Differences in Differentiation: Rising Variety and Markups in Retail Food Stores

James Brand*

May 5, 2021

Abstract

In this paper I show that consumers in food stores and supermarkets/hypermarkets became significantly less price sensitive between 2006 and 2017. At the median, across thousands of stores and products in nine large categories, estimated own-price elasticities have declined by 25% over this period. I argue that these changes are likely due in part to improved supply chain management, which has led stores to offer a larger variety of goods which better match consumers' individual preferences. I show that newer products are indeed more "niche" in this sense, and that other potential sources of rising differentiation including increases in quality and changes in consumer wealth play a smaller role. Markups implied by a monopolistic pricing rule suggest that the observed rise in differentiation was large enough to generate significant increases in firms' markups absent any changes in pricing behavior or competition.

*Office of the Chief Economist. Microsoft. email: jamesbrand@microsoft.com. I'm thankful to Dan Ackerberg, Jorge Balat, Eugenio Miravete, Yuya Sasaki, and Bob Town for their advice and comments. Thanks also to Pablo Varas for many helpful conversations. Researcher(s) own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

Retail supply chains have become substantially more advanced over recent decades. As large firms have invested in information technology (Shin and Eksioglu, 2014; Soliman et al., 2005) and have developed more agile vertically integrated distribution networks (Kuhn and Sternbeck, 2013), they have been able to more efficiently keep track of and move inventory between wholesalers, distribution centers, and stores. As a result, stores are now able to sell a much wider variety of products than in the past (Fernie, Sparks and McKinnon, 2010; Hortaçsu and Syverson, 2015; Consumer Reports, 2014). These trends have provided retail consumers with far more choices, thereby increasing the likelihood that a given consumer will find a product which closely fits their preferences (Baumol and Ide, 1956). Recent work by Neiman and Vavra (2019) detailing the rising concentration of household purchases over time is quite consistent with this intuition.

In this paper I discuss the ramifications of these facts from an industrial organization perspective, in order to understand the ways in which product selection, prices, and markups in retail food stores may have changed as a result of advances in supply chain management. Traditional discrete choice models of demand highlight the potential benefits of differentiated goods to multi-product firms such as retail food stores. In some models of demand, a large portfolio of goods can crowd the characteristic space, meaning consumers will be increasingly willing to substitute between goods as the size of their choice set grows. Offering highly differentiated goods is one way the firm can limit such substitution and maintain large markups even as they offer a growing portfolio of goods. In particular, as I demonstrate in a simple discrete choice model in Section 2, as the number of products sold in a store increases, firms have an incentive to offer products for which consumers have increasingly heterogeneous preferences, which I call "niche" products. Given a large enough menu of products offered, niche products offer the monopolist the opportunity to better match each consumer's tastes, thereby permitting the firm to better price discriminate and charge higher prices. The rise of agile, flexible, supply chains seems to offer a natural opportunity for firms to have made progress on these dimensions. Supply chains have become more demand focused (e.g. the "quick response" regime), meaning firms receive regular information about the real-time demand for each product. This provides the opportunity to better select the assortment of products offered in order to maximize profits. Accordingly, the hypothesis of this paper is that, as firms have changed the assortment of products stocked, products have become more differentiated in the eyes of consumers, and markups have risen as a result.

With this hypothesis in mind, the focus of the empirical work herein is in estimating consumer demand for many products in 2006 and 2017 across thousands of retail stores in the United States. With these estimates I provide partial answers to three questions: (1) how has differentiation between products within stores changed over time? (2) what evidence is there that firms have increased markups in response to these changes? (3) what role have niche products played in these changes? The first question is interesting in its own right, as significant changes in substitution patterns over this time horizon could point to changes in economic primitives of interest (e.g. the shadow value of wealth). The second addresses the extent to which firms have taken advantage of these changes by charging higher markups. The third provides a potential structural explanation for rising differentiation. Addressing the second and third questions requires placing some additional structure on consumer demand and/or firms' pricing behavior, but allows us to relate the rise of variety directly to a growing literature on rising markups and profits.

Answering these questions convincingly relies on credible estimates of demand for many products at a granular level. This requires product-level data for a reasonably long panel of firms. Most recent discussions of markups rely on firm- or establishment-level data to estimate markups via the "production approach," which requires strong assumptions even in settings with data on quantities sold by single-product firms (Flynn, Gandhi and Traina, 2019; Brand, 2019; Jaumandreu, 2018). In the much more common setting in which the researcher only observes revenue, instead of quantities and prices separately, recent work by Bond et al. (2020) indicate that little if anything can be learned about markups from these methods. This is one way in which the data herein are ideal, as they cover a large number of products and firms and contain product-level prices and sales data which can be used to calculate product-level markups for each firm.

The second issue is that an empirical model of demand for thousands of products which span many categories should be as flexible and general as possible while remaining tractable and computationally feasible. In order to ensure that my results are robust to a variety of assumptions on consumer preferences, I estimate demand for nine product categories in three ways in Section 5 (and twice more in Section 8). In the first, I follow the canonical approach in the spirit of Berry, Levinsohn and Pakes (1995) (BLP), which estimates structural utility parameters in a discrete choice framework. This model permits significant unobserved heterogeneity differentiating products and stores, but assumes that the distribution of consumer preferences (e.g. over horizontal product characteristics) is the same at every store.[1] In a second approach,

[1] My inclusion of store-product fixed effects does permit the mean preferences for goods to vary across stores.

I estimate separate structural demand parameters for every three-digit ZIP code, for every category in every year, using an approximation to a structural model as developed by Salanié and Wolak (2019). This results in hundreds of estimated structural preference parameters each year for every product category in my sample. In a third, I focus instead on estimating more than 100,000 constant elasticity demand curves at the store-category level in an approach very similar to that taken by DellaVigna and Gentzkow (2019) in their study of uniform pricing.

My results indicate that consumers are substantially less willing to substitute between products in 2017 than in 2006. My estimates from a BLP-style model imply that own-price elasticities have declined from -2.17 to -1.65 at the median across all products in my sample. I find a similar pattern with respect to cross-price elasticities and within category in all but one of the nine studied categories. Using estimates of preferences which vary at the three-digit ZIP code level, I show that these declines appear to be in part the direct result of changes in consumer utility functions, specifically in the effective disutility of price. Next I demonstrate that under a monopoly pricing rule the observed substitution patterns and prices are consistent with an increase in markups in most modules, both at the median and along much of the interquartile range. Finally, I provide evidence for two predictions of the simple model in Section 2 regarding the role of niche products in explaining these trends. First, I re-calculate own-price elasticities under a counterfactual scenario in which preference heterogeneity in 2017 is reduced to its 2006 levels. I find that this eliminates much of the observed changes in price elasticities. Second I show that, in all but one of the modules I study, newer products tend to be more niche than older products. Altogether these results provide evidence for the importance of rising differentiation and the introduction of more niche products in changing substitution patterns and markups.

This paper is closely related to the recent literature suggesting that the variable profits of the largest firms in the United States have been rising since the 1980s (Barkai, 2020; De Loecker, Eeckhout and Unger, 2020). Existing explanations for the underlying causes have been varied, ranging from rising concentration (Autor et al., 2020), to increasing monopsony power (Stansbury and Summers, 2020). In general, the dominant mechanisms will likely vary by industry. For example, as outlined by Ganapati (2021), concentration in a number of sectors has increased in recent decades. However, prices have declined and output has increased in these sectors over the same period, both of which are counter to standard concerns about rising market power. Similarly, although Hsieh and Rossi-Hansberg (2019) make a convincing case

Still, the distribution of preferences for price and the random utility shocks are assumed not to differ across stores.

that the cost reductions brought on by investment in IT played a major role in these changes in the service industry, Grieco, Murry and Yurukoglu (2020) find that costs have risen in the U.S. automobile industry as a result of cars growing larger, heavier, and more durable. My findings offer another example of the ways in which investment in IT may have caused markups to increase, and offer a mechanism through which this may have occurred.

The closest existing paper to mine is Neiman and Vavra (2019), who document the fact that household purchases have become more concentrated at the same time that store-level sales have become less concentrated.[2] As in this paper, they attribute these facts to firms offering more products, which allows consumers to find goods which are closely tailored to their tastes. Their model aggregates nicely and permits a discussion which covers many more products than are considered herein. However, this comes at the cost of doing the analysis at the product category level rather than at the store or product level, which aggregates over a substantial amount of variation, and abstracts from some of the heterogeneity and changes in substitution patterns which can be incorporated in studies of fewer products like the present paper. This paper compliments theirs - by focusing on a smaller number of products (and paying a much higher computational cost) I am able to describe the full distribution of price elasticities at thousands of stores and discuss in more detail the ways in which consumer preferences have changed over time.

2 Simple Model of Optimal Differentiation

To set the stage for the remainder of the paper, in this section I discuss the likely effects of a firm developing a supply chain which is increasingly demand driven and which requires significantly less inventory to be stored, as has been the case for many firms in recent decades. I model these features of the supply chain as an opportunity for firms to (i) keep track of and better respond to key features of demand and (ii) sell more products/varieties in total, and provide intuition which indicates that these new and improved supply chains may have made stores more likely to sell products which are increasingly differentiated and niche, in a sense I define below.

[2] Kiedaisch, Chai and Rohde (2018) also study changing consumption patterns and variety, and argue for the importance of including consumer-level preference heterogeneity in studies of the benefits of variety.

2.1 Monopolist's Problem

Consider a monopolist choosing a menu of products to offer from some large set. The firm knows that, after choosing which portfolio of products to stock, it will set prices in order to maximize short-run profits. The less willing consumers are to substitute between products or to purchasing nothing, the more surplus the monopolist can extract from consumers in the form of high prices. Knowing this, the firm would like to stock products which are "differentiated" in exactly this sense. This point is related to a now significant literature studying firms' endogenous choices of products, and product characteristics, to offer to consumers (Waldfogel, 2003; Draganska, Mazzeo and Seim, 2009; Fan, 2013; Sullivan, 2017).

