^{mc} \cdot SC_f + \gamma_d^{mc} \cdot \frac{stores_{ft}}{pop_t} + \varepsilon_{ft}^{mc} \quad (8)$$

These estimates can be combined to compute variable profits, which are constructed as:

$$\pi_{ft} = (p_{ft} - mc_{ft}) \cdot s_{ft} \cdot pop_t \quad (9)$$

where the dependence of price, cost and market share on the current state is suppressed for brevity. Weekly profits are scaled up to the yearly level to match the decision frequency of the discrete game.

The demand system is estimated on the observed data using standard methods (i.e. 2SLS using Hausman-Nevo style instruments for price). One limitation is that we only observe prices for a subset of firms in a single period (2004) so the first stage of the usual 2SLS procedure is performed on this subset alone while the second stage uses the full set of firms and periods. Once demand estimation is complete, the system is then inverted to recover estimates of marginal costs (for the period and firms for which prices are observed). These costs are then projected down onto the two state variables and extrapolated out to the full set of firms and periods. The construction of variable profits then follows directly.

4.2 Definition of Aggregate States

Recall that our approach to approximating the value functions, which then feed into the computation of the structural CCPs, involves the use of both aggregate state “moments” (for dimension reduction) and a regression forest (for projection and interpolation, as well as dimension reduction). The set of aggregate states was chosen to reflect the key components of the natural state space that best capture both profits and valuations.

There are a total of 16 aggregate states or “features” used in our analysis. The full list of these states is displayed in Table 3, though we highlight a few here to illustrate how these aggregate states help characterize each firm’s potential flow profits and, thus, inform their future actions. First, we include a measure of chain-level profit, since this is the most salient indicator of how well a chain performs within a market. We construct chain-level profits from the demand and cost parameters recovered in our static demand estimation. In practice, different chains may arrive at the same profit outcome for quite distinct reasons. For example, a large chain operating in a small market versus a small chain operating in a large market could each earn the same flow profit at a given point, but face quite different long-

term prospects. Therefore, we also construct “aggregate states” around how much better or worse the focal firm’s market position is compared to the best and worst quality levels of other competitors in the market. Similarly, we calculate the differential in the number of stores the focal chain currently has versus the market’s biggest and small chains. We include these measures, along with the average levels of state variables, the population size, and the market’s growth rate as the remaining aggregates.

4.3 Profit Projection

Recall that a key component of our approximation procedure involves training a regression forest to capture the shape of the value function (through the forest’s partition and weight structure), which is then used for the structural estimation of the dynamic cost parameters. As a precursor to this exercise, we must first project our estimates of variable profits to many candidate points in the state space, including a large number that fall outside of the support of the data. To do so, we first create a large grid of points that includes both the data and several perturbations off its support. We then compute the exact variable profit outcome for every point in the grid (by solving for the Nash equilibrium of each market configuration) and fit a regression forest to these observed outcomes. Because market shares also appear as one of the aggregate states, we perform a similar procedure using the computed shares as the focal outcome. Further details regarding the estimation procedures, the resulting estimates and the overall fit are provided in the Appendix.

4.4 Estimating the Reduced Form CCPs

Estimation of the CCPs

As noted earlier, a key component to reducing the computational burden of estimating dynamic games involves using reduced form estimates of the CCPs to describe rival firms’ actions in a given market (employed when evaluating the integral on the right-hand side of equation (5)). Recovering the CCPs - the probabilities of each action for each player given

Table 3: Description of Feature Variables Used in Regression Forest Approximations

Feature	Description
Avg. Opp. Quality	Average quality metric for competitors in the market
Avg. Opp. Stores	Average number of stores owned by competitor chains in the market
Density	Focal Chain's number of stores divided by the population of the market
Grow	The Growth Rate of the market
Market Share	Calculated market share from the static demand model
Num SC	Number of SuperCenter chains in the market
Num SM	Total Number of SuperMarket chains in the market
Own Q - max Opp. Q	Own Quality Metric minus the highest quality metric for a competitor in the market
Own Q - min Opp. Q	Own Quality Metric minus the lowest quality metric for a competitor in the market
Own Quality	Own Quality Metric, obtained from the static demand estimation
Own Store - max Opp. Store	Total Number of Stores with the chain minus the number of stores held by the largest chain in the market
Own Store - min Opp. Store	Total Number of Stores with the chain minus the number of stores held by the smallest chain in the market
Own Stores	Total Number of Stores with the chain
Pop	Number of people in the market, scaled by 10,000
Profit	Calculated yearly profit from the static demand model
SC Dum	Dummy Variable if the Chain is a SuperCenter

the current state of the market - allows us to forward simulate the state of a market given the starting point.

We estimate the reduced form CCPs using flexible multinomial logit models, performed separately by chain type (supermarket or supercenter) and activity status (entrant or incumbent). In each logit model, we include, as appropriate, the 16 aggregate states to maintain consistency between the reduced-form CCPs and the structure of our other auxiliary models. In the supermarket incumbent case, we have to condition the flexible multinomial logit on the current number of stores operated by the chain during period t . Recall that firms are allowed to open stores, close stores, do nothing, or exit the market entirely. Thus, we condition the actions a supermarket incumbent can take on the current number of stores they are operating within the market. For example, if the supermarket incumbent has 1 store, then their options are: open two stores, open one store, do nothing, and close one store. We add flexibility to the multinomial logit by including the full set of aggregate states, as well as a quadratic term for the current period profit outcome.

Training the Forests

The final component of the set-up stage is training the regression forest that serves as the basis of the value function approximation utilized in the nested fixed point estimation of the dynamic parameters. We start with a simulated set of outcomes that serve as an approximate measure of the ex-ante value function of firm j , $\hat{V}_f^\Gamma(\mathcal{M}_{jft}^{\text{agg}})$, where $\mathcal{M}_{jft}^{\text{agg}}$ is the collection of aggregate states detailed earlier. This serves as the target variable upon which to train our regression forest.

