

Taste versus Space: Demand Heterogeneity in the Grocery Industry*

Paul B. Ellickson^a, Paul L. E. Grieco^b, and Oleksii Khvastunov^b

^aSimon School of Business, University of Rochester

^bThe Pennsylvania State University

March 28, 2014

Abstract

The grocery industry is dominated by chains, both national and local, competing in product and geographic space. While grocery stores have historically been relatively homogeneous—offering a similar variety of products and differentiating primarily by geographic location—in recent years several firms have entered regional markets with alternative product offerings. These include “limited assortment” chains such as Aldi, “high-end” chains such as Whole Foods, and “supercenter” firms such as Wal-Mart. This segmentation suggests that product offerings represent an alternative avenue for differentiation beyond simple geography, which has traditionally been understood to drive market structure in this industry. Using a census of store level data on revenue, location, size, and chain affiliation, we assess the extent to which firms cater to observable consumer sub-groups and explore the demand-side implications of this segmentation. We develop a model of individual store choice based on census-tract level consumer demographics, that allows for rich substitution patterns between stores. These patterns are driven by both demographic variation and physical location. In addition to being of interest in its own right, this store choice model represents an essential ingredient to a supply-side model that can be used to recover estimates of the cost implications of location and product-line decisions.

Keywords: Grocery Retail; Supercenters, Store Choice, Demand Estimation.

1 Introduction

The grocery industry was long dominated by a handful of powerful chains. Leveraging the explosive growth of mass media and national brands in the 1950s and 1960s and the great migration to the suburbs, strong regional and sometimes national chains invested heavily in information technology and large-scale stores to efficiently stock a wide variety of products. By the early 1990s, this escalating investment in endogenous sunk costs (Sutton, 1991) led to a concentrated industry structure in which every regional market was dominated by a small number of large-scale chains (Ellickson, 2007). At the local level, the resulting equilibrium structure was remarkably symmetric, both in terms of market share and product positioning. The heavy investments required to stock a massive selection of products served as an effective barrier to entry that maintained the status quo.

*Very preliminary, comments welcome. Please send correspondence to paul.ellickson@simon.rochester.edu (Ellickson), paul.grieco@psu.edu (Grieco) or ouk111@psu.edu (Khvastunov).

However, this structure began to shift starting in the early 1990s. One source of disruption was the entry of Wal-Mart into grocery retailing through the introduction of “supercenters”—extremely large stores that sell both groceries and other mass merchandise. Wal-Mart had gained prominence as a mass-merchandiser using similar investments in IT, bandwidth, and distribution that grocery chains had employed, but on a much larger scale (Holmes, 2001). In hindsight, the pivot to supermarkets was natural for Wal-Mart, as they already owned a large network of big box stores and operated a highly efficient distribution network. Through the 1990s and early 2000s, Wal-Mart grew to be the nation’s largest supermarket firm by number of stores and total revenue, mainly at the expense of older, mid-tier supermarket chains who had failed to update their store profile and distribution infrastructure. As Wal-Mart grew, the competitive structure of the industry shifted in other ways as well. At one end, some firms like Aldi and Save A Lot have traded down market, focusing on consumers with lower willingness to pay for national brands and tighter travel constraints due to a reliance on public transport. At the other end, firms like Whole Foods, Trader Joes and Wegmans have pivoted up market to focus on wealthier consumers who crave more specialized offerings and have higher travel costs. As a result, what was once a relatively homogeneous industry, has appears today to have become far more differentiated and fragmented. Our first goal in this paper is to empirically document this fragmentation.

Once we establish the importance of differentiation within the grocery industry, an obvious question is what are the underlying reasons for such a shift in market structure? One possibility is that these (and other) firms have been forced into less profitable segments by the entry of a dominant, low cost competitor in the main segment (i.e., Wal-Mart). A second hypothesis is that consumer demographics have shifted in ways that make differentiation more profitable—either because consumers are sorting into less heterogeneous neighborhoods or because rising inequality has made marketing too “non-average” consumers more profitable. Relatedly, it may be that consumer tastes have grown more heterogenous over time. Of course, all of these mechanisms could be in play simultaneously.

The long-term goal of this research is to evaluate these hypotheses by first building a flexible model of store choice by heterogeneous consumers facing a differentiated set of products (this paper) and then nesting this inside an equilibrium model of location and format choice (subsequent research). We begin here by proposing and estimating a store-choice model of the entire US retail grocery industry, in which heterogeneous consumers choose where to allocate their shopping budget based on individual preference. To keep things simple, consumers are initially treated as differing along only two dimensions: income and location. We show how to exploit variation in the joint distribution of location and incomes across the US to recover consumer preferences for grocery bundles from alternative firms. Our goal is to then 1) document whether there are significant differences in consumer tastes across the income distribution, variation that was

assumed away in the classical endogenous cost framework and 2) to explore how those differences influence chains decisions to make more differentiated choices in positioning their stores (in geographic, size and product space).

Our estimates yield several interesting insights. First, we find that consumers indeed find travel costly, and this burden increases significantly with income, consistent with the increased opportunity cost of time. Second, we find that consumers have a strong preference for larger stores (as assumed in the earlier literature) and that this taste for size also increases with income. However, the size premium is not large enough to outweigh travel costs for most consumers. That is, a supercenter format such as Wal-Mart, which specializes in large stores, will find it easier to attract consumers from far away if those consumers are low income. We also find that shopping budgets increase with income—we find that, relative to the median income consumer, consumers at the 75th percentile of income spend roughly 300 additional dollars per person per year on groceries. Finally, for large chains, we use chain-level fixed effects for the and interactions to infer how consumers in different income groups view the quality of different chains, controlling for store size and distance. This quality measure reflects consumers opinion of chain level policies such as pricing, product mix, and the level of service. We find both significant differences in quality across chains and significant differences in quality perceptions of chains across income levels. Most starkly, Whole Foods is the highest quality chain for consumers in the 90th percentile of income, while it is the lowest quality chain for consumers at the 10th percentile of income.