We can see this intuition somewhat more precisely in, for example, a short run profit maximizing monopolist's first order condition in the form of price derivatives of demand. Toward this end, in Equations 1 and 2 I present first order conditions for single and multi-product monopolists, respectively. In each equation, p^* represents the (vector) of optimal prices, mc denotes marginal costs, and $s(p^*)$ is the realized demand in the form of market shares.

$$p^* - mc = -\left(\frac{ds}{dp^*}\right)^{-1} s(p^*) \qquad \text{(single product)} \tag{1}$$

$$p^* - mc = -J_p^{-1}(p^*)s(p^*) \qquad \text{(multi-product)} \tag{2}$$

The term $J_p(p^*)$ denotes a matrix of derivatives of demand with each (i, j) element equal to $\frac{\partial s_i(p^*)}{\partial p_j}$, making $J_p^{-1}(p^*)$ analogous to $(\frac{ds}{dp^*})^{-1}$ in Equation 1. According to these first order conditions, monopolists will be able to charge higher prices (and earn larger margins) by offering products for which consumers are less price sensitive, as measured by price derivatives of demand.[3]

Following this logic, the more differentiated are products, the more the monopolist can (i) set high prices on many goods without losing customers and (ii) price discriminate across consumers, ensuring that each consumer purchases a product with a large margin for the monopolist.[4] The firm can achieve its goals by offering goods differing on either vertical (e.g. quality) or horizontal (e.g. niche product characteristics) dimensions but the end result is the same. This brings us to the first prediction of the paper. Given the growing emphasis of retail

[3]Of course, the derivatives of demand which appear in these equations are functions, so defining "less price sensitive" precisely would require some additional notation.

[4]Note that the off-diagonal elements of J_p^{-1} are important largely because, to the extent that marginal costs differ across products, minimal cross-product substitution means that the monopolist can charge higher markups on goods with lower marginal costs.

supply chains on quick responses to real-time consumer demand, the increased efficiency of these supply chains, and the resulting rise in the number of products on store shelves, it seems natural to suspect that firms have had an increasing number of opportunities, and an increasing technological ability, to change the assortment of products they offer in order to increase differentiation and thereby increase markups.

2.2 Differentiation as Stocking Costs Decline

An important effect of improved supply chain efficiency is that firms are not required to keep large inventories of products within a store (Sparks, 2010). Instead, products are often kept at chain-owned distribution centers and delivered to stores in smaller frequent trips as inventories run low. This means that the effective costs of stocking additional products in a store have declined, as the firm need not have significant inventory space within the store in order to offer a new good. Naturally, this changes the nature of a monopolist's problem, in that it can now offer a wider selection of products. This raises the question: how will the attributes of newer products (those added due to reduced stocking costs) differ from those offered before these technological improvements?

Both horizontal and vertical differentiation may be beneficial to the firm, as both imply that consumers will be willing to pay at higher prices. In a sense, firms always wish to offer high quality goods. A sufficiently high quality good can induce all consumers to purchase at some price (which naturally depends on the strength of preferences for quality). Normally, however, we think of vertical differentiation as being costly. In empirical studies of manufacturing productivity, for instance, higher quality goods are assumed to require higher quality inputs (De Loecker et al., 2016). This may constrain the extent to which firms can increase vertical differentiation even as stocking costs decline, because wholesale costs for high quality goods will remain high.

To the contrary, introducing a new horizontal attribute like flavor, or changing product branding, may entail a much smaller cost. An important question then, is how the optimal level of horizontal differentiation differs with the number of products offered, and in particular how the characteristics of the marginal good differs as the firm's portfolio size grows. To study this question, consider a setting in which the utility a consumer i receives from purchasing a good j at price p_j is represented by

$$u_{ij} = -p_j + \epsilon_{ij}$$

For simplicity, let ϵ_{ij} be independent across goods for each consumer.[5] A consumer chooses to purchase one of J inside options or one outside option with $u_{i0} = 0$ by choosing the option which maximizes her utility. In this form, ϵ_{ij} is each consumer's willingness to pay for product j. Now suppose a monopolist can only afford to stock a single product (e.g. because doing so requires maintaining a large in-store inventory). In choosing the type of product it will sell, the monopolist must choose among two options:

$$\begin{aligned} \text{Staple} &: \epsilon_{ij} = 1.5 \text{ for all } i \\ \text{Niche} &: \epsilon_{ij} = 1 \text{ or } 2, \text{ each with probability } 0.5 \end{aligned}$$

In words, all consumers will pay 1.5 for any Staple good, while for Niche goods half of all consumers will pay 2 while half will only pay 1. For each consumer, ϵ_{ij} is independent across goods, meaning a consumer's willingness to pay for one Niche good does not predict their preferences for another.[6]

I call these "types" of products to emphasize that there are many goods of each type to choose from. Staple goods are meant to represent goods which many consumers are willing to purchase but for which no consumers are perfectly matched. Think of goods like Coca-Cola and Pepsi, which are ubiquitous in grocery stores and which are sold to a relatively large fraction of soda buyers. Niche goods represent goods which are more narrowly tailored to a subset of consumers' tastes. Continuing with sodas as an example, this might include things like Orange Vanilla Coke Zero. For many consumers, any of "Orange," "Vanilla," or "Zero" might make the product undesirable, but for some consumers these descriptors will make the product a better fit to their tastes. These consumers will be willing to pay more for the Niche good, and thus will be less likely to substitute to other products when the price of the Niche good increases (e.g. above the Staple good price). In this example, the firm will choose to offer a Staple product, as it can set a price of 1.5 and sell to all consumers, whereas a Niche product can only be sold to all consumers at a price of 1 or half of consumers at a price of 2.

In this model, this decision changes as stocking costs decline (i.e. as the number of products offered increases). If the monopolist can afford to stock two products instead of one, then it can either again sell to all consumers at $p = 1.5$ (by offering Staples), it can sell to 75% of consumers at a price of 2 (by offering Niche products), or it can offer one of each type of product and sell

[5] This is much stronger than necessary. As long as ϵ_{ij} are not perfectly correlated, adding more niche products becomes more desirable as J increases.

[6] The intuition here holds much more generally, as long as ϵ_{ij} are not perfectly correlated across goods.

to 50% of consumers at a price of 2 and the remaining 50% at a price of 1.5. The latter option maximizes profits, meaning the second product the monopolist stocks will be Niche. In this example, in fact, all products stocked beyond the first will be Niche. As stocking costs decline, then, two predictions follow: a growing share of stocked products will be Niche, and average prices will increase.[7]

3 Data

3.1 Weekly Scanner Data

The first dataset I use comes from the Nielsen Corporation through an agreement with the Kilts Center at the University of Chicago. This dataset contains product-level sales and average price data from 2006 to 2017 for thousands of stores and dozens of retail chains which span the United States.[8] These data are reported weekly and are exhaustive of all products sold. To make it into my working data set for a given product category, a store must be a "food store" or "mass merchandiser" and must appear in the Nielsen data selling at least one good in that category in the first 16 weeks of both 2006 and 2017. Focusing on food stores and mass merchandisers is natural, as these stores tend to be larger than other (i.e. drug and convenience) stores in the Nielsen data, and thus better represent the rise in variety being studied.

As for the second restriction, there have been a number of recent papers discussing changes in retail concentration and the entry of new, higher-productivity retail stores, which are not the focus of this paper (Foster, Haltiwanger and Krizan, 2006; Rossi-Hansberg, Sarte and Trachter, 2018; Smith and Díaz, 2020). By focusing on a balanced panel of stores for each category, we reduce the likelihood that these mechanisms explain the results herein.[9] Among these stores, I deflate all prices to 2017 dollars using the CPI.[10] I restrict all of my analysis to the first 16 weeks of the year, and to the 300 most popular goods (by the total number of purchases in the

[7] In this example, the monopolist never drops its single Staple good. This is because, for any monopolist with at least 2 products, offering at least one Staple good ensures that every consumer pays at least their willingness to pay for the Staple (i.e. all consumers purchase a good). Permitting a small number of consumers to dislike the Staple good solves this issue at the cost of making the example less clear.

[8] Use of this data set and the home scanner data below have become very common within economics. For a more detailed description of these data, see Hitsch, Hortacsu and Lin (2019).

[9] For example, if a large chain in the scanner data had expanded its geographic presence substantially since 2006 and had also reduced prices substantially, this could reduce estimated price elasticities in the aggregate.

[10] Note that, because we are most interested in price elasticities, using the wrong inflation measure (or, in fact, not deflating at all) should not impact our conclusions. Any multiple on all prices is included in the denominator and numerator of elasticities, meaning it will cancel out. I show this empirically in Appendix C.

year across all stores in my sample) in each product module each year, largely as a solution to computational issues.[11] In this paper, I specify a product as a unique Universal Product Code (UPC), which is the finest level at which the Nielsen data differentiates products within a year.[12] I present some summary statistics of this scanner data in the next section.

3.2 Nielsen Homescan Consumer Panel

The second dataset I use is the Nielsen Corporation's Homescan Panel data, which consists of a rotating panel of consumers who are asked to scan all of their purchases from retailers. This data is available to me though an agreement with the Kilts Data Center at the University of Chicago, and includes tens of thousands of households across the country and millions of transactions. I use these data for two central purposes. First, because consumers are followed through many purchases, they provide vital information about the frequency with which consumers choose the outside option (i.e. not purchasing from a given product category). This allows me to define market sizes (and therefore market shares). Second, I use these data to test for behavior which might bias my estimates of demand. I describe this test in Section 3.3.

To measure market sizes, for each product category in the panel I calculate the proportion of consumers who buy a product in that module per week out of the total number of consumers who buy from the category over the course of the year. I calculate this proportion for each week and take the average over the year. Then, I use this average value to scale all market shares (i.e. shares among inside options) calculated using the weekly scanner data (above). These calculations, while imperfect, are straightforward and intuitive. If 100 consumers purchase yogurt each year, but only 50 purchase in the average week, then perhaps something close to 50 consumers choose the outside option each week. For each store, for each category, I multiply this ratio (2 in this example) by the maximum number of units purchased in that store-category in any week of the year, which defines the market size for that store in that module. Many classic studies use more ad-hoc definitions of market size, such as population in a nearby region (e.g. the entire population of the United States, or of a ZIP code). As such, although defining market size is often a difficult decision in the estimation of structural demand models, this approach seems quite reasonable. I also present one set of demand estimates which are robust to misspecification of market sizes.

[11] I estimate store-level logit demand models for a subset of categories which include all products in the scanner data in Figure A.1. My results are unchanged by including these additional products.

[12] Technically, a UPC-version code combination is the finest level at which the data differentiates products, but this distinction will be of little importance to most of my analysis as I do not generally assume that all attributes of UPCs remain constant over time.

3.3 Selecting Product Categories

Products in the Nielsen data are divided into hundreds of categories called "modules" (moving forward, I use "categories" and "modules" as synonyms). These modules tend to define large but reasonable sets of substitutes. For example, refrigerated yogurt and fresh eggs are each a product module. Given the huge number of products included in the Nielsen scanner data, and the computational costs of some of the structural models estimated herein, studying all modules is unrealistic. To construct my sample of products to study, I begin with the 41 modules studied by DellaVigna and Gentzkow (2019) which are sold by mass merchandisers and/or food stores, a set which covers a large and diverse array of products. Even this is too many modules to include in this study, which raises the question of whether there is a principled way to select which categories to study among this subset.