These ex ante value function outcomes are approximated by a bootstrapped ensemble of regression trees, which can be represented as:

$$\hat{V}_f^\Gamma(\mathcal{M}_{jft}^{\text{agg}}) \approx \sum_{n=1}^N \alpha_n(\mathcal{M}_{jft}^{\text{agg}}) \hat{\lambda}_n \tag{10}$$

where

$$\alpha_n(X) = \frac{\sum_{b=1}^B 1(X \in R_{n,b})}{\sum_{n=1}^N \sum_{b=1}^B 1(X \in R_{n,b})} \quad (11)$$

and in which $1(X \in R_{n,b})$ is an indicator for being in region R_n (one of N partitions of the regressor space) in tree b , and $\hat{\lambda}_n$ is the simulated value for observation n that the forest is grown on initially. Here α_n is the data adaptive kernel, representing the number of times a given grid point described by the vector of $\mathcal{M}_{jft}^{\text{agg}}$ falls in a leaf, or region, with grid point n across all bootstrapped sample trees b in a forest of size B . The summation is then scaled by the total number of all grid points n in N that have appeared in the same region as $\mathcal{M}_{jft}^{\text{agg}}$ across a forest of size B . Thus, grid points that are closer in multi-dimensional space to those seen in $\mathcal{M}_{jft}^{\text{agg}}$ are given higher weight in creating the approximation of $\hat{V}_f^\Gamma(\mathcal{M}_{jft}^{\text{agg}})$ compared to those that have never resided in a leaf. Thus, this method also helps guard against the potential bias that may stem from outlying points on the simulated grid.

To mitigate overfitting bias and aid in convergence, our algorithm learns the “structure” of the value function outside the dynamic estimation. We start with simulating 2,400 markets for the supermarkets and 1,200 markets for the supercenters. In each market, we define one firm as the “focal” firm used to define each grid point. For each such firm (grid point), we calculate its profits in the current period and include the current guess at the structural cost parameters.¹⁰ Then, for each action, we compute the CCPs of all firms in the market, including potential entrants, and forward simulate actions of the other firms in the market, holding fixed the focal firm’s current action choice. We forward simulate out 10 periods, conditional on each potential action taken by the focal firm. In the 11th period, we assume the focal firm’s current profits become a perpetuity. Using the fixed discount rate, we discount all the cash flow streams back to the current period to approximate the continuation value of the specific choice. Therefore, we can then combine the current period profit, the cost of actions, and the approximate continuation value of each choice into an initial

¹⁰We start with initializing these costs to be zero and estimate the cost parameters in the second stage accordingly. For our final estimation, we update these costs with our prior estimates and proceed accordingly.

approximation of the choice specific value function for a given focal firm outside of estimation. We average the computed value function over 20 simulations. The regression forest is then grown on this approximate value function surface, with the full set of 16 aggregate states as our feature variables. Once estimated, we hold this structure constant across the final stage of estimation and, as outlined earlier, use the partitions and subsequent sample weights during our value function iteration/interpolation stage.

Finally, we enumerate and retain the set of one-period-ahead states for use in computing the integration over rival actions. Given the reduced form CCPs calculated earlier, we are able to probabilistically enumerate the occurrence of each possible future state for a given market. We retain any market combination that has a probability of occurring of at least .5% and discard the rest to reduce the computational burden.¹¹

4.5 Recovery of Dynamic Parameters

Given the various steps completed in the “Set-up” stage, we now proceed with estimating the structural parameters of the dynamic game. In so doing, we recover several cost parameters, conditional on both the chain-type (supermarket or supercenter) and its activity status in the market (entrant or incumbent). In addition, we parameterize the scale parameter as a function of market population and chain-type and recover the parameters indexing this component.

Our estimation procedure targets seven key structural parameters: opening costs of a store for an incumbent supermarket; closing costs of a store for an incumbent supermarket; opening costs of a store for an incumbent supercenter; entry cost, scaled by population, for an entrant supermarket; entry cost, scaled by population, for an entrant supercenter; and the two parameters that index the scaling factor. We collect these parameters into a single vector, θ .

Our final structural estimation occurs in the following steps. For a given value of θ_i , we

¹¹For the states retained; we update their probability of occurrence by re-scaling so all future probability for a given focal firm at time t sum to 1.

compute the flow profit for each point contained within our value function approximation grid. This requires updating the current guess of the value function based on the current value of θ and keeping the current period profit and beliefs regarding rivals' actions fixed. Once a flow profit is calculated with the new value of θ , we then solve for a new approximation of the value function for each point in our simulated grid using value function iteration and interpolation via the Bellman equation (equations (6) and (5)) along with the previously-calculated sample weights. Once the maximum difference between the current and prior iterations value function approximation is within a tolerance level, $1e^{-3}$ in our estimation, we stop updating the value function. We now have the mapping of future states to the guesses on our simulated grid of 2,400 points for the supermarkets and 1,200 points for the supercenters.

With these approximations of the value function defined for our grid points and a mapping between those points and all one-period-ahead states in hand, we then compute the choice specific value function for each action a chain j can take. These choice specific value functions are then a combination of: the current period profit for j , the cost of their choice (a component of θ), and the approximation of the value function based on the firm's choice. This allows us to compute the implied CCPs using equation (4).

We then apply a pseudo-maximum likelihood routine to match our implied CCPs to the actual choices in the data. If convergence is met through the use of an optimizer, the estimation stops, and θ are recovered sufficiently. If not, we start back at the initial part of this stage and cycle through the steps once again. Given that we use auxiliary functions to approximate constructs in our second-stage estimation, we apply the bootstrap, by markets, to compute standard errors.