To explore how this variation in tastes impacts optimal store positioning, we present several comparative static exercises. First, we compute a cross-store substitution matrix to illustrate how a uniform increase in the quality of a given chain (or a uniform decrease in its price) affects its own revenue and that of its rivals. These patterns are driven by both geography and the degree to which the stores compete for the same customers in demographic space. We find patterns consistent with segmentation by income that accord with stylized facts from the industry. Second, we calculate distance and income elasticities for firm revenue. With respect to distance, we find evidence of significant heterogeneity across firms. Firms at the high end of quality are most damaged by greater travel distances, while big box stores are hurt least. Again this accords with conventional wisdom and lends face validity to the framework’s ability to reflect the substitution patterns contained in the data. Finally, with respect to income, we again find substantial heterogeneity. We are able to decompose the income elasticity into a wealth effect—the increase in grocery budgets due to higher income—and a substitution affect—the change in relative preference for stores due to a change in income. We find that most of the variation is due to a substitution effect, namely the degree to which greater income drives consumers to make different store choices, as opposed to a wealth effect, which would drive them to spend more on groceries.

	mean	sd	min	max	p5	p95	Observations
Size	30.01	19.06	1	200	6	67	33823
Revenue	307.27	291.07	39	2525	60	950	33823

Table 1: Store Level Information: Size and Revenue

Overall, our results quantify some of the primary drivers of consumer choice in determining where to shop. Our future goal is to use this structure to disentangle the demand side drivers of differentiation (e.g. changing consumer demographics) from the cost side components (e.g. new store formats and distribution technology).

The following section introduces the data we use for this study, and gives some preliminary information about grocery chains. Section 3 introduces the store choice model, and discusses identification and estimation. Our empirical results are presented in Section 4. To conclude, Section 5 considers the implications of our results and presents directions for future research.

2 Data

The data on grocery revenues, locations, store and chain characteristics are drawn from the Trade Dimensions TDLinX dataset for calendar year 2006. Trade Dimensions collects information on every supermarket and grocery store operating in the United States, classifying as a supermarket any “full-line food store” whose annual sales exceed two million dollars (this cutoff is the government and industry standard). Data on store level sales volume is imputed using a proprietary scheme that incorporates store level transaction data for a subset of the full universe of stores (we will account for this source of measurement error in our empirical framework). We also observe the full ownership structure of each firm, allowing us to tie individual stores to either a high level holding company or a smaller collection of co-branded stores that operate under a single banner.

Table 1 provides summary statistics for the full set of 33,823 stores. The average store sells just under \$16 million of groceries per year, with largest stores topping out at over \$100 million (note that the minimum sales level of \$2 million reflects the definitional cutoff noted above). In terms of selling area, the average store is just over 30 thousand square feet, while the largest super centers can be well over 100 thousand square feet. Both size and sales include a sizable amount of variation about their respective means, reflecting differences in both the age of stores and regional variation in zoning, land availability and consumer preferences.

There is also a wide variation in firm size, as shown in Table 1. While the average firm (chain) includes about five stores, the distribution is highly skewed, with a few very large chains and a large number of sole

	mean	sd	min	max	p5	p95	Observations
Stores per chain	4.55	42.35	1	2117	1	8	7424

Table 2: Chain Size

proprietorships. While 25% of the stores belong to single-store firms, there are over 200 firms that operate at least 10 stores (the industry definition of a chain), 116 that operate at least 20, and 39 that operate more than 100. The top 4 chains each operate over 1000 stores.

To account for the role of brand, we will make use of chain fixed effects. However, the large number of multi-store firms makes this impractical at the moment. Therefore, we group firms by chain type and include a full set of fixed effects only for those firms that are in the largest class. In particular, we impose the following firm classification:

- small chains and independents - firms that own less than 10 stores in total.
- medium chains - firms with more than 9 but less than 160 stores in total.
- large chains - firms with 160 or more stores in total.

Table 3 provides summary statistics for the full set of large chains, ordered by number of stores. Wal-Mart is by far the largest chain, both in terms of the number of stores and average sales. It also owns the largest stores, due to both its large scale super center format and the relatively young vintage of its store profile. Note that the size of Wal-Mart stores reflects the size of the grocery sales floor only, not including the mass-merchandise portion of the supercenter. The largest chains also include the low end limited assortment chain Save A Lot, and mid level chains Food Lion and Kroger, as well as more upscale variants like Publix and Safeway. Notably, there is a large amount of variation both across and within firms in the distribution of store sizes and store level revenues. Some firms, such as Food Lion, include a fairly standardized store profile, while others, such as HEB, offer a far more heterogenous set of outlets.

Data on consumer demographics are drawn from the U.S. Census. The 2010 Census Gazetteer files include characteristics of census tracts such as geographic location, population and average income.¹ Table 4 contains population and income statistics for 72,539 census tracts included in the contiguous United States. While census tracts are designed to be somewhat homogeneous in terms of population, there is still a great deal of variation. In addition, there is substantial heterogeneity in terms of average income within a census tract. This heterogeneity is key to our identification strategy exploits variation in store revenues when stores are located near high income tracts relative to low income tracts.

¹Census tracts are defined by the decennial census, we opt to use the 2010 tract-level data as it is the closest census to the 2006 store-level data. While this introduces some measurement error, we believe that population dynamics are small enough that this error is small.

Chain	Number of stores	Mean Revenue	Sd. Revenue	Mean Size	Sd. Size
Wal Mart Supercenter	2,117	1015.14	299.36	63.37	10.44
Save A Lot	1,176	103.32	43.29	12.74	5.97
Food Lion Store	1,174	174.01	54.15	27.75	3.81
Kroger Store	1,067	478.73	201.18	41.82	10.46
Albertsons Store	962	361.67	126.56	50.04	11.35
Safeway Store	921	414.10	167.54	35.20	9.73
Publix Super Market	882	418.20	136.44	38.72	7.62
Aldi Food Store	803	75.66	15.96	12.80	2.53
Winn Dixie	516	252.13	69.54	46.17	8.03
Super Stop & Shop	324	685.88	213.73	47.46	9.46
Piggly Wiggly	295	98.06	47.09	17.13	6.48
Giant Food Store	287	568.29	201.92	37.65	10.43
H E B	261	768.39	351.22	47.92	19.25
Trader Joes Market	255	304.02	70.68	9.44	1.87
BI-LO	235	207.98	67.63	36.10	7.23
Ralphs Grocery	226	338.72	136.59	33.08	6.62
Giant Eagle Market	216	520.72	248.97	42.63	15.84
Hy Vee Food Store	198	419.79	182.93	38.35	14.74
Ingles Supermarket	192	211.56	94.04	40.88	10.96
Jewel Osco	192	660.29	216.28	58.48	10.84
Smart & Final	183	150.82	36.05	15.56	3.99
Meijer Supermarket	176	818.32	220.89	59.07	7.33
Vons Market	176	367.47	146.15	29.65	10.15
Whole Foods Market	171	520.47	176.44	26.96	8.83
Food City	162	265.34	97.21	29.91	11.09
Stater Bros Mkt	162	388.27	54.26	24.22	3.77
SuperTarget Center	162	525.62	120.40	60.60	7.40