The approach I take in this vein is to exclude modules which are likely to yield biased estimates due to a particular form of dynamic behavior. One issue which has been raised in the literature on demand estimation, particularly in retail, is that estimating demand in settings in which consumers tend to keep an inventory of products will often lead to biased estimates of price elasticities. This is largely due to the fact that, if consumers can keep a large enough inventory, they can decide to buy many units when a product is on sale and none otherwise. This type of behavior can generate demand estimates that imply very large short run price elasticities, even if long run price elasticities are small (Hendel and Nevo, 2006). With this in mind, only product categories in which consumers do not keep a large stock of products or target sales should be included in the main sample in this paper.[13] In order to exclude any modules in which consumers exhibit this type of behavior, I conduct a test which draws on the arguments in Hendel and Nevo (2006): if consumers are unable to store goods, then a consumer's purchase amount should not predict the time until their next purchase. For each module, I estimate a regression motivated by this intuition using the consumer panel data and test the null hypothesis that the previous purchase amount does not predict the time between purchases (more details are included in Appendix B). I then keep only the nine modules with the largest p-values for this test (all greater than 0.5). I list these nine modules in Table 1 as

[13]Of course, another option would be to introduce a dynamic model of demand instead of excluding some modules. The computational costs of dynamic demand models are often severe, and would make much of the analysis herein infeasible.

well as the number of stores in my sample selling each category.[14,15]

Table 1: Modules Without Evidence of Storage

Module Description	# Stores
FRUIT DRINKS - OTHER CONTAINER	10994
SOUP-CANNED	10911
COOKIES	10995
PIZZA - FROZEN	7528
ICE CREAM - BULK	7393
ENTREES - REFRIGERATED	5613
YOGURT - REFRIGERATED	5742
FRESH FRUIT - REMAINING	4733
LIGHT BEER (LOW CALORIE/ALCOHOL)	4138

Note: Modules chosen following procedure in text. Full table including all considered product modules can be found in the Appendix.

4 Motivating Empirical Facts

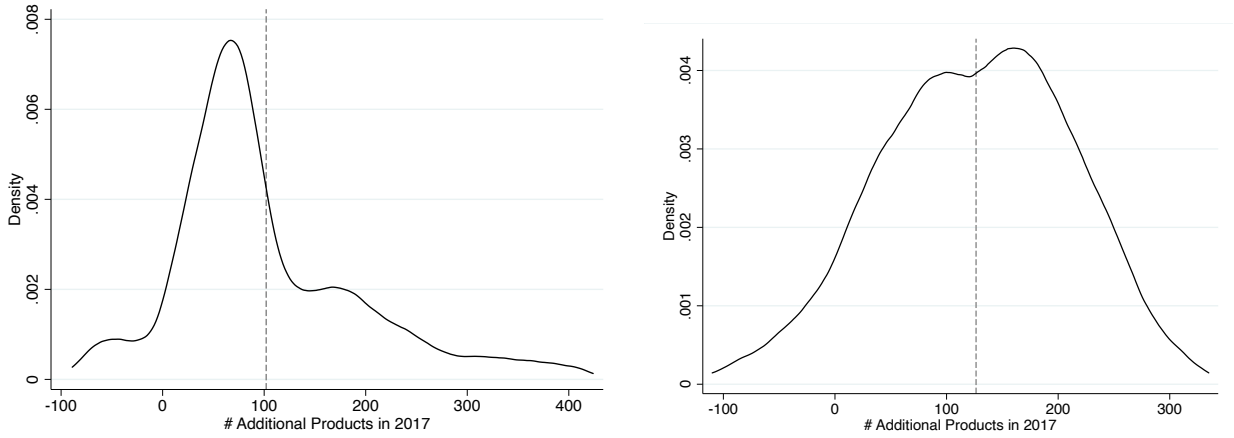
One of the most significant changes in retail over recent decades has been a drastic increase in the use of information technology to manage supply chains and keep track of inventory. The growing use of IT in supply chain management has been part of a trend in retail firms toward demand-driven supply, often following dominant frameworks such as just-in-time/quick response and efficient consumer response. See Fernie, Sparks and McKinnon (2010), for example, for a full review of these changes with a focus on the United Kingdom. These approaches tend to involve vertically integrating and passing data up the supply chain, such that inventories are restocked at a rate which closely matches actual demand in downstream stores (Fernie and Azuma, 2004). In the retail fashion industry, for example, these methods have brought on significant cost reductions and faster production, leading to a quicker turnover in stores' inventory as consumer tastes change (Bhardwaj and Fairhurst, 2010). In grocery retail stores, and especially in the largest stores, the rise of IT technologies like RFID has made it easier to track inventory across levels of the supply chain (Shin and Eksioglu, 2014).

[14]Though most categories are given intuitive names, two which may require clarification are "Fresh Fruit - Remaining" and "Fruit Drinks - Other Container." The former contains the large majority of fruits in grocery stores, and appears to only exclude a small handful of popular fruits (e.g. apples and oranges). The latter, on inspection, appears to contain most fruit-flavored drinks, including sports drinks and pure fruit juice blends.

[15]I also show that my results are not dependent on this particular approach to selecting categories by calculating some results for additional product categories in Figure A.1.

According to a *Consumer Reports* article from 2014 citing numbers from the Food Marketing Institute, the average number of products in a supermarket increased from less than 9,000 in 1975 to almost 47,000 by 2008 (Consumer Reports, 2014). The Nielsen data for the categories I study, which cover the following decade, are consistent with the continuation of this trend. In the first week of 2006, the average store sold 385 of the products in my sample. Stores selling all nine categories sold on average 804 products. In 2017, these numbers were 488 and 930, respectively. These means mask a significant amount of heterogeneity, which I demonstrate in Figure 1. For each store in my sample I calculate the difference in the number of products sold (across all nine modules) between 2006 and 2017. In Figure 1(a) I plot the empirical density of these differences across all stores in my sample, and in Figure 1(b) I do the same after restricting the sample to stores which sold all nine modules at least once. The mean is more than 100 new products in each figure, as denoted by the vertical grey lines, but many stores have increased their product selection in these modules by 200 or more products. Both figures indicate that, even among this small subset of modules, the number of products sold in many stores in my sample has increased substantially between 2006 and 2017.[16]

Figure 1: Number of Additional Products in Sample Modules in 2017



(a) (b)

Note: Constructed from Nielsen scanner data. Figure (a) includes stores all stores in my sample, and Figure (b) restricts the sample to stores which sold a product in each of the nine sample modules at least once in my sample. Top and bottom 1% of the distributions have been trimmed.

Next, we should discuss the ways in which the (price and non-price) characteristics of products sold have changed over time. The first piece of evidence in this direction comes from the United States Department of Agriculture Economic Research Service (USDA ERS). In Figure

[16] I also show the full distribution of the number of products sold separately by module in Figure A.5.

2 I reproduce two tables constructed by the ERS which describe a subset of characteristics of new products as indicated by the Mintel Global New Products Database (GNPD), via "tags" indicating product characteristics. The Mintel GNPD is a large database of products which claims to add approximately 40,000 new products each month and covers a significant fraction of products sold in over 80 countries. In Figure 2(a) , we can see that according to this database the number of new products each year marked as "kosher," "gluten free," "organic," and "low/no/reduced allergen" has grown dramatically, often by a factor of 2 or more. Figure 2(b), which presents these trends as a fraction of all new products in the GNPD each year, demonstrates that the number of products with many characteristics (e.g. "GMO free") have grown not only in levels but as a proportion of new products.

The Mintel data are meant to cover dozens of countries, meaning they are not necessarily representative of the stores in my sample. The data Nielsen provides regarding the characteristics of all products sold by retailers in their data are much less detailed, meaning the extent to which we can match these trends in the Nielsen data is very limited. What can be said is that the number of organic products in the scanner data has more than quadrupled between 2006 and 2017, and the number of unique flavors has more than doubled. At the same time, the total number of flavored products has actually declined, meaning the number of unique flavors per flavored product has increased substantially over time. Though more detailed product characteristics would be greatly beneficial both here and in estimating demand, it is encouraging that these two measurable dimensions of product characteristics are consistent with the increasing prevalence of new horizontal product characteristics.[17]

We will close this section with a discussion of prices and the size of the outside option over time. Summary statistics of price in this setting can be difficult to interpret, as consumer price indices often include grocery items like those in my sample. Thus, we should expect that changes in real prices should on average be small. Moreover, as inflation has been relatively small, nominal prices should also have been relative constant over time. Still, for completeness, in Table 2 I calculate the market share-weighted (i.e. sales weighted) price in every store-week, and present the average across stores for each module.[18] There appears to be little to

[17]For exact numbers associated with these claims, see Table A.1. I construct this table by matching annual version files to the master file containing the full set of products in the scanner data. I consider a product as having been sold in a given year if the versions file for that year can be matched to the file containing all products.

[18]Any choice of weights to aggregate so many prices has its flaws. Weighting by market shares confounds changes in consumer behavior with prices. Uniform weights on all prices would place too much weight on less popular products, which make up a large fraction of unique products sold but a small fraction of purchased goods. Cost weights would be ideal but I lack the data for such an approach.

Figure 2: Characteristics of New Products, USDA ERS

Number of new product introductions in the top 10 product claim categories for 2009-16

Tag or claim*	2009	2010	2011	2012	2013	2014	2015	2016
Kosher	4,159	6,164	5,606	5,386	7,570	7,942	7,423	8,985
Low/no/reduced allergen	1,325	2,215	2,228	2,250	3,930	4,828	4,914	6,552
Gluten free	1,121	1,942	1,994	2,000	3,609	4,550	4,534	6,123
Ethical-environmentally-friendly package	1,329	2,892	2,806	2,903	4,254	4,268	4,239	5,056
No additives/preservatives	2,068	2,993	2,647	2,524	3,544	3,549	3,471	4,591
Social media	1	3	0	808	2,118	2,599	2,871	3,994
GMO free	297	340	550	567	1,352	1,993	2,685	3,732
Organic	1,445	1,548	1,332	1,279	2,097	2,084	2,313	3,011
Microwaveable	1,724	2,279	1,827	1,706	2,531	2,530	1,749	2,287
Ease of use	903	1,610	1,319	1,401	2,133	2,062	1,700	2,287
Total new product claims	32,300	47,303	45,090	44,374	64,133	67,001	63,320	79,779

*A new product may have multiple tags or claims.
Source: Mintel GNPD.