5 Estimation Results

We now discuss our structural estimates, starting with the parameters that govern the static portion of the problem (demand, marginal costs and variable profits). We then turn to the dynamic parameters that characterize strategic investment.

5.1 Static Parameters

The parameter estimates of the conditional indirect utility function are as follows

$$\widehat{u}_{ift} = 1.64 + .063 \cdot SC_f + 5.60 \cdot \frac{stores_{ft}}{pop_t} - .049 \cdot p_{ft}, R^2 = .49$$

(.107)
(.013)
(.073)
(.001)

and include 18,889 observations over 1,762 firms. The corresponding marginal cost components are

$$\ln(\widehat{mc}_{ft}) = 4.32 - .237 \cdot SC_f - .301 \cdot \frac{stores_{ft}}{pop_t}, R^2 = .27$$

(.012)
(.015)
(.051)

and include 651 observations.

The parameters of the utility function indicate that consumers have a mild preference for the supercenter format and a strong preference for chains with a greater density of stores. The negative price coefficient corresponds to an average own elasticity of -3.53 and an average price-cost margin ($\frac{p-c}{p}$) of 28.3%. Note that these implied margins are closely in line with existing estimates from several studies.¹²

Substantively, the cost estimates imply that supercenters enjoy a strong cost advantage in the product market: they face marginal costs that are on average 23.7% below those of the non-supercenters. This advantage is somewhat mitigated by the impact of store density on marginal costs, since non-supercenters generally operate a larger number of stores in a

¹²Using detailed data on store level prices and costs for a U.S. grocery chain observed from 2004 to 2007, Stroebel and Vavra (2019) report average margins of roughly 31%. This corresponds quite closely with data from the Census' Annual Retail Trade Survey, according to which average gross margins for grocery stores (NAICS 4451) over our period are 28.5%. Using the Dominick's Finer Foods dataset for 1989-1994, Montgomery (1997) reports an average product-level gross margin of 25%. Finally, using data from the BLS, Nakamura (2008) reports store level margins of 28.3%.

given market.

5.2 Dynamic Parameters

Recall that the deterministic portion of the cost of opening new stores is given by $\theta_o^{C^{open}}$, which depends on the chain’s format o . A chain that opens one store receives $\theta_o^{C^{open}}$, while a chain that opens two receives $2 \cdot \theta_o^{C^{open}}$. Similarly, chains that close stores receive a payoff $\theta_o^{C^{close}}$, which is similarly incremented from one to two. Finally, firms that exit receive $\theta_o^{C^{close}}$ for each store they close when exiting (i.e., one increment for each remaining outlet).¹³ These costs are augmented further if the chain is an entrant, as they have to pay not only the cost of opening stores, but also an upfront entry cost. This entry fee reflects the cost of the firm setting up a network within the region and is scaled by the population (size) of the market. These entry costs are further differentiated by the chain’s format, either supermarket or supercenter.

We estimate these dynamic cost parameters on a subset of the data, including only markets with less than 350,000 consumers and fewer than 8 firms in operation. This restriction is imposed for two reasons. First, the set of large markets is very sparse and heterogeneous: the US has relatively few “large” cities. Pooling the large markets together with the more common mid-size cities would likely lead to misspecification. The second reason is simply numerical tractability: enumerating the possible actions of 8 or more players with continuous variables remains intractable. As a result, our model focuses on 114 markets across a common time horizon. This results in over 6,000 observations in total for incumbents, where supermarkets naturally appear much more frequency than supercenters. All coefficients, except those attached to the scaling factor, are in \$10 million dollar increments. All standard errors are computed by drawing 100 bootstrapped samples (at the market-level), where the first stage is kept fixed, and we re-sample the observations used in the dynamic estimation. The table below contains the full set of parameter estimates.

¹³Due to data limitations, we cannot estimate a closing/exit cost for the supercenters. Therefore, these are set to zero in our analysis.

Table 4: Estimates of Dynamic Structural Parameters

Variable	Estimate	Std. Error	T-stat
Supermarket Build One Store	-22.498	0.524	-42.915
Supermarket Close One Store	3.162	0.257	12.281
Supercenter Build One Store	-14.046	0.871	-16.132
Supermarket Entry Cost (scaled by population)	-0.16	0.037	-4.371
Supercenter Entry Cost (scaled by population)	-2.812	0.28	-10.05
Sigma - Log Population Size Coefficient	0.543	0.013	40.728
Sigma - Supercenter Dummy Variable	0.432	0.13	3.332

For the supermarket chains, the cost of an incumbent opening a single store is approximately \$225 million dollars (SE 5.24). If a supermarket chooses to close a store, they recover approximately \$31.6 million dollars (SE 2.57), or rather, the salvage cost of a supermarket is approximately one-seventh the cost of opening. This salvage value represents both the cost of liquidating the store, selling off any remaining items for salvage, and the discounted cash flow savings of not employing the staff for that location.

For the supercenter firms, the cost of an incumbent opening a single store is approximately \$140 million dollars (SE 8.71). Note that this cost is over \$70 million dollars lower than their supermarket counterparts. This superior cost position may reflect the fact that supercenters are often built on already established outlets as part of a conversion process. This pre-existing structure represents a considerable cost advantage when opening a supercenter. Recall that supercenters also face a much better variable cost position, revealing that their margins on products sold are substantially higher than supermarkets.

Part of the supercenter cost advantage comes from meticulous planning and upfront costs established when the supercenter enters a market. We see that the cost per 10,000 people for a supermarket to enter a new market is approximately \$1.6 million dollars (SE 0.37). On the other hand, the supercenter's fixed cost of entry is approximately \$28.1 million dollars per 10,000 people (SE 2.80), a much larger outlay. This upfront investment can be rationalized as the supercenter building their network and distribution hub within the region. Once established, they can exploit the cost savings borne from this upfront, fixed costs to enjoy better margins and less costly store openings.