Table 3: Size and sales volume of large chains

3 Model

To explore the importance of consumer and firm heterogeneity in driving market structure, we develop a model of individual store choice. Our framework extends the approach proposed by Holmes (2011) to include competing stores of rival firms. To keep things simple, we allow consumers to differ only by location and income, although additional dimensions of heterogeneity could easily be incorporated into the model. Since location is the primary driver of store choice, a consumer’s utility for shopping at a given store will vary with distance to the store. In addition, we assume that consumers only consider shopping at stores within a 15 mile radius of their home, which determines their choice set. Income affects consumer spending in two

	mean	sd	min	max	p5	p95	N
Population	4227	1979	0	37452	1584	7535	72539
Average Income	27183	14061	150	219661	11725	52340	71487

Table 4: Census tract population and income

ways, first through the overall budget that they allocate to groceries and second through the stores in which they choose to shop. This will allow us to capture the tendency for a consumer’s share of budget devoted to food at home to fall with income and the mix of stores they visit to change as well (i.e. away from Kroger and toward Whole Foods). To estimate the model, we use information from the 2000 US Census data on tract-level income and population. Within each census tract, we model a representative consumer dividing her grocery budget across stores depending on her tastes for the characteristics of that store (including the overall brand). This generates store-tract level revenue predictions. By aggregating these revenue predictions across tracts for each store, we can match the model’s predicted revenue for a store to the revenue observed in the data. We use this relationship to estimate the model parameters via nonlinear least squares.

3.1 Consumers’ Store Choice

We assume the existence of a representative consumer at every census tract, indexing consumers according to the tract t in which they reside.² Every grocery consumer in tract t is therefore endowed with the same income inc_t . We assume that the consumer’s weekly grocery budget is a function of their income,

$$b_t(\lambda) = \lambda_0 + \lambda_1 \log(inc_t) \tag{1}$$

This relationship is a simple reduced form of the consumers’ household budgeting problem capturing the fact that grocery expenditures are increasing and concave in income³. Consumers allocate their budget according to a discrete-choice random utility model. We assume the consumers choice set C_t consists of all stores located within 15 miles of the consumer, $C_t = \{s : d_{st} \leq 15\}$, where d_{st} is the distance from the consumer’s location to the store location.

For each unit of expenditure i , consumers choose the store that provides the highest utility, where utility of an expenditure at store s for the representative consumer from tract t is a function of consumer characteristics, store characteristics, and an idiosyncratic preference shock for that expenditure,

$$u_{sti} = \tau_0 d_{st} + \tau_1 d_{st} \log(inc_t) + \gamma_0 size_s + \gamma_1 size_s \log(inc_t) + \beta_0^{f_s} + \beta_1^{f_s} \log(inc_t) + \epsilon_{sti}. \tag{2}$$

The first two terms in the firms utility function, $\tau_0 d_{st} + \tau_1 d_{st} \log(inc_t)$, capture the consumer’s disutility of traveling over distance d_{st} . We allow this disutility to be a function of income, reflecting the fact that wealthier consumers may have higher travel costs. One reason for this relationship is the opportunity costs

²It would be conceptually straightforward to allow for unobserved heterogeneity at the census tract level, resulting in a random coefficients model in the spirit of Berry et al. (1995).

³For computational reasons we are using deviation of log of income from it’s mean instead of log of income in the estimation routine. In order to make intuition more transparent, we are stick with simpler set up in the model section.

of high-income consumer’s time. However, a second possibility is that wealthier consumers live in higher density locations in which it is more costly to travel. In principle, additional controls could separate these two effects (e.g., by controlling for population density as in Holmes (2011), or measuring distance by “travel minutes” rather than miles).

The next two terms, $\gamma_0 size_s + \gamma_1 size_s \log(inc_t)$, account for differences in utility due to the size of the selling area store s (measured in thousands of square feet). Both the supermarket and super center formats were introduced to accommodate the explosive growth in the number of products carried per store, which coincided with the rise of national brands, the increasing importance of mass media advertising, and the adoption of advanced distribution technology. Since larger stores generally provide a wider array of products, we expect consumers to prefer larger stores, all else equal. This is consistent with the vertical model of competition proposed by Ellickson (2007), in which firms invest in distribution technology to efficiently stock a wider array of products. Moreover, larger stores are less like to stock out of any given product (Matsa, 2011), providing another reason for consumers to prefer larger formats. We allow the taste for size to vary by income, as one might expect higher income customers, who have more money to spend at a store, to have a higher preference for stores that offer greater variety. Of course, larger stores are more expensive to build and operate and tend to be more remotely located. The natural tension between travel time and access to variety is accounted for in the model.

Next, we consider chain-level drivers of utility, $\beta_0^{f_s} + \beta_1^{f_s} \log(inc_t)$. Here f_s is the identity of the firm (i.e., chain) who operates store s . Since we do not observe prices or the particular set of products on offer, the chain effect is aimed at capturing the broad pricing and quality strategy of the firm, which we assume is set at the chain level. That is, some firms may choose to market themselves as limited-assortment, low-price grocery stores, while others may choose to be “boutique” stores that offer expensive, high-end products. Other, more conventional chains, may choose to serve a broad segment of the market. These different strategies will appeal to different consumers in heterogeneous ways. For example, Ellickson and Misra (2008) find evidence that supermarket firms use distinct price and positioning strategies to target different consumer segments depending on purchase size. While our approach does not control for pricing and quality decisions that are specific to the individual store, we believe these are second order to the average policy set by the chain.⁴ The chain effect accounts for the fact that the basket of goods received when you purchase a dollar’s worth of goods at Whole Foods is different from the basket of goods you receive from a dollar’s expenditure at Aldi, and that the utility tradeoff between these two expenditures varies by income. In practice, we consider

⁴For example, one could imagine adding a “store level quality shock” to the model which firms could observe and use in endogenously adjusting the stores quality-price offerings. While there is some evidence that firms do so, there are also strong branding and efficiency reasons for them to maintain uniformity. While exploring such a model is an interesting avenue for future work, given that we do not have explicit data on individual store prices or product offerings, and given that marketing campaigns enforce a large degree of heterogeneity in chain policies, we do not consider this approach in this paper.

chain effects for the largest 27 chains (based on store counts) listed in Table 3. As noted above, we aggregate the remaining chains with 10-160 stores into a single “medium chain” category and small firms with fewer than 10 stores into a small chain category, which includes independents.⁵

Finally, ϵ_{sti} is an idiosyncratic preference shock for each expenditure, assumed to be distributed type-I extreme value, and independent across expenditures and stores. This shock accounts for both heterogeneity across consumers within the tract, and heterogeneity across expenditures for the same consumer. As a result, the representative consumer will make some purchases at all stores in the choice set, and we can derive the share of grocery expenditures that consumers from tract t spend at store s by integrating over ϵ_{sti} ,

$$p_{st}(\tau, \gamma, \beta) = \frac{e^{u_{st}}}{\sum_{q \in C_t} e^{u_{qt}}}, \quad (3)$$

where $u_{st} = u_{sti} - \epsilon_{sit}$. In principle, we could allow for unobserved heterogeneity to depend on store or tract characteristics, which would be equivalent to allowing for random coefficients in our utility framework—as in Berry et al (1995). However, using the joint distribution of income and location already accommodates a substantial amount of heterogeneity and, as we will see, allows for the calculation of rich substitution patterns across chains.⁶ Notice that the store choice model does not include an outside good. The possibility of not purchasing groceries is already accounted for in the size of the grocery budget (1), while the store choice model above determines the share of grocery expenditure that goes to store s relative to total grocery expenditure for consumers in tract t .