(a)

Percent of new product introductions in the top 10 product claim categories for 2009-16

Tag or claim*	2009	2010	2011	2012	2013	2014	2015	2016
Kosher	12.9	13.0	12.4	12.1	11.8	11.9	11.7	11.3
Low/no/reduced allergen	4.1	4.7	4.9	5.1	6.1	7.2	7.8	8.2
Gluten free	3.5	4.1	4.4	4.5	5.6	6.8	7.2	7.7
Ethical-environmentally-friendly package	4.1	6.1	6.2	6.5	6.6	6.4	6.7	6.3
No additives/preservatives	6.4	6.3	5.9	5.7	5.5	5.3	5.5	5.8
Social media	0.0	0.0	0.0	1.8	3.3	3.9	4.5	5.0
GMO free	0.9	0.7	1.2	1.3	2.1	3.0	4.2	4.7
Organic	4.5	3.3	3.0	2.9	3.3	3.1	3.7	3.8
Microwaveable	5.3	4.8	4.1	3.8	3.9	3.8	2.8	2.9
Ease of use	2.8	3.4	2.9	3.2	3.3	3.1	2.7	2.9

*A new product may have multiple tags or claims.
Source: Mintel GNPD.

(b)

Note: Tables are reprinted from the USDA ERS website, and were constructed by USDA ERS using data from the Mintel Global New Products Database. Downloaded by the author from https://www.ers.usda.gov/topics/food-markets-prices/processing-marketing/new-products/

no pattern in changes in prices over time. Other than in the Yogurt and Remaining Fruits categories, changes in average prices are quite small relative to 2006 prices.[19] This indicates

[19]Price dispersion is also important to document here, as increasing dispersion could potentially indicate

that, to the extent that markups have risen, much of this growth must be due to cost reductions.

Finally, in columns 2 and 4 of Table 2 I present the average (across stores and weeks) of the market share of the outside option (i.e. no purchase within the category) for each module. If a module had become much more desirable on average over my sample period, or if products had become significantly higher quality, the size of the outside option would be smaller, and substitution to the outside option would generally be weaker. If this were the case, estimates of changes in substitution patterns over time might capture this feature in addition to the introduction of niche goods. As is the case with prices, the average size of the outside option has changed little, except in the Yogurt and Remaining Fruit categories. The similarity of columns 2 and 4 implies that this type of differentiation is unlikely to drive results below, though I return to a brief discussion of vertical differentiation at the end of Section 8.

To summarize, in this section I have shown that the number of products sold in retail stores in my sample has increased significantly over recent decades, at the same time as retail supply chains have become increasingly streamlined and demand focused. Over a similar period, the number of products with new, or previously rare, characteristics has also grown globally, and within the Nielsen data to the extent that this can be measured. To the contrary, prices and the size of the outside option have moved relatively little over time. Thus, although simplistic measures of horizontal characteristics are consistent with rising variety and/or the nicheness of retail food products, evidence for changes in vertical attributes and for rising prices is much weaker. If we find rising differentiation, then, the primary focus will be on horizontal differentiation. Moreover, if pricing patterns over my sample period are consistent with rising markups, that trend most likely to be due to declining marginal costs.

5 Structural Model of Demand

In order to measure the extent to which products have become more differentiated over time we require estimates of the extent to which consumers are willing to substitute between products. More precisely, we need estimates of the own- and cross-price elasticities of demand of many products. Toward that end, I estimate a number of models of demand which permit flexibility in differing directions, most of which fall into the following class of demand systems. I assume that, after choosing a store at which to shop, consumers select among the available

rising price discrimination even in the absence of changes in average prices. For example, if a firm introduced a very high and a very low quality good, it could set a high price for the former and a low price for the latter in equilibrium. Thus, differentiation would generate larger price dispersion. I show in Figure A.2 that this does not appear to have been the case.

Table 2: Share Weighted Price and Outside Option Size, by Module and Year

	2006		2017	
	s_0	Weighted Price	s_0	Weighted Price
Fruit Drinks	0.942	2.13	0.941	1.88
Soup	0.931	1.53	0.942	1.57
Cookies	0.910	2.22	0.914	2.20
Pizza	0.941	3.47	0.942	3.59
Ice Cream	0.939	3.86	0.946	3.75
Entrees	0.960	5.85	0.958	5.46
Yogurt	0.902	0.89	0.877	1.43
Remaining Fruit	0.973	4.43	0.913	3.70
Light Beer	0.942	11.56	0.945	12.08

Note: Selected summary statistics for nine modules. The "Weighted Price" column presents the average (across stores and weeks) of the market share-weighted price among goods in my sample, and s_0 presents the average size of the outside option market share across stores and weeks.

products within a product module to maximize indirect utility functions of the form

$$\begin{aligned} u_{ijsw} &= \alpha_{i(z)} p_{jsw} + \xi_{jsw} + \epsilon_{ijsw} \\ \alpha_{i(z)} &\sim N(\bar{\alpha}_{(z)}, \sigma^2_{\alpha,(z)}) \\ u_{i0sw} &= \epsilon_{i0sw} \end{aligned}$$

I use the indices i, j, s, and w to denote consumer, products, store, and week respectively. The parenthetical subscript (z) refers to three-digit ZIP codes. Consumers are differentiated by two terms: α_i and ϵ_{ijt}. The coefficient $\bar{\alpha}_{i(z)}$ represents the disutility of price for consumer i and is modeled as a normal distribution. The subscripts (z) are used to indicate that in some results, which I describe more in future sections, I permit both the mean and variance of α_i to differ in each three-digit ZIP code. Consumers also have idiosyncratic preferences ϵ for each product in each week (including the outside option), which are distributed according to a Type 1 Extreme Value distribution with scale parameter σ. I also permit all consumers in in each store to observe an unobserved product- and market-specific characteristic ξ_{jsw} which is not observed by the econometrician. In this model, market shares for each product take the following form

$$s_{jsw} = \int \frac{\exp(\frac{\alpha_{i(z)} p_{jsw} + \xi_{jsw}}{\sigma_{(z)}})}{1 + \sum_k exp(\frac{\alpha_{i(z)} p_{ksw} + \xi_{ksw}}{\sigma_{(z)}})} d\alpha_i \tag{3}$$

This model permits product differentiation to change between 2006 and 2017 through three channels. First, the variance of ϵ may change over time. The associated scale parameter σ is the key parameter controlling horizontal differentiation and therefore the "nicheness" of products. The larger is σ, the larger the variance of ϵ (relative to p_{jst}), and the less responsive is demand for a given product to changes in price. This can be seen directly in Equation 3, where clearly larger values of σ reduces the effective coefficient on price ($\frac{\alpha_i}{\sigma}$). In the fullest generality, the variance of ϵ might differ across products within a year, as some products are naturally more broadly appealing (i.e. less niche) than others. With the large number of products in my data, estimating such a model is infeasible. Although I estimate a model with more than one scaling parameter in Section 8, for now I will think of σ variance as capturing year-level average preference heterogeneity across products. Changes in this variance will be particularly important if more niche products have been introduced over time.

Products can also be vertically differentiated by the unobserved product characteristic ξ, which permits products to be weak substitutes even conditional on price and preference shocks. In most of the demand models I estimate, ξ will take the form

$$\xi_{jsw} = \bar{\xi}_{js} + \bar{\xi}_{(s),(z),w} + \Delta\xi_{jsw} \tag{4}$$

where $\bar{\xi}_{js}$ denotes average preferences for product j (relative to the outside option) at store s, and $\bar{\xi}_{(s),(z),w}$ represents the average preferences for the outside option in a given week, which I also let vary either at the store or three-digit ZIP code level. The former permits consumers, on average, to rank products differently at each store, which might occur for a number of reasons including geography-specific advertising, consumers sorting across chains, or heterogeneous willingness to pay for quality (e.g. due in part to wealth). The latter is a nonparametric way to capture seasonality in average preferences for all goods within a category. This seems particularly important for seasonal goods like Ice Cream, but may also play a role in other categories.

Finally, the distribution of α_i is crucial in determining consumers' willingness to trade money for utility. If consumers at the stores in my sample are significantly wealthier in real terms in 2017 relative to 2006 (and relative to grocery prices) then they may be less price sensitive, all else equal. This would be captured in my estimates largely through $\bar{\alpha}$ and σ_α, where the former measures average preferences for price and the latter measures preference heterogeneity within a market. In some estimates, I will also permit the distribution of α_i to differ by geography (three-digit ZIP code or store), thereby capturing differences in wealth and

other demographics by geography to the extent that they impact price consumers' sensitivities.[20]

6 Estimation

6.1 BLP

Estimation of demand systems like the baseline model herein has been studied by a substantial literature since the original "BLP" paper (Berry, Levinsohn and Pakes, 1995). Current best practices are implemented in a new Python package called *pyblp*, written by Chris Conlon and Jeff Gortmaker (Conlon and Gortmaker, 2019). This package provides empirical IO researchers with an easily customizable way to estimate this standard demand model, and implements cutting edge numerical tools to speed up estimation. Though the literature describing the estimation of this type of model is long and in-depth, for completeness I describe the procedure below as well.

We can rewrite the random utility model in Section 5 as follows:

$$u_{ijsw} = \delta_{jsw} + \sigma_\alpha \nu_i + \epsilon_{ijsw} \tag{5}$$

where $\delta_{jsw} = \bar{\alpha} p_{jsw} + \xi_{jsw}$ is the common component of utility derived from product j in store s in week w, ν_i is a standard normal random variable, and ϵ is an idiosyncratic error term which follows a Type 1 Extreme Value distribution with scale parameter 1. Although I normalize the scale parameter of ϵ in each module-year, by letting both the price coefficients α_i and the unobserved characteristics ξ vary across modules and time I am implicitly letting the variance of preferences vary as well. This is because the scale of utility and the variance of ϵ are not separately identified without some additional structure (to see this, note that doubling α_i, ξ_{jsw}, and the scale parameter σ leaves Equation 3 unchanged). Changes in estimates of $\bar{\alpha}$ should therefore be interpreted as some combination of changes in the disutility of price and changes in the variance of ϵ. In Section 8 I perform a decomposition which measures the relative importance of these two components of utility in explaining changes in price elasticities over time.