Finally, we see there is heteroscedasticity with respect to population size. As the population increases, there is much higher variation in the flow payoffs of the firms, the coefficient on the (logged) population is .543 with standard error 0.013. We also note that supercenters have a larger scale parameter than supermarkets, holding the population of a market fixed, as the dummy variable for supercenters is .432 with standard error 0.13.

Before proceeding with a counterfactual analysis to uncover how the new format has impacted the market overall, it is worth discussing the degree to which our estimates of these structural costs coincide with expectations. Given that we have normalized fixed operating costs to zero, the cost per store should be interpreted as the store entry cost plus the present discounted value (PDV) of fixed costs for the lifetime of the store (Aguirregabiria and Suzuki, 2014). Recall that the estimated value for supermarkets is \$225 million, while the estimate for supercenter firms is a considerably smaller \$140 million. Note that these are averages, not realizations conditional on choice. Factoring in the scale parameter and the relative frequency with which entry occurs, the typical cost paid would be closer to \$170 million for supermarkets and \$106 million for supercenters. The average scrap value for a supermarket is a more modest \$33 million, with the typical price received closer to \$73 million. Finally, de novo entry adds \$100 million to the cost of building the initial stores for supermarkets, and \$500 million for supercenters.

How reasonable are these numbers? Note that the average entering supermarket makes approximately \$13.7 million in revenue per year. Assuming a 30% gross margin and a 5% discount rate, this translates to a PDV of about \$82 million, about half as large as the value reported above. The 75th and 90th percentiles for the PDV are \$120 and \$162 million, respectively, which are closer. For supercenters, the mean, the 75th, and the 90th percentiles are \$170 million, \$200 million, and \$220 million, respectively. The scrap value is only identified for supermarkets (since supercenters rarely closed during our data window). Note that closed stores are typically much lower revenue than newly opened ones. For example, the PDV for the average closing supermarket is \$57 million, and the 25th percentile

is \$47 million.

6 Counterfactuals

From a policy perspective, it is informative to consider how the grocery market would appear today if the supercenter format had never existed. Many articles have established that Wal-mart, and by extension, the supercenter format results in a shift in the dimensions of competition. One salient point in these articles is that smaller firms tend to suffer more from the existence of supercenters than larger firms. Another is that incumbent supermarkets exhibit little reaction to Walmart’s low price position. Using the results of our estimation above and conducting counterfactual analysis, we can extrapolate how markets would have evolved if the supercenter format had not existed across our 13-year panel.

To compute valid counterfactual comparisons, we solve for a new set of equilibrium behaviors that reflect a world in which the supercenter format did not exist. In practice, this means re-solving the full set of structural CCPs for all firms. To do so, we follow a procedure that mimics how the estimation took place but with three main differences. First, we consider a smaller set of “representative” markets (and a modified grid of points on which to fully solve the model). Second, we now solve the full doubly nested problem of optimal dynamic behavior coupled with equilibrium best response. Third, we hold the structural parameters fixed at their estimated values.

We start by separating the markets’ population size into 50,000 person increment blocks and simulating 500 grid points for 5 representative market size. For each grid point, we perform the same exercise as in our set-up step, namely, generating an initial guess of each firm’s approximate value function. With the initial set-up, we use the CCPs of the data to initialize the forward-looking beliefs. These CCPs are potentially contaminated by the presence of supercenters. Thus, after we complete the counterfactual step once, we do so again using the approximate counterfactual CCP’s to re-simulate the beliefs and update our

counterfactual CCPs to eliminate the presence of supercenters.

Once there is a defined grid for each population block, we grow a regression forest to describe the response surface of the value function in that population range. Last, we proceed in enumerating the one-period ahead states associated with each point in the grid. The sum total of these steps is the collection of aggregate states, the forest associated with the value function response surface, and the full set of one-period ahead states.

Using the structural estimates and above constructs, we then estimate the choice specific value function of each observed firm based on the current guess at the value function. The choice specific value functions then naturally translate into a calculated, choice-specific probabilities for a given firm in a given market composition. We then, recursively update the probabilities for all firms in the market, and any potential entrants, until the maximum difference between the full set of choice specific probabilities before fully updating and the current updated values is within a tolerance, $1e^{-7}$.

Once each of the selected markets has reached an equilibrium in the probability space, we calculate a new approximation of the value function for each firm. We compare the vector of calculated value function values of the current iteration to the last iteration. If the maximum difference between the last iteration function approximations and the current iterations are within a $1e^{-5}$ tolerance, we stop. Otherwise, we proceed at the start in finding an equilibrium in the probability space.

Once convergence is achieved between successive iterations of the value function approximation, we collect all the choice specific probabilities from all firms across the five population blocks. These probabilities serve as the left-hand side variables, with the right-hand size variables being our aggregate states. Using non-linear least squares, we map the counterfactual choice specific probabilities to the aggregate states for parsimony with our original CCP estimates. We complete this whole process twice to ensure that our counterfactual CCPs contain no trace of bias from the actions taken under the premise that supercenters exist.

Armed with both actual and counterfactual CCPs, we then forward simulate 13 periods

for each of our representative markets and report aggregate statistics for the final period. In total, we forward simulate each market 250 times and then compute summary statistics of interest. The summary estimates are presented in the series of tables below. We begin by examining the evolution of the overall market structure. Next, we show the differences in pricing and profitability. Third, we highlight changes in share, quality, and consumer surplus. Last, we comment on the differences in supermarket behavior, given the two scenarios.