3.2 Estimation with Store Revenue Data

We estimate the model using our census of store-level revenues for all firms operating in the US. Stores collect revenue from consumers’ expenditures on groceries. Given the model parameters $\theta = (\lambda, \tau, \gamma, \beta)$, it is easy to calculate the predicted revenue that store s receives from tract t ,

$$\hat{R}_{st}(\theta) = p_{st}(\tau, \gamma, \beta) b_t(\lambda) n_t.$$

where n_t is the total population living in tract t . Store s collects revenue from all tracts where it is included in the choice set (i.e. all tracts within 15 miles of its location). Therefore, predicted total revenue for store

⁵Recall that within the small chain category, the majority of firms are single-store chain, while the second largest category operates only 2 chains.

⁶Moreover, since we do not directly observe tract-level revenue shares, identification of unobserved heterogeneity would rely on aggregation of these preferences over tracts. While this heterogeneity may be identified in principle, the flexibility of the current model suggests there is little to gain from this approach, and has the benefit of being directly tied to observable demographic variables.

s is,

$$\hat{R}_s(\theta) = \sum_{t \in L_s} R_{st}(\theta) \quad (4)$$

where $L_s = \{t : s \in C_t\} = \{t : d_{st} \leq 15\}$ is the set of tracts in the catchment area of store s . We will compare the model-generated revenue predictions to the revenues recorded in the data. However, because they are imputed from a random sample of store level scanner data, the Trade Dimensions data are measured with error. To account for this error (and the resulting mismatch between the revenue predicted by the model and that which is observed in the data), we assume that the observed revenues for each store are perturbed by a multiplicative shock,⁷

$$R_s = e^{\eta_s} \hat{R}_s(\theta).$$

We assume that η_s is mean zero and independent of the exogenous variables. We then estimate θ using nonlinear least squares,

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \sum_s \left(\log(\hat{R}_s(\theta)) - \log(R_s) \right)^2 \quad (5)$$

It is straightforward to show that this estimator is consistent and asymptotically normal, with the standard variance-covariance matrix implied by the nonlinear least squares objective function.

3.3 Identification

Before turning to the results, it is useful to briefly discuss how the model parameters are identified from the data. First, as in all discrete choice models, utility is only identified up to a location normalization. Adding a unit to each element of u_{st} produces identical expenditure shares as the original model. The standard normalization is to normalize one element of the choice set to zero, usually the outside good. Since there is no outside good in our model, and since all stores have nonzero distance and size characteristics for most tracts, we instead normalize the fixed effect of small stores, β_0^{small} to zero. Therefore, all elements of $\beta_0^{f_s}$ are relative to small stores. Moreover, because income does not vary within tracts, we must also normalize β_1^{small} to zero by the same logic.⁸ It is of course still possible that the utility consumers derive from small stores, holding characteristics fixed, varies with income. That said, our estimates of $\beta_1^{f_s}$ should be understood as the difference in the slope of the utility-income relationship for store f_s relative to small stores. While this normalization does have implications for the calculation of consumer surplus (it is difficult to compare surplus across income levels without stronger assumptions), it is without loss of generality when considering substitution patterns, income elasticities, and other outcomes relevant to store profits.

⁷We have also estimated the model where measurement error enters via an additive shock, the qualitative results are similar.

⁸A model that allowed for within-tract variation in income would not need to make this normalization, however data on tract-level income distributions is not available to our knowledge.

We further assume that $\epsilon_{it,s}$ and η_s are independent of stores' locations and size as well as consumers' locations and income. Critically, we assume that consumers take store locations as given, and that consumers' perceptions of stores pricing, quality and product policies are set at the chain level, rather than the store level. This allows us to control for these policies through chain fixed effects. Of course, it is possible that chains adjust their pricing policies store by store, based on local demographics. While there is some evidence that they do so, we view this a second order for two reasons. First, grocery pricing policies are a combination of thousands of prices, and it is not realistic to believe that consumers calculate store level price indices for each store. It is much more likely they have a rough perception of the price differences across chains. Second, grocery stores usually do not set prices at the store level, but instead set the same price across broad "pricing zones." The impetus for these zones is so that stores can jointly market their products (for example, through newspaper circulars) to an area that is wider than a given store's catchment area. This suggests that it is not efficient for chains to set policies at the level of the individual stores. While "pricing-zones" are not nation-wide, it is reasonable to believe that within-brand variation in pricing and product offerings across a pricing zone is less important than across-chain variation in these policies within the same zone.

Having discussed the assumptions which allow us to assume store and tract characteristics are exogenous, identification exploits the richness of the geographic distribution of stores and consumers. For example, consider the variation in distance. Varying the distance between a tract and a store alters the share of expenditure at that store relative to others in the tract's choice set, $p_{st}(\tau, \gamma, \beta)$. This will be reflected in the store's revenue relative to others in the same choice set, all of which are observed. A similar logic can be used to identify the parameters relating to store size and the firm fixed effects. Considering the effect of income is only slightly more involved. Income alters the substitution pattern between stores, but also impacts total grocery expenditure from the tract. This second effect is captured through the tract grocery budget, $b_t(\lambda)$. As income increases, the total revenue at all stores within the tract's choice set will increase. The income increase will also adjust the substitution pattern across stores, as discussed earlier, but these will net to zero across all stores in the choice set. Thus, the relationship between income and expenditure is identified through the total grocery expenditure across the tract's choice set.

4 Empirical Results

The store choice/revenue model allows us to look at three different dimensions of grocery competition: consumer's preferences over stores, chain heterogeneity in product characteristics, and chain-level competition within markets. We consider each in turn.

Budget, λ_0	32.281*** (0.239)
Budget x $\log(inc_t)$, λ_1	10.643*** (0.679)
Distance, τ_0	-0.137*** (0.007)
Distance x $\log(inc_t)$, τ_1	-0.169*** (0.014)
Size, γ_0	0.027*** (0.0004)
Size x $\log(inc_t)$, γ_1	0.006*** (0.0013)

Standard errors in parentheses.