For each guess of the "nonlinear" parameter σ_α, which again captures heterogeneity in the

[20] Some readers may have a preference a nested logit model in which all inside goods constitute a nest. The fixed effects I include accomplish some of the goals of this approach, and I find that nested models are much less likely to converge given my inclusion of product-store fixed effects. In Figures A.3 and Figures A.4 I show one case in which a nested logit model does converge (for Light Beer) with UPC and week fixed effects. The results are similar to those in the next section.

disutility of price across consumers, we can construct estimates of market shares

$$\hat{s}_{jsw} = \int \frac{\exp(\delta_{jsw} + \sigma_\alpha \nu_i)}{1 + \sum_k exp(\delta_{ksw} + \sigma_\alpha \nu_i)} d\nu_i \tag{6}$$

I use this mapping iteratively (each iteration denoted by k) to solve for δ using the following fixed point problem[21,22]

$$\delta_{jsw}^{k+1} = \delta_{jsw}^{k} + log(s_{jsw}) - log(\hat{s}_{jsw})$$

where $\hat{s}_{jsw}$ is the model predicted vector of market shares. After δ_{jsw} converges for each market we can recover α by regressing δ_{jsw} on price (usually instrumenting for price). This regression also generates estimates of ξ_{jsw}, which can be used to enforce the following moment conditions to identify α and σ.

$$\mathbb{E}[\Delta\xi_{jsw} Z_{jsw}] = 0, \tag{7}$$

where $\Delta\xi_{jsw}$ and Z_{jsw} denote the residuals of ξ_{jsw} and the vector (p_{jsw}, p_{jsw}^2) (respectively) after absorbing the fixed effects described in Section 5. In Section 6.4 I offer arguments that price serves as a valid instrument in this and following models, and I also present one set of estimates for the model in Section 6.2 which instead use the average price of good j in nearby stores as a cost-shifter (assuming wholesale costs are spatially correlated) in Appendix Figure A.9.

6.2 FRAC

One important critique of the preceding model, especially in this setting in which I observe so many stores around the country, is that I do not permit the distribution of α_i to differ geographically. In reality, consumer price sensitivity likely substantively differs across regions of the country. The most direct extension of the preceding model to permit geographic heterogeneity would be to estimate a different pair of distributional parameters α and σ in each three-digit ZIP code (the finest level of geographic information in my data). This would require estimating hundreds of parameters via GMM (jointly or separately by geography), which is a complicated computational problem. In particular, after attempting this for some modules, the problem is that the contraction mapping necessary to deal with preference heterogeneity (σ_α) is still relatively time consuming even when the sample is small. As a result, estimating BLP

[21] Following the frontier of the literature (and the default in *pyblp*), this problem is solved via the *squarem* algorithm (Varadhan and Roland, 2008; Reynaerts, Varadha and Nash, 2012).

[22] I calculate the integral in Equation 6 using Gaussian quadrature.

separately by three-digit ZIP code is too computationally burdensome to attempt here.[23]

Recently, Salanié and Wolak (2019) introduced what they call the Fast, "Robust," and Approximately Correct (FRAC) procedure for the estimation of mixed logit demand models. Their paper shows that, just as logit models (i.e. models with no random coefficients) can be estimated via an appropriately specified linear IV regression, random coefficient models like the one considered here can be *approximated* by a very similar regression. Their work offers the following approximation, similar to the frequently used logit inversion, for the demand model herein

$$log\Big(\frac{s_{jsw}}{s_{0sw}}\Big) = \bar{\alpha}p_{jsw} + \sigma^2_\alpha K_{jsw} + \xi_{jsw} + O(\sigma^4), \tag{8}$$
$$\text{where} \quad K_{jsw} = \Big(\frac{p_{jsw}}{2} - e_{sw}\Big)p_{jsw}$$
$$e_{sw} \equiv \sum_{j=1}^{J} s_{jsw}p_{jsw}$$

To estimate this model, one need only construct the regressor K in Equation 8 and run an IV regression (instrumenting for p and K) which will yield estimates of σ^2_α and $\bar{\alpha}$.[24] Absorbing fixed effects as in the preceding BLP model is straightforward, and without imposing additional constraints there is even a closed form solution for the parameters of interest. Thus, introducing geographic heterogeneity becomes much more feasible under this approximation. We can estimate such a model by estimating the following approximately correct regression for each three-digit ZIP code:

$$log\Big(\frac{s_{jsw}}{s_{0sw}}\Big) \equiv y_{jsw} = \bar{\alpha}_z p_{jsw} + \sigma^2_{\alpha,z} K_{jsw} + \xi_{jsw} \tag{9}$$

where $\bar{\alpha}_z$ and $\sigma^2_{\alpha,z}$ are again the mean and variance of random coefficients on price, which I now allow to differ in each three-digit ZIP code in my sample. With between 500 and 800 three-digit ZIP codes represented in each module, this approximate model permits drastically more heterogeneity in price elasticities across markets than the preceding model, at the cost of

[23]The route often taken in studies of random coefficient models like the one here is to include demographic characteristics in the utility function. In some initial work (not included here) I found that the sample of consumers in the household panel data is insufficient to estimate the store-level demographic distribution precisely. One could aggregate demographics to the three-digit ZIP code level, but this involves a high computational cost and is not necessarily more flexible than my FRAC estimates, which are robust to misspecification of the distribution of random coefficients.

[24]Note that, even when prices are exogenous, OLS estimates of Equation 8 will be biased in general, as the constructed regressor K includes market shares which are endogenous by construction.

requiring the estimation of far more parameters.

In estimating this FRAC model in this setting, I often find estimates of $\bar{\alpha}_z$ which are positive and $\sigma^2_{\alpha,z}$ which are negative, seemingly because p_{jsw} and K_{jsw} are very highly correlated (nearly co-linear in some ZIP codes). The former is inconsistent with basic economic theory, and the latter is nonsensical. Broadly speaking, I find one or both of these issues in 10% (or more) of the ZIP codes in my samples. There are a number of ways to deal with these estimates, some more ad-hoc than others. The approach I take is to convert the original regression to a constrained GMM problem, in which I constrain the mean price coefficient $\bar{\alpha}_z$ to be nonpositive and the variance $\sigma^2_{\alpha,z}$ to be nonnegative.

In practice, I estimate this constrained FRAC model separately for each three-digit ZIP code, dropping ZIP codes with fewer than 500 total observations in a module-year. I include product-store and (ZIP code-)week fixed effects, which I absorb in advance using the methods (and author-written code) introduced by Somaini and Wolak (2016). Let y^+, p^+, and K^+ denote the vectors $log\left(\frac{s_{jsw}}{s_{0sw}}\right)$, p, and K after absorbing this set of fixed effects. Further, let $Z^+ \equiv (p^+, p^{2(+)}, p^{3(+)})$ denote the vector of polynomial terms in price (after absorbing fixed effects) which I use as instruments.[25] Estimation for each three-digit ZIP code then simply involves searching over $\alpha_z, \sigma_{\alpha,z}$ to solve

$$\min_{\alpha_z,\sigma_z} \hat{\mathbb{E}}\left[\Delta\xi_{jsw} Z^+_{jsw}\right]' W \hat{\mathbb{E}}\left[\Delta\xi_{jsw} Z^+_{jsw}\right] \tag{10}$$
$$s.t. \quad \alpha_z \leq 0$$
$$\sigma^2_{\alpha,z} \geq 0$$
$$\Delta\xi_{jsw} \equiv (y^+_{jsw} - \alpha_z p^+_{jsw} - \sigma^2_{\alpha,z} K^+_{jsw})$$

Where $\hat{\mathbb{E}}$ denotes the sample mean and W is a weighting matrix which is initially set to $(Z'Z/N)^{-1}$ and is updated in a second step following the standard two-step efficient GMM procedure. I find that enforcing these constraints performs very well, and the constraints on $\bar{\alpha}_z$ are almost never binding in most product modules. I do find that estimates of $\sigma_{\alpha,z}$ are very close to zero in some modules.[26] After estimation, I calculate own- and cross-price elasticities

[25] Results are similar if the quadratic and cubic terms in p_{jt} are dropped and alternative instruments such as $(p_{jt} - \bar{p}_t)^2$ and $(p_{jt} - \max_j p_{jt})^2$ are used. These instruments, under the maintained assumption that prices are exogenous, are similar in spirit to the "differentiation IVs" offered by Gandhi and Houde (2019). I also present estimates of α_z which instrument for p_{jt} using the average price of good j in nearby stores in Appendix Figure A.9.

[26] Another reasonable approach would be to apply an empirical Bayes procedure to the distribution of estimates of $\bar{\alpha}_z$ and $\sigma^2_{\alpha,z}$. Empirical Bayes methods are motivated by the fact that extreme estimates of parameters

using the standard mixed logit formulas, and the relevant integrals are approximated using 200 Halton draws.

6.3 Constant Elasticity

One reasonable critique of mixed logit models is that the structural interpretation of their estimates relies on correctly specifying the utility function and the size of each market. These are strong assumptions, meaning it may be fruitful to consider estimates of substitution patterns which make different assumptions, as a sort of robustness check. To this end, I also consider an approach similar to that taken by DellaVigna and Gentzkow (2019), a recent paper highlighting the extent of uniform pricing in retail settings which also uses the Nielsen data. Instead of estimating a random utility model, we can target the store-module demand function via the following regression for each store and module

$$log(s_{jsw}) = \eta_s log(p_{jsw}) + \bar{\xi}_{js} + \Delta\xi_{jsw} \tag{11}$$

where η_s is the store- and module-specific price elasticity, $\bar{\xi}_{js}$ are UPC-store fixed effects.[27] As in my other estimates, these regressions are estimated separately by year and by product module. Though targeting this demand function directly has some aforementioned advantages, note that this specification does not directly address substitution between products. Thus, we cannot calculate any estimates of cross-price elasticities from this model.[28] Still, estimates of η_s can inform us of the extent to which price increases lead to declines in quantity demanded while imposing a different structure than that required by logit models. As shown in Section 4, substitution to the outside option has not increased in general over the sample period, meaning that changes in η will largely be due to changing cross-product substitution patterns. To the extent that estimates of η_s support the conclusions of the other estimated models, this

are likely the result of estimation error. One can regularize these estimates, shifting them closer to the mean. In early work I found this approach to be much less successful at eliminating nonsensical results than imposing constraints.

[27] I also estimate a version of this model which controls for the average of competing prices in each store-week-module combination. That is, I estimate each η_s in a regression of the form $log(s_{jsw}) = \eta_s log(p_{jsw}) + \beta_s log(\bar{p}_{-jt}) + \bar{\xi}_{js} + \Delta\xi_{jsw}$. Appendix Figure A.11 presents the results of this exercise, which are very similar to my results corresponding to the model in this section.