Table 5: Overall Changes in Market Structure

Metric	No Walmart	With Walmart
Number of Chains	5.297	4.883
Total Number of Stores	14.916	15.163
Total Market Density	0.94	0.945
Average Number of Stores per Chain	2.862	3.319
Average Density per Chain	0.189	0.206
Population	16.714	16.714
Number of Supercenters	0	0.763
Number of Supermarkets	5.297	4.12

We are interested in both how supercenters have affected market-level competition and consumer surplus. In congruence with other studies that have focused on a more reduced form approach, we find that supercenters cause the remaining supermarket firms to be larger and, thus, more profitable on average. The overall store-level density is almost the same in both scenarios. The major difference comes from the number of chains in the market. In the absence of a supercenter, there are approximately .4 more chains within a market on average, 5.297 chains on average in the markets where supercenters do not exist versus 4.883 chains in the markets where supercenters exist. That extra chain is, generally, of a smaller size, as the average number of stores per chain in the scenario where supercenters do not exist is 2.862 versus 3.319 in the other scenario. Last, we see that supercenters have not yet fully populated all the markets, with an average number of supercenter chains being .763. Thus, even the threat of a supercenter entry can stave off other supermarkets from entering a market, which explains why the number of supermarket chains in a market is 4.12 versus 5.297 when supercenters do not exist.

Table 6: Overall Changes in Profitability and Pricing

Metric	No Walmart	With Walmart
Total Market Profits	11.129	12.244
Avg. Profit by Chain	2.19	2.799
Avg. Profit by Store	0.72	0.799
Avg. Market Price	94.29	92.263
Minimum Price	92.279	84.48
Maximum Price	95.916	96.542
Avg. Margin of Chains	0.246	0.258
Minimum Margin	0.224	0.226
Maximum Margin	0.288	0.313

Beyond how the market structure has changed with the supercenter format’s presence, we can also identify the effect on pricing, profitability, and margins within markets. In the scenario where supercenters exist, the overall market profitability is higher than when they do not exist by approximately \$11 million dollars (\$122.4 million versus \$111.3 million). A more stark difference occurs at the firm level, where the chain level profits are \$ 6 million dollars higher (approximately \$28 million dollars per period versus \$22 million dollars). One major contributing factor is the increase in store density outlined earlier, since density has the benefit of increasing the attractiveness of a chain’s product. The expansion of a chain has another effect; it reduces their marginal costs. In the scenario where supercenters exist, the average and maximum margins in the market are higher (25.8% and 31.3%, respectively). Some of this increase is due to supercenters having a higher average margin than supermarkets, yet, the tendency of the remaining supermarket chains to be larger influences this number as well. Due to better margins, we see the average weekly basket price in the market is substantially lower when supercenters exist, \$92.26 versus \$94.29. The lower number is primarily driven by the supercenters charging a lower price, \$84.48, as the average minimum price in a market. More interesting is that some supercenters are now allowed to charge higher prices, where the maximum price in a market is \$96.54 versus \$95.92. Vertical differentiation brought about by the format change allows the larger, more high-quality supermarkets to actually increase their prices over the competition and make more profits.

This finding is in stark contrast to most conventional wisdom, as we would normally expect all prices to drop due to the presence of supercenters in the market.

Table 7: Overall Changes in Consumer Surplus, Quality, and Share

Metric	No Walmart	With Walmart
Total Market Share	0.528	0.537
Avg. Share by Chain	0.108	0.118
Share of the Outside Good	0.472	0.463
Avg. Chain Quality	0.073	0.041
Minimum Chain Quality	-0.592	-0.601
Maximum Chain Quality	0.639	0.606
Consumer Surplus	15.883	16.439
Consumer Surplus by Chain	3.248	3.647
Consumer Surplus by Store	1.179	1.22

Over time, we see that the overall share of chains within a market is relatively unchanged. This is a factor of population growth being faster than the rate at which firms open new stores. That said, the inclusion of a supercenter in the scenario increases the total market share coverage of chains by almost 1% more (53.7% versus 52.8%). Given that there are fewer chains in the scenario with supercenters, each chain serves almost 1% more of the market, where the average chain level share is 11.8% versus 10.8%. We also see that the two scenarios' quality levels are virtually the same, with the no supercenter scenario having an average quality level of .073 versus .041. In terms of consumer surplus, the ultimate metric that codifies how the consumer market views the inclusion of the new technology, we see that the scenario with supercenters wins out with a market-level consumer surplus of 16.439 versus 15.883. The combination of lower average prices and a more favorable format (the supercenter) in a market wins out compared to more chain options and slightly higher quality.

Last, we consider just supermarkets in both scenarios to see how they behave differently in the presence of the supercenter format. On average, the stores that survive a supercenter's entrance are marginally lower quality but substantially larger. Even being larger, there is virtually no change in the supermarket chain's average price in either scenario. The

Table 8: Overall Changes in Supermarket Behavior

Metric	No Walmart	With Walmart
Avg. Price	94.29	94.236
Avg. Chain Profit	2.19	2.731
Avg. Chain Share	0.108	0.113
Avg. Chain Margin	0.246	0.249
Avg. Chain Number of Store	2.862	3.52
Avg. Chain Quality	0.073	0.042

margins are a little better for the supermarkets in the scenario where supercenters exist, due to having more stores on average, and thus scale economies. The combination of the pricing remaining relatively unchanged, a slightly better margin and average share results in an overall increase in supermarkets' average profitability, upon remaining in a market to compete with supercenters. These chains are large enough to compete directly with supercenters. On the other hand, smaller, higher-quality chains quickly exit the market, as these offerings cannot compete in the new environment.

7 Conclusion

This paper proposes and estimates a model of dynamic competition in the supermarket industry. Using recently developed two-step estimation techniques, we recover the structural parameters governing demand, pricing decisions, incremental investment, and entry costs. We then use these parameter estimates to evaluate policies aimed at eliminating Supercenters.