Table 5: Consumer parameters

4.1 Model Parameters

The parameters of the consumer’s budget and utility functions are presented in Table 5. All parameters have the expected signs and are highly significant. In the estimation, we have demeaned $\log(inc_t)$, so that the constants represent the coefficients at the average logarithm of income across tracts. In practice, this is approximately the median income in the population. The first two terms (λ ’s) capture the consumer’s allocation of money to groceries. According to the model, the median income person spends roughly 32 dollars on groceries weekly (note that this number is per person, not per household). Not surprisingly, this budget is increasing in log income. Note that the relationship is concave due to the using of the log function. The budget parameters allow us to present grocery budget distribution over entire population of US which is given in Figure 1. The distribution implies that upper quartile spends roughly 300 dollars more per person on groceries per year than the median person. The second two terms (τ ’s) reflect the impact of distance, measured in miles. Consumer’s dislike travel, and so are more likely to spend money at stores closer to their homes. This distaste increases significantly with income, consistent with higher income consumers facing higher opportunity costs of travel. Finally, consumer’s prefer larger stores (the γ terms) and this preference also increases with income. Size is measured in units of thousands of square feet of sales floor. Interestingly, consumers appear to dislike travel much more than they like scale. At the median income, consumers are indifferent between making a store 1000 square feet larger or 0.15 miles closer to their homes. This relationship grows even steeper as income increases.

The remaining parameters of the model are the revealed preference estimates of chain quality for the 27 large chains and the medium chain category. These are presented in Table 6. Recall that we are normalizing

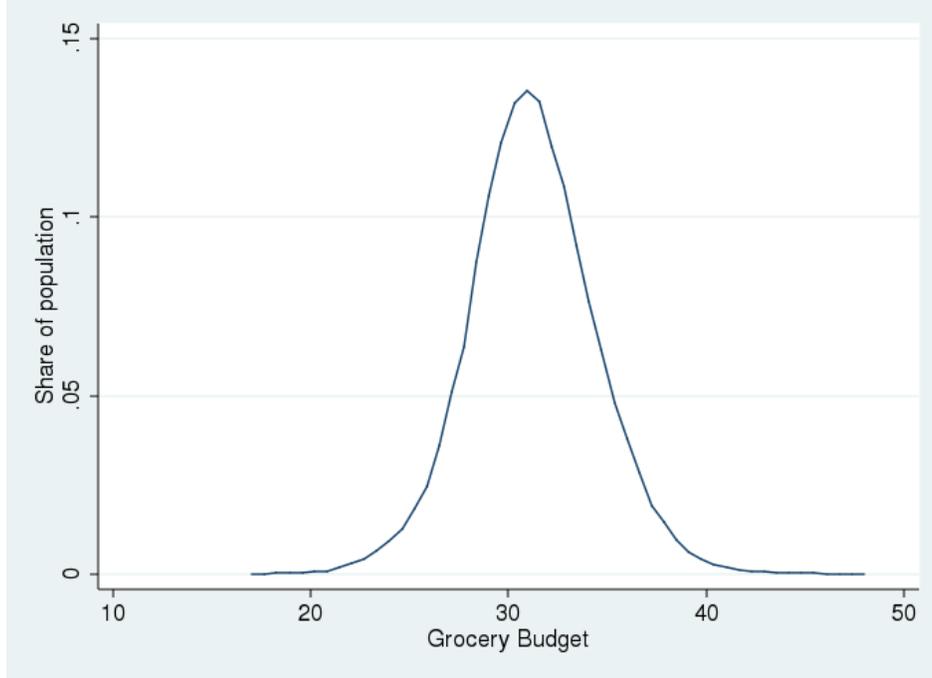


Figure 1: Distribution of Grocery Budgets

relative to small stores, so all coefficients should be interpreted relative to this baseline. This is true of both the constants and the $\log(\text{inc}_t)$ interactions. That is, the slope of a chain's quality with income is relative to the change in quality of a small store. This normalization is without loss of generality when considering functions of store revenue (such as the semi-elasticities and elasticities presented below), but it does have implications when calculating consumer surplus, since the utility of a small store (controlling for size and distance) is normalized to zero for all income levels. Relatedly, the fact that chain preferences are allowed to vary with income complicates the interpretation of the intercept. Because we have demeaned $\log(\text{inc}_t)$, the intercept reflects the baseline tastes of the (roughly) median income consumer, relative to the small stores. The main takeaway from this first column is simply the large degree of heterogeneity: chains vary dramatically in their perceived baseline level of quality, even after controlling for variation in the size and location of stores. Because the revealed preference quality measure reflects all chain-level differences in store policies (i.e., pricing, marketing, staffing, layout, etc), it is difficult to draw a connection between these quality levels and more concrete quality measures (such as stock-out rates, or average margins). The interactions with income are more straightforward, creating a natural separation of firms that clearly target the wealthy (e.g. Whole Foods, Giant and Publix) and those that instead target the less affluent (Walmart, Stop and Shop, Meijer). While many of the slopes are not statistically significantly different from the small store category, the joint hypothesis that all of the slopes are zero is strongly rejected with a p-value of less than .0001.

	Constant	Standard errors	Log Income Interaction	Standard errors
Medium Chain	0.335***	0.018	0.146*	0.064
Albertsons Store	0.196***	0.047	-0.252	0.218
Aldi Food Store	-0.209*	0.071	-0.045	0.462
BI-LO	0.222	0.146	-0.381	0.551
Food City	0.449*	0.165	-0.662	0.509
Food Lion Store	0.191***	0.057	-0.261	0.265
Giant Eagle Market	0.758***	0.102	-0.219	0.789
Giant Food Store	0.893***	0.159	0.268	0.453
H E B	0.943***	0.106	-0.159	0.295
Hy Vee Food Store	0.709***	0.121	-0.281	0.839
Ingles Supermarket	0.089	0.155	-0.286	0.715
Jewel Osco	0.668***	0.092	-0.334	0.517
Kroger Store	0.759***	0.047	-0.245	0.233
Meijer Supermarket	0.860***	0.104	-0.892	0.675
Piggly Wiggly	0.037	0.144	0.224	0.457
Publix Super Market	0.711***	0.065	0.073	0.324
Ralphs Grocery	0.475***	0.089	-0.155	0.434
Safeway Store	0.695***	0.047	0.037	0.163
Save A Lot	0.108*	0.054	-0.072	0.232
Smart & Final	0.075	0.123	-0.274	0.478
Stater Bros Mkt	1.064***	0.162	-0.424	0.742
Super Stop & Shop	2.611***	0.106	-9.718***	0.433
SuperTarget Center	0.308	0.172	0.105	0.803
Trader Joes Market	1.062***	0.151	-0.066	0.505
Vons Market	0.742***	0.090	-0.574	0.438
Wal Mart Supercenter	0.805***	0.037	-0.295*	0.141
Whole Foods Market	-0.178	0.261	2.515***	0.328
Winn Dixie	0.043	0.080	-0.467	0.386

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: Chains parameters

As is common in discrete choice models, the utility parameters are difficult to interpret directly. In the remainder of this section, we use the model to document the impact of heterogeneity through a variety of well-defined elasticities and marginal effects.