[28] If this model is interpreted in the Berry, Gandhi and Haile (2013) framework, then the inverse demand function includes only a product's own price and market share, meaning cross-price elasticities are identically zero. This interpretation also indicates that structural interpretation of this demand model requires a strong form of symmetry, as a more general model (even without cross-substitution) would imply that η should differ for every product.

strengthens the findings herein significantly.[29]

Estimating this model yields more than 100,000 estimates of η_s (across all stores and modules). Although I drop stores with fewer than 200 total observations in a module-year it is likely, given the relatively small sample used to estimate each η_s, that the most extreme estimates are due in large part to estimation error. Comparisons involving these noise-ridden estimates are likely to overstate changes over time, as some estimates in 2006 are less precise than their 2017 counterpart (largely due to the growth in the number of products, i.e. observations, per store). To alleviate some of this concern, I employ an empirical Bayes shrinkage procedure, following Morris (1983) and as implemented in Chandra et al. (2016)[30]. This procedure "shrinks" estimates which are imprecisely estimated toward the module-year mean of estimates, and is similar to many recent empirical papers which estimate a large number of parameters of interest (e.g. fixed effects) (Abaluck et al., 2020; Hull, 2018). All proceeding discussions and figures involving estimates of η_s refer to estimates after applying this shrinkage procedure unless otherwise specified.

6.4 Price Endogeneity

In estimating all models of demand in this paper I treat prices as exogenous with respect to changes in market- and product-level unobservable (to the econometrician) determinants of demand. As much of the IO literature has focused on developing good instruments for price, some time should be spent to explain why prices in the retail sector seem unlikely to be affected by the traditional critique, that prices are set by firms in response to (or anticipation of) changes in consumer demand which are observed by the firm and not by the econometrician (i.e. $\Delta\xi_{jsw}$).

To justify the use of price as an instrument, I operate under two assumptions: (i) prices are set before the demand shock $\Delta\xi_{jsw}$ is realized, and (ii) firms are unable to predict these shocks in advance. With regard to (i), evidence suggests that menu costs are non-trivial (Stamatopoulos, Bassamboo and Moreno, 2020), meaning stores are unlikely to be able to respond quickly to small changes in demand. On (ii), recent work by Hitsch, Hortacsu and Lin (2019) (using the Nielsen data as well) indicate that even store-level differences in demand within a chain may be difficult to precisely estimate given the data that firms are likely to have. Given this finding, weekly variation in product-specific demand seems significantly less likely to be well estimated

[29] To estimate this model, I absorb product-store fixed effects from prices and market shares using the *reghdfe* command in Stata (Correia, 2016) and then run the user-written *regressby* (by Michael Droste) to estimate Equation 11 at every store for each module. I calculate heteroskedasticity-robust standard errors which I adjust manually to account for absorbed degrees of freedom.

[30] I implement this approach via the *ebayes* Stata command written by Adam Sacarny.

by managers. Note also that I include product-store fixed effects in all estimates and either store-week or ZIP code-week fixed effects in my structural estimates, and that price only needs to be uncorrelated with the residual $\Delta\xi_{jsw}$ after absorbing these fixed effects. I am therefore only assuming that prices are exogenous with respect to changes in demand within a product at a given store. This means that firms can price according to store-specific demand for each product, as well as (in my mixed logit estimates) store-level changes in demand across weeks (e.g. seasonality in demand for ice cream), but assumes that firms do not price according to changes in demand for a product in a given week. This argument is similar to that made in Hitsch, Hortacsu and Lin (2019), who also estimate a model of demand treating prices as exogenous conditional on a similar set of fixed effects.

There are a number of other pricing behaviors which could also make prices exogenous conditional on the included fixed effects, even if the timing assumption I make is invalid. Existing evidence indicates that many retail chains appear to price nearly uniformly across stores within a week. Whether this is indicative of uniform or zone pricing on the part of retail chains (DellaVigna and Gentzkow, 2019; Adams and Williams, 2019) or uniform *markups* over geographically homogeneous wholesale costs (Butters, Sacks and Seo, 2020), this significantly limits the extent to which prices can be priced in response to market demand shifters. In a different vein, Conlon and Rao (2019) present evidence that retailers selling alcohol in Connecticut tend to set prices at a small number of discrete prices (e.g. those ending in ".99"), and that the majority of prices that change in response to tax increases move by whole dollar increments. Anecdotally, this is true more broadly, and the extent of endogenous price responses is naturally limited to the extent that firms maintain this discrete pricing behavior in the categories I study.[31]

There is also evidence that retail stores do not respond by changing prices in responses to changes in local competition. For example, although Arcidiacono et al. (2019) find a significant effect of superstore entry on incumbents' revenue, they find no such effect on price. That is, even when the entry of a new competitor decreases revenue at a given store by more than 10%, that store tends not to decrease its prices in any meaningful way. They conclude that the most likely explanation for this pair of findings is that stores follow pricing strategies which are independent of the demand they face. In particular, they argue (drawing on evidence from

[31]For example, many changes in demand which the firm might respond to are likely small enough that the optimal price changes by less than a dollar. If firms are constrained (e.g. by norms) to end prices in 0.99, then the firm may not respond at all to such a demand shock. More broadly, the firm can only fully respond to the very specific shocks to consumer demand which imply that the optimal price is one of the acceptable discrete values.

Eichenbaum, Jaimovich and Rebelo (2011)) that so-called "cost-plus" pricing is most likely. The key remaining source of potential endogeneity comes from inventory behavior on the part of consumers, which I assume has been addressed by the procedure in Section 3.3.

7 Results

In this section I offer two broad sets of results. First, I show that own- and cross-price elasticities have declined over time across a broad range of products and in each of the models of demand I have estimated. Second, I calculate markups implied by monopoly pricing at the module-store-week level. These hypothetical markups are not meant to be estimates of true retailer markups. Rather, this calculation is a helpful way to learn the general magnitude of the effect these changes in demand may have had on retailers' variable profits to date. Given the substantial recent literature on retail pricing, any assumed pricing strategy seems likely to be misspecified. These markups simply offer a way to measure whether the substitution patterns shown here are sufficient to imply increases in markups under a simple model.

7.1 Demand Estimates

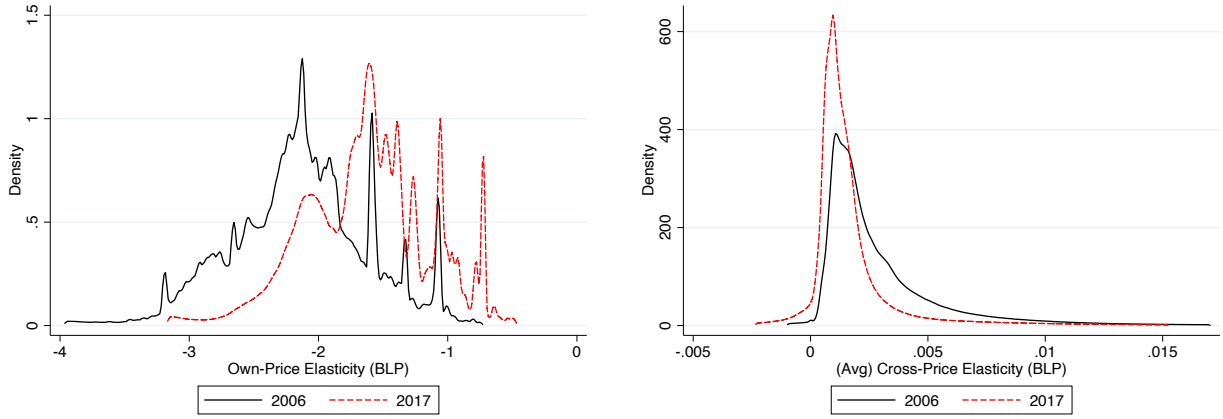
7.1.1 BLP

In Figure 3(a) I plot my first set of results, which are kernel densities of the distribution of price elasticities across all stores and products (in all nine modules) in my sample for 2006 and 2017. These distributions, which are composed of more than 150 million product-store-week observations, indicate a substantial shift to toward zero over this decade. Although much of the distribution in 2006 is to the left of -2, by 2017 we can see that the distribution has shifted significantly to the right. In columns 1 and 3 of Table 3, I present estimates of various quantiles of each of these distributions, ranging from the 5^{th} percentile to the 95^{th}. At each quantile, price elasticities are noticeably smaller in 2017 (in absolute value) than in 2006. On average, price elasticities declined from -2.16 to -1.62, a change of 25%, and very similarly at the median.[32]

Cross-price elasticities show a similar shift toward zero. In Figure 3(b), I calculate the average cross-price elasticity for each product in each market, and plot the distribution of these averages in 2006 and 2017. Dealing with averages smooths the distributions significantly,

[32]One reasonable concern is that many estimated own-price elasticities, especially in 2017, are less than 1 in absolute value (this is true in all results which follow as well). In standard models, this implies that firms are not short-run profit maximizing. To this point, one should note that Hitsch, Hortacsu and Lin (2019) estimate tens of millions of price elasticities across thousands of products using the Nielsen data and also find that many price elasticities are less than 1.

Figure 3: Distributions of Estimated Own- and Cross-Price Elasticities (BLP)



Note: BLP estimates of own-price elasticities across all 9 modules in the first 16 weeks of 2006 and 2017. Top and bottom 1% of each distribution have been trimmed.

meaning some detail is lost, but it necessary due to the disk space which would be required to store the full set of cross-price elasticities for all observations. Consistent with the notion that differentiation has increased in some way, cross-price elasticities have shifted significantly toward zero in 2017 relative to 2006, meaning changes in prices of goods are much less likely to induce substitution between products in 2017 than even a decade before. Figures 3(a) and (b) together demonstrate the headline result of this paper: differentiation appears to have substantially increased across many stores and products in the United States.

These results, while stark, aggregate over much of the variation in the data. One appeal of the level of detail provided by the Nielsen scanner data is that we can study whether substitution patterns have changed similarly across different groups of products. The median own-price elasticities in each module-year can be found in columns 1 and 3 of Table 4. In all but one of the nine modules I study (Ice Cream), I find that own-price elasticities have shifted toward zero substantially over the sample period.[33] Of the modules which see an increase in differentiation, the median decline ranges from approximately 11% (Frozen Pizza) to more than 30% (Fruit Drinks and Yogurt). Recall that estimates for each module in each year are the result of structural BLP estimates of demand which are estimated entirely separately from each other (18 estimated demand models in total). Thus, the changes we see in these substitution patterns in many modules are identified solely by within-module variation over time. Thus, the agreement

[33]Notably, among all 18 estimated BLP models, my only numerical issue arises in my 2006 estimates for Ice Cream, where I find that $\hat{\sigma}_\alpha$ touches the lower bound I enforce in estimation.

of all but one module is strong initial evidence that price sensitivity has declined broadly.[34] I now move on to the FRAC model, which relaxes the assumption that the distribution of preferences is homogeneous in all stores.