Our analysis finds that in the absence of Wal-mart, the average market grows in terms of number of firms, but the overall average density is reduced. The presence of Wal-mart creates a blockade that dissuades entry of new firms and actually causes supermarkets to exit the market. However, with the exclusion of the possibility of entry of Wal-mart, we find strong evidence that prices increase, not only the average basket price, but the maximum basket price observed in the market as well. Therefore, Wal-mart can be seen as a boon to

consumer surplus on the dimension of price, but not necessarily on the dimension of store choice.

Further, in this paper we outline a new estimation methodology, leveraging advances in machine learning, to estimate models that include dynamics. Utilizing data-adapted kernel weighted forests, we are able to calculate a continuation value in our dynamic setting, despite the richness of the data and large state space. We propose that this technique can offer future researchers in this area an avenue to pursue other questions using dynamics that were not possible with previous techniques.

In this paper, we focus on one sole dimension of competition, changing quality through store density. However, one could imagine that chains have a few different marketing levers that could be pulled to differentiate themselves. One such avenue for future research is allowing firms to simultaneously change by enhancing their position via density of stores and advertising to change their unobserved quality signal. By allowing chains both of these options, we hypothesize that there may be other avenues that firms are using to compete with the dominant technology, beyond the escalation of chain size.

References

- Aguirregabiria, V., A. Collard-Wexler, and S. P. Ryan (2021). Dynamic games in empirical industrial organization. In *Handbook of Industrial Organization*, Volume 4, pp. 225–343. Elsevier.
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75(1), 1–53.
- Aguirregabiria, V. and J. Suzuki (2014). Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics* 12(3), 267–304.
- Arcidiacono, P., J. Blevins, P. Bayer, and P. B. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *Review of Economic Studies* 83(3), 889–931.
- Arcidiacono, P. and P. B. Ellickson (2011). Practical methods for estimation of dynamic discrete choice models. *Annu. Rev. Econ.* 3(1), 363–394.
- Arcidiacono, P., P. B. Ellickson, C. F. Mela, and J. D. Singleton (2020). The competitive effects of entry: Evidence from supercenter expansion. *American Economic Journal: Applied Economics* 12(3), 175–206.
- Atkin, D., B. Faber, and M. Gonzalez-Navarro (2018). Retail globalization and household welfare: Evidence from Mexico. *Journal of Political Economy* 126(1), 1–73.
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75(5), 1331–1370.
- Barwick, P. J. and P. A. Pathak (2015). The costs of free entry: An empirical study of real estate agents in greater Boston. *Rand Journal of Economics* 46(1), 103–145.

- 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.
- Bellman, R., R. Kalaba, and B. Kotkin (1963). Polynomial approximation—a new computational technique in dynamic programming: Allocation processes. *Mathematics of Computation* 17(82), 155–161.
- Benitez-Silva, H., G. Hall, G. J. Hitsch, G. Pauletto, and J. Rust (2000). A comparison of discrete and parametric approximation methods for continuous-state dynamic programming problems. Working Paper: Yale University.
- Berry, S. (1994). Estimating discrete choice models of product differentiation. *Rand Journal of Economics* 25(2), 242–262.
- Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica* 81(3), 1003–1037.
- Crawford, G. and M. Shum (2005). Uncertainty and learning in pharmaceutical demand. *Econometrica* 73(4), 1137–1173.
- DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in US retail chains. *The Quarterly Journal of Economics* 134(4), 2011–2084.
- Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in IO. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization*, Volume 3, Chapter 30, pp. 1887–1966. Amsterdam: Elsevier.
- Dunne, T., S. Klimek, M. J. Roberts, and Y. Xu (2013). Entry and exit in geographic markets. *Rand Journal of Economics* 44(3), 462–487.
- Ellickson, P. B. (2006). Quality competition in retailing: A structural analysis. *International Journal of Industrial Organization* 24(3), 521–540.

- Ellickson, P. B. (2007). Does Sutton apply to supermarkets? *Rand Journal of Economics* 38(1), 43–59.
- Ellickson, P. B. and P. L. Grieco (2013). Wal-mart and the geography of grocery retailing. *Journal of Urban Economics* 75, 1–14.
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62(1), 53–82.
- Foster, L., J. Haltiwanger, and C. J. Krizan (2006). Market selection, reallocation, and restructuring in the US retail trade sector in the 1990s. *The Review of Economics and Statistics* 88(4), 748–758.
- Fowlie, M., M. Reguant, and S. P. Ryan (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy* 124(1), 249–302.
- Handbury, J. (2021). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Econometrica* 89(6), 2679–2715.
- Hausman, J. and E. Leibtag (2007). Consumer benefits from increased competition in shopping outlets: Measuring the effect of Wal-Mart. *Journal of Applied Econometrics* 22(7), 1157–1177.
- Hitsch, G. J., A. Hortacsu, and X. Lin (2021). Prices and promotions in US retail markets. *Quantitative Marketing and Economics* 19, 289–368.
- Hollenbeck, B. (2017). The economic advantages of chain organization. *The RAND Journal of Economics* 48(4), 1103–1135.
- Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica* 79(1), 253–302.
- Hothorn, T., B. Lausen, A. Benner, and M. Radespiel-Tröger (2004, 01). Bagging survival tree. *Statistics in medicine* 23, 77–91.

- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60(3), 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(2), 265–289.
- Houde, J.-F., P. Newberry, and K. Seim (2023). Nexus tax laws and economies of density in e-commerce: A study of Amazon’s fulfillment center network. *Econometrica* 91(1), 147–190.
- Judd, K. L. (1998). *Numerical methods in economics*. MIT press.
- Kalouptsi, M. (2018). Detection and impact of industrial subsidies: The case of Chinese shipbuilding. *The Review of Economic Studies* 85(2), 1111–1158.
- Kalouptsi, M., P. T. Scott, and E. Souza-Rodrigues (2021). Identification of counterfactuals in dynamic discrete choice models. *Quantitative Economics* 12(2), 351–403.
- Keane, M. P. and K. I. Wolpin (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence. *Review of Economics and Statistics* 76(4), 648–672.
- Lagakos, D. (2016). Explaining cross-country productivity differences in retail trade. *Journal of Political Economy* 124(2), 579–620.
- Meinshausen, N. and G. Ridgeway (2006). Quantile regression forests. *Journal of Machine Learning Research* 7(6).
- Montgomery, A. L. (1997). Creating micro-marketing pricing strategies using supermarket scanner data. *Marketing Science* 16(4), 315–337.
- Nakamura, E. (2008). Pass-through in retail and wholesale. *American Economic Review: Papers and Proceedings* 98(2), 430–437.

- 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(2), 373–399.
- Pesendorfer, M. and P. Schmidt-Dengler (2007). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75(3), 901–928.
- Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5), 999–1013.
- Rust, J. (1994). Estimation of dynamic structural models, problems and prospects: Discrete decision processes. In C. Sims (Ed.), *Advances in Econometrics: Sixth World Congress*, Volume 2, Chapter 4, pp. 119–170. Cambridge University Press.
- Ryan, S. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica* 80(3), 1019–1061.
- Smith, H. (2004). Supermarket choice and supermarket competition in market equilibrium. *Review of Economic Studies* 71(1), 235–263.
- Stroebel, J. and J. Vavra (2019). House prices, local demand, and retail prices. *Journal of Political Economy* 127(3), 1391–1436.
- 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(5), 1763–1803.
- Thomassen, O., H. Smith, S. Seiler, and P. Schiraldi (2017). Multi-category competition and market power: A model of supermarket pricing. *American Economic Review* 107(8), 2308–2351.

Wager, S. and S. Athey (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523), 1228–1242.

8 Appendix

8.1 Details of Profit and Share Projection Procedures

Given the approximation of profits and market shares from the prior step, we then use machine learning to create the necessary auxiliary functions. The calculation of chain-level profits would require solving the Nash equilibrium in prices for each chain within a specific market configuration, which is computationally burdensome. To alleviate this burden, we use a regression forest trained on the estimates from the prior step to then construct profits for all market combinations observed in the data that are used in our final estimation. Using a regression forest gives us an advantage in this stage as it allows for non-parametric relationships between our 16 “approximation states” to enhance the prediction accuracy of both profit and market share.

In the subsection, we enumerate all the features used in our two regression forests, one for chain market share and the other for chain profits. Given that we use a regression forest to link the state aggregates to the calculated profits and share, we cannot directly show a set of coefficients. We present two tables below summarizing the importance ratings and fit statistics for each model in place of a set of coefficients. These tables show evidence that the regression forests can map the state features to the outcome variable sufficiently well. The importance ratings measure each feature’s frequency within the first four levels as a splitting variable for a tree. Thus, variables with greater frequencies are deemed to be more important as they help minimize the gap between predicted and observed variables more efficiently.

In both cases, the fit statistic is 98% or greater. Of note, each predictive model has a different ordering of features based on the importance rating. For the market share regression forest, the top three features based on the variable importance measure are: density, own quality, and the difference between how many stores the focal chain has versus the largest competitor chain in the market. Density largely determines the utility of the firm’s product

Table 9: Market Share Regression Forest Feature Variable Importance Ratings

Feature Variable	Importance Rating
Density	0.5868
Own Store - max Store	0.1753
Own Quality	0.0955
Own Store - min Store	0.0382
Own Q - max Q	0.0363
Own Stores	0.0291
SC Dum	0.0118
Opp Stores	0.009
Num SM	0.0064
Own Q - min Q	0.0051
Opp Quality	0.0022
Pop	0.0021
Num SC	0.0011
Grow	0.001
r squared	0.9794

Table 10: Profit Function Regression Forest Feature Variable Importance Ratings

Feature Variable	Importance Rating
Market Share	0.6114
Own Stores	0.1912
Own Store - min Store	0.0796
Pop	0.0516
Own Store - max Store	0.0197
Density	0.0154
Own Quality	0.0121
Own Q - max Q	0.0072
Num SM	0.0034
Opp Stores	0.0028
Grow	0.0024
SC Dum	0.0015
Own Q - min Q	0.0014
Opp Quality	0.0002
Num SC	0.0001
r squared	0.9936

in the static demand model. However, density on its own does not take into account the relative size of the focal firm compared to the competition; thus the differential in chain size is an important determinant in predicting market share. Last, the recovered quality metric recovered from the static demand analysis further refines market share prediction.

In the profit regression forest, the top three features based on variable importance are: market share, own stores, and the differential between how many stores the firm has compared to the smallest chain in the market. Market share is the largest determinant as it directly affects the chain's product market profitability. However, the share prediction is further refined by the focal chain's size, both in terms of the number of stores it owns and how many more than it compared to the smallest chain in the market. Also of note, the market population is the fourth most important variable, which is another variable that directly impacts the chain's profitability.

From both of these exercises, we see that the total number of stores within the chain, either directly or indirectly through the density measure, impacts the share and firm profitability. This finding gives evidence that our dynamic model is focused on a variable, the number of stores in the chain, that is important to chain managers. Had the chain size not been one of the most important variables, choosing this as the action variable in our analysis would be questionable. However, we see that the chain's size directly impacts the prediction of chain-level profitability.

8.2 Generation of grid points to generate the value function response surface

In both the main estimation and counterfactuals, we rely on generation of a set of grid points to grow the random forest that describes the surface of the value function. A point on the grid represents a given focal firm within a market of other players. As such, we simulate from the data not only characteristics of the focal firm (type, store density, and quality), but also other players in the same market.