4.2 Chain Competitive Effects

To get a clearer sense of the substitution patterns across firms, we compute own and cross elasticities for the large chains based on increasing the level of firm quality implied by the firm fixed effects. In particular, consider the revenue of a chain f :

$$R_f = \sum_{s \in F_f} \sum_{t \in L_s} (\lambda_0 + \lambda_1 \log(\text{inc}_t)) p_{st} n_t \quad (6)$$

Where F_f is the set of stores owned by firm f and, as defined earlier L_s is the set of tracts with store s inside their choice set. Then, the derivative of chain f 's revenue with respect to a chain g 's fixed effect (not interacted with income) is the following:

$$\begin{aligned} \frac{\partial R_f}{\partial \beta_0^g} &= \sum_{s \in F_f} \sum_{t \in L_s} (\lambda_0 + \lambda_1 \log(\text{inc}_t)) \frac{\partial p_{st}}{\partial \beta_0^g} n_t = \sum_{s \in F_f} \sum_{t \in L_s} (\lambda_0 + \lambda_1 \log(\text{inc}_t)) n_t \sum_{q \in C_t} \frac{\partial p_{st}}{\partial u_{qt}} \frac{\partial u_{qt}}{\partial \beta_0^g} = \\ &= \sum_{s \in F_f} \sum_{t \in L_s} (\lambda_0 + \lambda_1 \log(\text{inc}_t)) n_t \left(\mathbf{1}\{f = g\} p_{st}(1 - p_{st}) - \sum_{q \in C_t; q \neq s} \mathbf{1}\{q \in F_g\} p_{st} p_{qt} \right) \end{aligned}$$

Recall that C_t is the choice set of consumers who live in tract t . The second line makes use of the well known derivative of the logit share formula. It can then be shown that:

$$\frac{\partial R_f}{\partial \beta_0^g} = \frac{\partial R_g}{\partial \beta_0^f}, \quad \forall f, t$$

This symmetry is driven by the lack of an outside good. In short, grocery budgets are fully determined by income, and a change in store quality impacts how much share of grocery budgets is attracted to that store relative to others. Thus, changes in store quality generate flows of revenue across chains which must net to zero.

In this manner, we can treat chain fixed effects as proxies for the prices or quality of services offered by a given firm. As a result, the derivative of the chain revenue with respect to another chain's fixed effect reflects the level of competition between chains. This means that any gain in revenue will come at the expense of other chains and we can say which chain will be hurt the most. Note that the main competitor status (i.e. who is hurt the most) will be mostly driven by geography, since the closer two firms' stores are to each other the more they can steal revenue from one another. That said, it will also be affected by the characteristics of chains - similar chains will compete more closely with one another.

Because these derivatives are denominated in dollars, and of course in part driven by the size of the chains involved, it is useful to instead consider the change in terms of a percentage of a firm's revenue. Table 7 contains semielasticities of chains with respect to fixed effects, i.e., $\frac{\partial \log(R_f)}{\partial \beta_t}$, which is simply the above formula divided by R_f , for each firm with respect to itself and its two strongest competitors. Intuitively, these are the percentage change in revenue of firm f for a differential increase in the quality of firm g . For example, a Δ increase in Albertson's chain fixed effect will increase its revenue by $84.7\Delta\%$. On the other hand, Albertson's is most affected by an adjustment in Wal-Mart's quality. The same Δ increase in Wal-Mart Supercenter's fixed effect will decrease Albertson's revenue by $13.2\Delta\%$.

Several interesting patterns emerge from Table 7. With respect to own elasticity, the largest values

correspond to the firms that are commonly thought to offer highest levels of quality (e.g. Whole Foods, Trader Joe’s), while the lowest values correspond to firms like Wal-Mart and HEB that operate super centers or firms like Food Lion and Bi-Lo that target the poor. A supply-side interpretation of this result might be that the shadow cost of quality is higher for “high-end” firms like Whole Foods who already put a lot of resources towards providing quality. Turning to the cross elasticities, it is striking just how large a shadow Wal-Mart casts. It is the largest competitor of 15 of 27 chains, and the second largest competitor of an additional 4. Part of this is clearly driven by its enormous scale and national presence, but it also reflects its close proximity in product space to many of these conventional chains. Recall that Wal-Mart entered a relatively undifferentiated industry. In particular, its impact is strongest on either mid-tier southern chains (Bi-Lo, Food Lion, Ingles and Winn-Dixie) or firms that also operate super centers (HEB, Meijer, Hy Vee). In contrast, Wal-Mart is relatively insulated from competition from others (rivaled only by Stop and Shop, which has a near monopoly presence in many New England markets). This is consistent with Wal-Mart’s enormous cost advantage. More broadly, the chains that are hurt the least by their largest rivals are the firms that are generally thought to have significant market power, either due to regional monopoly (Giant Eagle, Jewel Osco, Super Stop and Shop) or their isolated position in product space (Whole Foods, Trader Joe’s). Overall, it appears that the choice model is able to replicate the rich substitution patterns that we believe exist in the market and which are a necessary input to a store location model.

Figure 2 provides an alternative illustration of the patterns present in Table 7. Wal-Mart is the main competitor of 15 out of 27 large chains. The second most influential chain is Albertsons, which is the main competitor for 5 large chains. It is worth noting that these chains are much smaller in terms of the total revenue.

4.3 Distance and Income elasticities

We now turn from competition between firms to an analysis of the drivers of how firms locate relative to consumers. The two key elements of consumer heterogeneity in our location and income. We can further unpack the impact of distance and income on firms by computing elasticities for each. We start with distance and then consider income. Recall the utility function from earlier,

$$u_{sti} = \tau_0 d_{st} + \tau_1 d_{st} \log(\text{inc}_t) + \gamma_0 \text{size}_s + \gamma_1 \text{size}_s \log(\text{inc}_t) + \beta_0^{f_s} + \beta_1^{f_s} \log(\text{inc}_t) + \epsilon_{sti}. \quad (7)$$

where $s \in C_t$. We want to construct a measure of how chain-level revenue responds to increasing the distance to consumers, say by building more remote stores in suburban locations (we’ll turn to income next). The distance elasticity can be constructed in the following steps:

Chain	Semielast.	Main Competitor	Semielast.	Second Competitor	Semielast.
Albertsons Store	0.847	Wal Mart Sup.	-0.132	Safeway Store	-0.084
Aldi Food Store	0.993	Wal Mart Sup.	-0.168	Jewel Osco	-0.067
BI-LO	0.787	Wal Mart Sup.	-0.271	Food Lion Store	-0.126
Food City	0.664	Wal Mart Sup.	-0.211	Safeway Store	-0.069
Food Lion Store	0.704	Wal Mart Sup.	-0.225	Kroger Store	-0.056
Giant Eagle Market	0.572	Wal Mart Sup.	-0.111	Save A Lot	-0.031
Giant Food Store	0.667	Safeway Store	-0.137	Wal Mart Sup.	-0.050
H E B	0.404	Wal Mart Sup.	-0.211	Kroger Store	-0.034
Hy Vee Food Store	0.545	Wal Mart Sup.	-0.209	SuperTarget Center	-0.027
Ingles Supermarket	0.708	Wal Mart Sup.	-0.221	BI-LO	-0.085
Jewel Osco	0.653	Aldi Food Store	-0.032	Meijer Supermarket	-0.031
Kroger Store	0.610	Wal Mart Sup.	-0.189	Meijer Supermarket	-0.051
Meijer Supermarket	0.716	Kroger Store	-0.181	Wal Mart Sup.	-0.107
Piggly Wiggly	0.816	Wal Mart Sup.	-0.258	Winn Dixie	-0.055
Publix Super Market	0.532	Wal Mart Sup.	-0.163	Winn Dixie	-0.093
Ralphs Grocery	1.002	Albertsons Store	-0.166	Vons Market	-0.119
Safeway Store	0.692	Albertsons Store	-0.079	Wal Mart Sup.	-0.067
Save A Lot	0.926	Wal Mart Sup.	-0.198	Kroger Store	-0.073
Smart & Final	1.077	Albertsons Store	-0.141	Ralphs Grocery	-0.107
Stater Bros Mkt	0.929	Albertsons Store	-0.199	Ralphs Grocery	-0.128
Super Stop & Shop	0.424	Wal Mart Sup.	-0.011	Trader Joes Market	-0.006
SuperTarget Center	0.902	Wal Mart Sup.	-0.178	Publix Super Market	-0.069
Trader Joes Market	1.073	Safeway Store	-0.120	Albertsons Store	-0.098
Vons Market	1.051	Albertsons Store	-0.199	Ralphs Grocery	-0.141
Wal Mart Sup.	0.455	Kroger Store	-0.045	Publix Super Market	-0.028
Whole Foods Market	1.123	Safeway Store	-0.117	Giant Food Store	-0.051
Winn Dixie	0.779	Publix Super Market	-0.265	Wal Mart Sup.	-0.239

Table 7: The impact of quality on own and cross substitution

1. Store-tract level elasticity: $\eta_{st} = \frac{\partial R_{st}}{\partial dist_{st}} \frac{dist_{st}}{R_{st}} = (\lambda + \lambda_i \log(inc_t)) n_t (\tau_0 + \tau_1 \log(inc_t)) \frac{dist_{st}}{R_{st}}$;
2. Store level elasticity: $\eta_s = \sum_{t \in L_s} \eta_{st} \frac{R_{st}}{R_s}$;
3. Chain level elasticity: $\eta_f = \sum_{s \in F_f} \eta_s \frac{R_s}{R_f}$.

So the final chain level elasticity represents the revenue weighted impact of distance to consumers for chain f . The estimated elasticities are presented in the first column of Table 8. The distance elasticities are presented first in column 2. Note that the firm that faces the highest distance elasticity is Whole Foods, reflecting their upscale focus on ultra high-income consumers who we have found have a high disutility of distance. More broadly, the firms with elasticities below -.9 have a clear urban focus (Target, Trader Joes, Ralphs) or target high income consumers (Whole Foods, Publix, Giant Foods). In contrast, firms with elastics above -.7 are almost exclusively super centers (Wal-Mart, HEB, Hy Vee, Super Stop and Shop) who tend to exploit an explicit large-scale, large catchment area strategy. Thus, there is strong evidence of a clear segmentation in strategy between chains.

Turning now to income, note that since our model has a two stage structure in which consumers choose a budget and then allocate that budget across stores, we can decompose the revenue impact from income into two distinct effects: wealth and substitution. In particular, the income elasticity involves similar steps as above but yields two terms:

1. Store-tract level elasticity:

$$\nu_{st} = \frac{\partial R_{st}}{\partial inc_t} \frac{inc_t}{R_{st}} = \frac{\partial R_{st}}{\partial \log(inc_t)} \frac{1}{R_{st}} = \frac{\lambda_1 p_{st} n_t}{R_{st}} + \frac{(\lambda_0 + \lambda_1 \log(inc_t)) n_t}{R_{st}} \frac{\partial p_{st}}{\partial \log(inc_t)}$$

2. Store level elasticity: $\nu_s = \sum_{l \in L_s} \nu_{st} \frac{R_{st}}{R_s}$

3. Chain level elasticity: $\nu_f = \sum_{s \in C_f} \nu_s \frac{R_s}{R_f}$

The first term of ν_{st} is the wealth effect, additional amount of expenditure a consumer in tract t spends on groceries at store s as income rises, holding substitution patterns fixed. The second term represents the substitution effect, how a change in income effects how the consumer divides their grocery budget among stores, holding total expenditure fixed. As with the distance elasticities, we can aggregate these effects to the store and chain levels.

The estimated income elasticities, and the decomposition into wealth and substitution effects, are presented in the final three columns of Table 8. Looking at the overall effect, we again see a large degree of heterogeneity. While most firms benefit from an increase in their consumer's income (perhaps not surprisingly) there are a few at the extreme low end that do not. On the other hand, some firms benefit increases in income a lot more than others (e.g. Whole Foods). In general, the firms that target the wealthy (Publix, Target, Giant) have the highest values. The next two columns decompose this overall effect into wealth and substitution effects. Not surprisingly, the wealth effect is positive for all, though it is smallest for Whole Foods. This makes since given the diminishing returns to greater income built into the budget function, the fact that Whole Foods is located in the wealthiest suburbs, and the tendency for the highest income consumers to substitute away from groceries toward restaurants and prepared foods. In contrast, some of the highest values are obtained for firms that actively target the poor (Save A Lot, Food City, Winn Dixie) these stores's consumers see the largest increase in grocery expenditure as a result of an increase in income. The final column reveals that most of the variation in income elasticity is driven by the substitution effect. Here the results are quite striking and intuitive. Positive values are associated with the firms the most narrowly target the rich (Giant, Publix, Target, and especially Whole Foods), while the most negative values are associated with urban stores aimed at a lower income demographic, primarily in California. Overall, the substitution patterns reinforce the claim that these firms are targeting quite distinct consumer segments,