Table 3: Own-Price Elasticities and Implied Markups

	2006		2017	
BLP	Own-Price Elasticity	Markup	Own-Price Elasticity	Markup
5%	-3.13	-1.19	-2.47	-10.00
25%	-2.60	1.72	-2.00	1.88
Median	-2.17	1.98	-1.65	2.36
75%	-1.75	2.40	-1.27	3.20
95%	-1.06	4.62	-0.72	8.31
FRAC				
5%	-5.94	-1.85	-3.65	-8.42
25%	-3.09	1.64	-2.30	1.72
Median	-2.34	1.92	-1.77	2.32
75%	-1.74	2.47	-1.29	3.45
95%	-0.93	5.94	-0.73	11.33

Note: Shown are the mean, median, and other quantiles of the own-price elasticity and implied markup distribution for all BLP and FRAC estimates, pooling all nine modules in my sample.

7.1.2 FRAC

The preceding estimates, while able to approximate flexible substitution patterns, do make the strong assumption that both the mean and variance of preferences for price are identical across locations. Further, because price elasticities represent a combination of revealed preferences and equilibrium prices, it is uncertain whether changing preferences (i.e. changing differentiation) plays a critical role in these estimates. Now I move to my second set of demand estimates, i.e. those from the FRAC approach, which will address both of these concerns to some degree. In many ways, the results of this section are very similar to those in the preceding section.

The central way in which these FRAC results differ from BLP estimates is that I have estimated preferences independently in many geographies. In Figure 4 I present the distribution of my estimates of the mean of random coefficients on price in the utility function ($\bar{\alpha}$) across

[34] Though omitted here, for each module, I also calculate the average own-price elasticity estimate within each UPC in each year, and plot the distributions of these means across products in both years. The full set of module-specific plots of this form are shown in Figure A.6.

Table 4: Median Price Elasticities and Implied Markups, by Year and Module

BLP	Own-Price 2006	Markups 2006	Own-Price 2017	Markups 2017
Fruit Drinks	-2.47	1.79	-1.66	2.26
Soup	-2.09	2.09	-1.50	3.14
Cookies	-2.01	2.16	-1.52	3.06
Pizza	-2.59	1.80	-2.31	1.89
Ice Cream	-1.88	2.22	-2.01	2.20
Entrees	-1.91	2.19	-1.60	2.46
Yogurt	-2.16	2.15	-1.40	2.48
Remaining Fruit	-2.38	1.75	-1.87	2.47
Light Beer	-3.36	1.49	-2.93	1.62
FRAC				
Fruit Drinks	-2.26	1.82	-1.55	2.20
Soup	-4.00	1.97	-2.41	2.90
Cookies	-2.08	2.01	-1.59	2.55
Pizza	-2.31	1.88	-2.22	1.90
Ice Cream	-1.93	2.01	-1.87	2.28
Entrees	-1.92	2.09	-1.55	2.42
Yogurt	-2.28	1.89	-1.44	2.13
Remaining Fruit	-2.40	1.73	-2.04	2.18
Light Beer	-3.13	1.49	-2.23	1.67

Note: Shown are the medians of the full distribution of own-price elasticities and implied markups for 2006 and 2017, calculated separately for each module studied. Estimates shown separately for BLP and FRAC estimates.

three-digit ZIP codes in each year. These figures show us that, in each product module, estimated average price coefficients are significantly smaller in 2017 than in 2006. Note that these are estimates of utility parameters, and thus provide a new result relative to the preceding BLP estimates. Because the densities in 2017 (blue, dashed) are often significantly to the right of the densities in 2006 (black), this implies a shift in the effective disutility of price over the span of the sample. As in my BLP estimates, each density in this figure is the result of an entirely separate estimation procedure, making their relative agreement quite stark. This is even more true here, where each density is the result of hundreds of (approximate) mixed logit demand estimates.

As described in Section 6.1, these estimates of $\bar{\alpha}$ confound two potential changes. In estimating these each of these FRAC models, the variance of ϵ is normalized to $\frac{\pi^2}{6}$. This means that *estimated* changes in $\bar{\alpha}$ may be the result of changes in σ (which scales the variance of ϵ)

or changes in $\bar{\alpha}$ itself. The latter could be the driving force if consumer wealth has increased substantially relative to prices in these categories, or if the demographic composition of consumers at the stores in my sample has changed significantly. The former, to the contrary, would indicate that products are, on average, more niche in 2017 than in 2006 (because preferences in 2017 have larger variance). For now, these are indistinguishable, but in Section 8 I demonstrate that preference heterogeneity specifically plays a significant role.

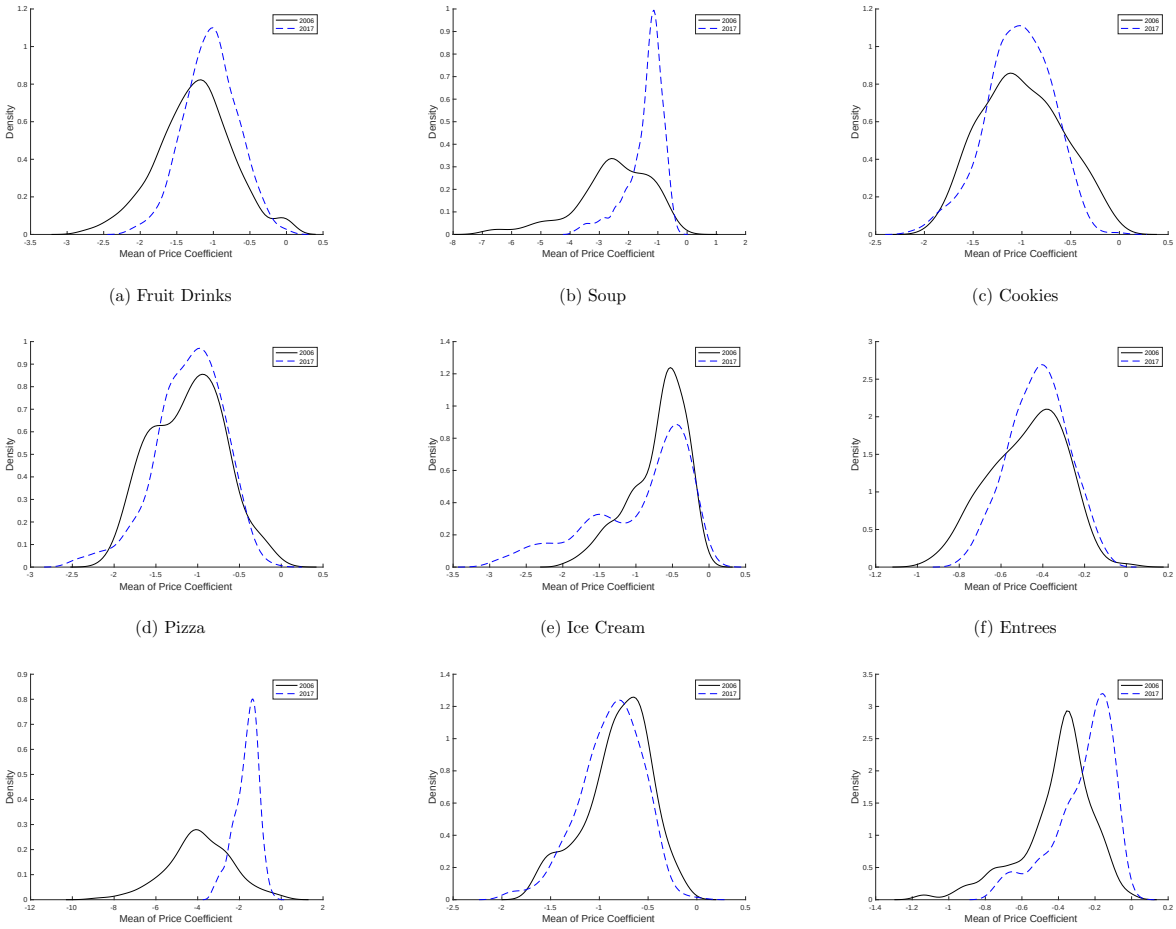
Now turning to estimates of differentiation, the story is quite similar to that in the preceding section. In Figure 5 I plot the distribution of own-price elasticities across all nine modules after trimming the top and bottom 1% in each year. Again, I find that this distribution has shifted significantly to the right in 2017 relative to 2006, and estimated own-price elasticities are similar in magnitude to my BLP estimates in both years. At the median, these elasticities have declined from -2.34 to -1.77, and a similar decline is also found at the mean and at a number of percentiles of the distribution, as described in columns 1 and 3 of Table 3. I also present median own-price elasticities for each module in each year in Table 4, which demonstrates that own-price elasticities have declined in all modules. The distribution of cross-price elasticities (Figure 5(b)) also exhibits a noticeable shift toward zero.[35,36]

Although I have estimated the FRAC model to permit geographic preferences heterogeneity, up to this point it is unclear whether the heterogeneity I find in my estimates represents differing preferences or estimation error. To be specific, one may worry that the variance, and in particular the long left tail, of own-price elasticities and estimated utility parameters here largely represents estimation error around BLP estimates from the preceding section, due to the relatively small sample size contributing to each FRAC estimate. If this were driving all of the estimated differences, FRAC estimates of $\bar{\alpha}$ would be uncorrelated across modules. I show in Table A.2 that this is not the case. To the contrary, estimates are generally significantly correlated across modules in both years. This seems to imply that preferences for price, and/or the nicheness of products, do differ significantly across geographies. Whereas this geographic heterogeneity can be quite difficult to incorporate in the standard BLP estimation procedure, it is easily incorporated in FRAC.

[35] I also calculate UPC-level average own-price elasticities for each module-year, which I omit here to save space but display in Figure A.8.

[36] One remaining concern with the estimates herein may be that either 2006 or 2017 are anomalous years in the data. I show in Appendix C that, in the Yogurt category, these patterns over time follow a clear trend, making this story unlikely.

Figure 4: Distribution of Estimates of ZIP Code-Level Price Coefficients



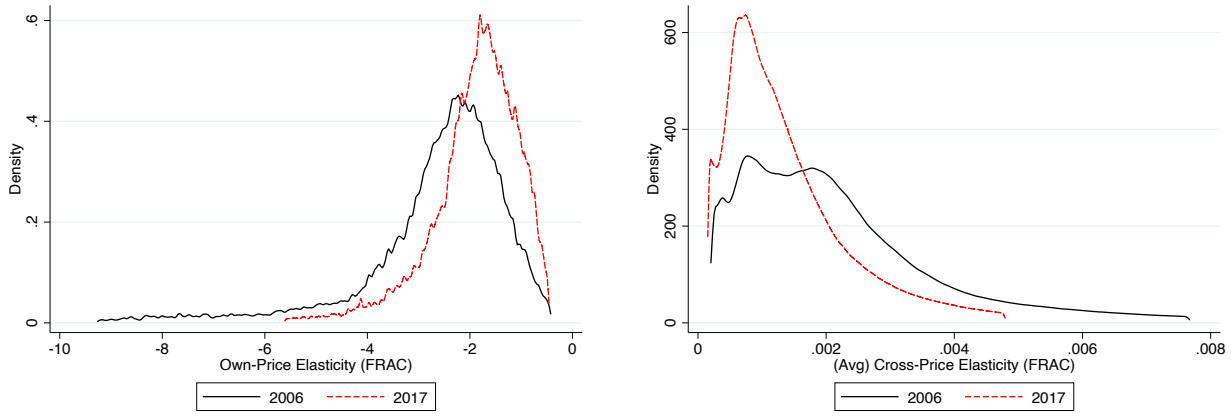
(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Distributions of estimated mean utility coefficient on price, in 2006 (black, solid) and 2017 (blue, dashed). Estimates come from solving the constrained least squares problem in Section 6.2 separately for each three-digit ZIP code, module, and year.