Following Sweeting (2013), we initialize this process by randomly sampling markets from the data. This ensures that our grid is supported by observable outcomes and not generate hypothetical grid points outside of normal market evolution. Once a market is drawn, one firm, depending on the forest that is being grown (supermarket or supercenter) is designated as the focal firm. All other firms in the market aide in generating the aggregate states.

After the initial draw of the market, we perturb the market by randomly assigning firms to do nothing, open, or close a store. This step also includes the focal firm. Further, we randomly allow new entrants, both supermarket and supercenters in the case of the main estimation grid, to enter the market. Additionally, we adjust the population of the market by its growth rate to add additional perturbations to the market. These incremental changes occur 4 times to give additional saturation in the state space for the forest to grow on and provide a finer grid for interpolation during estimation.

After collecting this initial set of grid points, we then need to create an approximation of the value function to serve as the dependent variable for the forest. Using the CCPs from the data (or from the counterfactual CCPs depending on the use of the grid), we forward simulate the actions of all the players in the market, and potential entrants, for 10 periods. In each simulated period, we calculate and store the current profits for the focal firm. After the 10th period, we calculate the profits one last time and assume these profits are generated into perpetuity. Armed with these values, we discount each payout to the present day as an approximation of the value function for the firm at the initial states. We forward simulate 20 times and average the estimates of the value function to create our final dependent variable to anchor the growth of the forest that describes the response surface of the value function based on the set of aggregate states.

For the grid used in estimation, we assume that the structural costs of actions are 0 in the first round of generating the forest. In the second round of estimation, we include the estimated structural parameters of the firm's actions to create a closer approximation of the value function.

8.3 Enumeration of estimation steps

Here, we provide a comprehensive overview of the steps taken to estimate our structural parameters. To achieve the goal of estimating the structural parameters of our model, we proceed with the following steps:

- 1 Estimating the yearly chain-level profit and market share
 - 1.1 Using the pricing data from 2004, apply a Hausman-Nevo style 2SLS estimation to obtain the market-level demand parameters
 - 1.2 Given the prior estimation, invert the system to obtain estimates of coefficients for the marginal cost function
 - 1.3 Combining the market-level demand parameters and marginal cost function, using the first order conditions to estimate the basket prices, marginal cost, and market share of each firm in a given market for all observations in our data set.
 - 1.4 Using the price, marginal cost, and market shares, construct the per-period profit in a market by multiplying the profit margin with its share, the population size of the market, and the number of weeks in the year.
- 2 Estimating conditional choice probability functions
 - 2.1 Using the data in our application, generate four different multinomial logits, where the independent variables are the set of aggregate states
 - 2.1.1 A model for incumbent supermarkets, a model for incumbent supercenters, a model for entrant supermarkets, and a model for entrant supercenters
- 3 Estimating profit and market share functions
 - 3.1 To construct aggregate states for the one-period ahead state-space, we need a flexible approximation of profit and market share
 - 3.2 Using the estimates from Step 1 (profit and market share), train an optimally tuned random forest on each outcome variable separately
 - 3.2.1 Each forest uses as feature variables the aggregate states of the market
- 4 Generating the value function response function via random forest
 - 4.1 Start by drawing 480 markets for supermarkets and 240 for supercenters (a market at a given year) and randomly determine one firm in the market as the focal firm.
 - 4.2 Perturb each selected market by randomly adding or removing stores from either the focal firm or other firms in the market, as well as changing the population by its growth rate. Also, randomly have supermarket or supercenters enter the market with randomly drawn quality metrics.

- 4.3 For each focal firm (now 2400 supermarkets and 1200 supercenters), forward simulate the evolution of the market to obtain the approximation of the value function for each point on our grid.
- 4.4 Create an optimally-tuned random forest with the dependent variable being the approximation of the value function obtained in the prior step and the feature variables being the aggregate states calculated for each point on the grid.
 - 4.4.1 This forest is stored as an object for us to obtain the data adaptive kernels.
- 5 Enumeration of one-period ahead state space (for both the value function approximation grid and the observations used in estimation)
 - 5.1 For each chain, enumerate all possible actions (adding stores, closing stores, or doing nothing) by competitors (and potential) entrants in the market, conditional on their type and market constraints.
 - 5.2 For each action of a chain, calculate its individual probability of taking that action using the aggregate states of that chain and the CCPs.
 - 5.3 Multiply each independently calculated probability of a given action with all the other enumerated possibilities to obtain the probability of the focal firm seeing the market evolve to that state one period from now.
 - 5.4 Retain any one-period ahead state where the probability of occurring is greater than .5%. Calculate the total probability of seeing all the retained one-period ahead states and re-weight the observations so they sum to 1.
 - 5.5 For all retained one-period ahead states, calculate the aggregate states. Use these aggregate states with the random forest obtained in 4.4 to get the vector of data-adaptive kernels that marry the one-period ahead states with the grid obtained in 4.2.
- 6 Estimation of structural parameters θ
 - 6.1 Start with a guess of θ
 - 6.2 For θ , start with the simulated grid
 - 6.2.1 For each point on the grid, calculate the current estimate of the value function.
 - 6.2.1.1 The value function is calculated by matrix multiplication of the current guess of each grid point's value function estimate and the data adaptive kernel matrix.
 - 6.2.2 Compare this iteration's estimate of the value function with the last iteration's
 - 6.2.3 If the maximum difference between the value function estimates is less than $1e^{-5}$, stop. Otherwise, return to the calculation step.
 - 6.3 Using the converged value of the value function for each grid point, proceed with estimating the choice specific value function for each observation in our data set.
 - 6.4 Translate the choice specific value functions into the logit probabilities for each firm in the market. Retain the probability associated with the firm's action.

6.5 Optimize the overall objective function, using Nelder-Mead, of the pseudo-maximum likelihood function using the log sum of the retained probabilities

Once θ is obtained, complete steps 4 through 6 again to obtain the final estimates of θ . Standard errors are obtained via bootstrap.