	Distance Elasticity	Income Elasticity	Wealth Effect	Substitution Effect
Small Chains	-0.804	0.425	0.338	0.087
Medium Chains	-0.961	0.662	0.323	0.339
Albertsons Store	-0.862	0.202	0.332	-0.130
Aldi Food Store	-0.874	0.233	0.339	-0.106
BI-LO	-0.759	0.142	0.351	-0.209
Food City	-0.666	-0.122	0.364	-0.486
Food Lion Store	-0.820	0.149	0.342	-0.193
Giant Eagle Market	-0.897	0.240	0.335	-0.095
Giant Food Store	-1.167	0.420	0.304	0.116
H E B	-0.656	0.338	0.356	-0.018
Hy Vee Food Store	-0.647	0.189	0.335	-0.146
Ingles Supermarket	-0.762	0.295	0.351	-0.056
Jewel Osco	-0.985	0.115	0.329	-0.214
Kroger Store	-0.873	0.246	0.338	-0.092
Meijer Supermarket	-0.797	-0.204	0.345	-0.550
Piggly Wiggly	-0.564	0.543	0.367	0.177
Publix Super Market	-0.927	0.442	0.328	0.114
Ralphs Grocery	-0.912	0.259	0.333	-0.074
Safeway Store	-0.985	0.247	0.315	-0.068
Save A Lot	-0.774	0.316	0.350	-0.033
Smart & Final	-0.811	0.015	0.346	-0.330
Stater Bros Mkt	-0.777	0.069	0.352	-0.284
Super Stop & Shop	-0.647	-2.021	0.366	-2.387
SuperTarget Center	-1.031	0.521	0.319	0.202
Trader Joes Market	-1.011	0.181	0.316	-0.135
Vons Market	-0.787	-0.132	0.345	-0.477
Wal Mart Supercenter	-0.652	0.296	0.346	-0.050
Whole Foods Market	-1.167	2.343	0.259	2.084
Winn Dixie	-0.740	0.034	0.348	-0.315

Note: Income Elasticity is a sum of Wealth and Substitution effects

Table 8: Distance and Income Elasticities

and that both income and location are critical for understanding store choice.

Finally, as the largest player in the market, Wal-Mart is worthy of special attention. Wal-Mart’s substitution effect is very close to zero, suggesting that it is not exclusively targeting low income consumers, but that it instead serves a broad segment of consumers. This is consistent with the Wal-Mart pursuing a strategy of attracting a large number of consumers in order to leverage significant economies of scale (Basker et al., 2012).

5 Conclusion and Directions for Future Research

We believe our results clearly establish that there is substantial heterogeneity in consumer preferences for groceries, and this heterogeneity provides firms an incentive to diversify. Both income and location play an

important role here, and the model could easily be extended to include additional demographic variables. An interesting question is why differentiation has intensified recently. We see three possible explanations: (1) Consumers preferences have become more heterogeneous; (2) While consumers preferences have stayed roughly the same, the distribution of consumers has changed, and the result is it is more profitable for stores to offer differentiated products; (3) The incentives of stores to differentiated have not changed, but new technologies have permitted the entry of new store types, while existing stores have repositioned to account for this change in market structure.

We see our model as an important first step to establishing which of these three possibilities is driving the change in market structure in this industry. We plan to extend our work in two directions. First, by estimating the model for earlier years (prior to Wal-Mart's expansion), we can distinguish between changes in demographics and tastes, and further pin down the degree to which these factors drove the firm response. Second, by characterizing the revenue response of firm location, size and quality, our model provides a key ingredient towards estimating the essential revenue side components driving firms location decisions. Revealed preference could then identify the cost side implications.

References

- Basker, E., S. D. Klimek, and P. H. Van (2012). Supersize it: The growth of retail chains and the rise of the “big box” store. *Journal of Economics & Management Strategy* 21(3), 541–582.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63(4), 841–890.
- Ellickson, P. B. (2007). Does sutton apply to supermarkets? *RAND Journal of Economics* 38(1), 43–59.
- Ellickson, P. B. and S. Misra (2008). Supermarket pricing strategies. *Marketing Science* 27(5), 811–828.
- Holmes, T. J. (2001). Barcodes lead to frequent deliveries and superstores. *RAND Journal of Economics* 32(4), 708–725.
- Holmes, T. J. (2011). The diffusion of wal-mart and the economies of density. *Econometrica* 79(1), 253–302.
- Matsa, D. A. (2011). Competition and product quality in the supermarket industry. *Quarterly Journal of Economics* 126(3), 1539–1591.
- Sutton, J. (1991). *Sunk Costs and Market Structure*. Cambridge, MA: MIT Press.

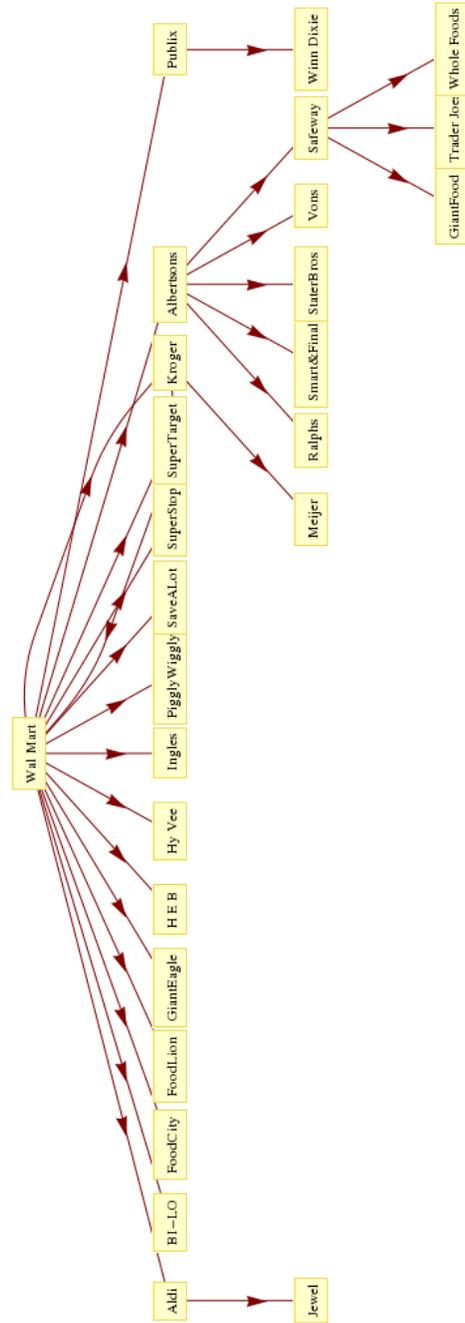


Figure 2: Main competitors