Figure 5: Distributions of Estimated Own- and Cross-Price Elasticities (FRAC)



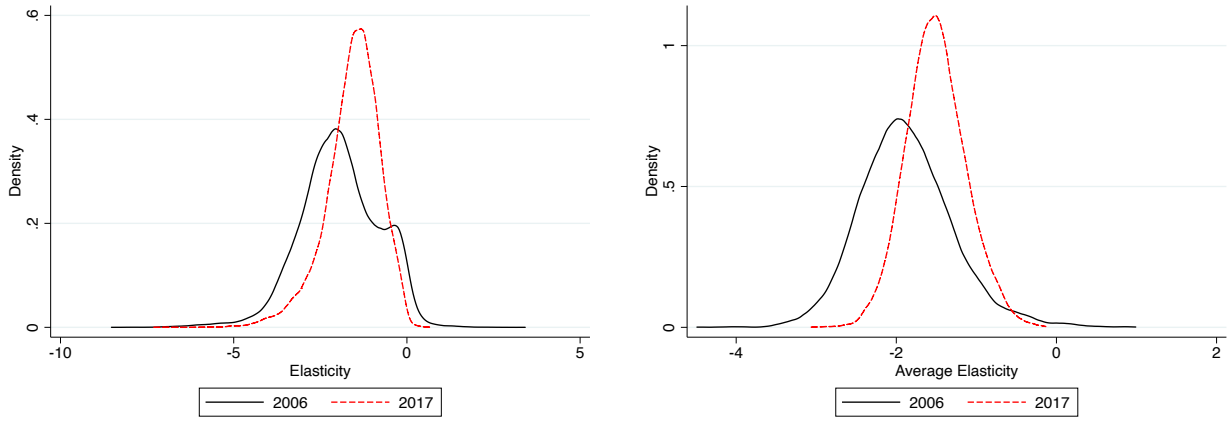
Note: FRAC estimates of own- and (average) cross-price elasticities across all 9 modules in the first 16 weeks of 2006 and 2017. Top and bottom 1% of own-price elasticities have been trimmed, as have the top and bottom 5% of average cross-price elasticities in each year.

7.1.3 Constant Elasticity

I now present my estimates of the constant elasticity demand model, for which I estimate more than 100,000 store-level price elasticities. As mentioned in the estimation section, I apply an empirical Bayes shrinkage procedure to these estimates before plotting them here. The distributions of these (post-shrinkage) estimates are shown in Figures 6(a) and (b).[37] In Figure 6(a) I plot the full distribution of module-store-level estimates of price elasticities in 2006 and 2017. In order to demonstrate that no single product module is driving these results, in Figure 6(b) I calculate store-level averages of estimated elasticities and plot the distribution of these averages across stores. A figure replacing store-level means with medians looks very similar, and my results are qualitatively unchanged when I also control for the average of competing prices in each regression (results shown in Figure A.11). As in the two structural models, these estimates imply that price elasticities have shifted substantially to the right over the sample period. Furthermore, the magnitudes of price elasticities in the two periods are quite similar to estimates of price elasticities from FRAC (Figure 5). This is encouraging, as these two models make very different assumptions about consumer choice (as discussed in the preceding section).

[37] Table A.3 presents the mean and selected quantiles of the distribution of pre- and post-shrinkage estimates. The two distributions are generally very similar.

Figure 6: Distributions of Price Elasticities (Constant Elasticity)



(a) (b)

Note: Figure (a) presents the distribution of estimates of η_s, the store-module-level price elasticity. Figure (b) presents the distribution of store-level averages of these estimates (i.e. averaging across the modules offered by a store).

7.2 Implied Markups

In this section I use the preceding demand estimates to calculate markups which would be implied under the assumption of monopoly pricing at the store-module level. Although in general I am hesitant to impose the full structure of a supply side model, assuming a model of supply side behavior can give us some idea of the magnitude of changes in markups over my sample period. For simplicity, I treat firms as pricing monopolistically at the store-module-week level, taking into account all cross-effects between products in a module. This is not meant to represent structural estimates of true markups at the store level. Rather, I think of this exercise as providing a sense of the potential magnitude of increasing profits due to the changes in demand which I have documented in this section. Given that we have already seen that prices have changed little on average in most modules, these estimates could reasonably be interpreted as one measure of the magnitude of cost reductions over time.

7.2.1 BLP

Under monopolistic pricing, optimal margins satisfy the following equality

$$p^* - mc = -J_p^{-1}(p^*)s(p^*) \tag{12}$$

where p^*, mc and s are vectors of product-store-week-level prices, marginal costs, and market shares, and J_p^{-1} is the inverse of the Jacobian of store-level market shares with respect to prices. Because prices are known and the demand function has been estimated, this equation implies estimates of marginal costs for each product in every store. Implied markups ($\frac{p}{mc}$) can then be calculated from those marginal costs. I use my estimates from BLP-style models to construct these estimated markups separately for each product module in each year, and present them first in Figure 7. In this figure I plot the distribution of implied markups from the 25^{th} to the 75^{th} percentile in 2006 (black, solid) and 2017 (red, dashed). I focus on this small range of the distribution because many points outside of this range are either negative or so large in absolute value as to be uninformative in either direction. This is one reason why these estimates should not be taken as estimates of true retailer markups. Still, within the interquartile range Figure 7 indicates significant growth in implied markups over the decade. At the median, implied markups have increased from 1.95 to 2.36 (see columns 2 and 4 of Table 3), an increase of 20%. At the 95^{th} percentile implied markups have nearly doubled from 4.62 to 8.31.[38]

My demand estimates (which underly implied markups) are quite consistent with other estimates of price elasticities in the Nielsen data (Miller and Weinberg, 2017).[39] It is also worth comparing these results to estimates in De Loecker, Eeckhout and Unger (2020), who present estimates of markups from Compustat, which covers publicly traded retail firms, and from the Census of Retail which contains the universe of retail firms. The retail industry they study is much broader than the retail food industry herein, and their methods and data differ substantially, as they do not observed product-level prices or sales. Although their Compustat estimates indicate very small increases in markups in this industry over the past 40 years, their estimates from the Census indicate a sharp increase in markups at the mean from slightly above 2 to slightly above 3 between 2002 and 2012 alone (Figure 12(c) therein). Further, these estimates imply that at the 95^{th} percentile markups have rise from between 3 and 4 to nearly 7. Clearly, my estimates tend to agree more with the Census estimates. The most notable difference, which is a key point of this paper, is that all of the changes in implied markups I estimate are the result of *within-firm* (and in fact within-establishment) changes in variety and consumer preferences, and do not require any changes in competitive behavior or market power (in the traditional sense). Thus, although my estimates match this set of estimates by

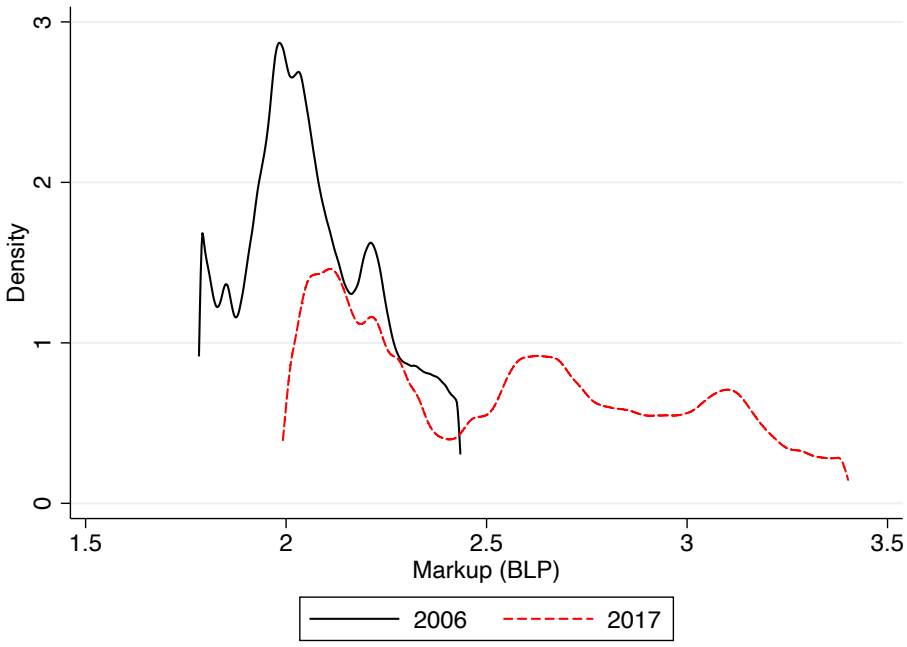
[38] In Appendix Figure A.7, I plot the distribution of markups implied by my BLP demand estimates separately by module.

[39] The apparent differences in headline numbers between, for example, my results in the Light Beer category and estimates in (Miller and Weinberg, 2017) appear to come from the difference in the number and characteristics of goods in our samples. When I restrict my Light Beer sample to one more similar to theirs, I find that my estimates of price elasticities are very similar.

De Loecker, Eeckhout and Unger (2020) reasonably well, they provide an alternative explanation with very different policy implications.

As was the case with the demand estimates in the preceding section, these results are not being driven by a small subset of modules. To demonstrate this, for each module I divide implied markups into 100 equally sized bins in each year, take the average of all markups within each bin, and plot these binned averages in 2017 against the same bin in 2006.[40] Again I focus only on the interquartile range of each distribution. These plots are essentially quantile-quantile plots, meaning that if the distribution of implied markups were unchanged over this decade, I would find that all points lie along the plotted 45-degree line. If parts of the markup distribution have shifted to the right over time, these points should be above the 45-degree line. Figure 8 tells a clear story. In all but the Ice Cream category, estimated markups are substantially larger in 2017 than in 2006 along nearly the entire interquartile range.

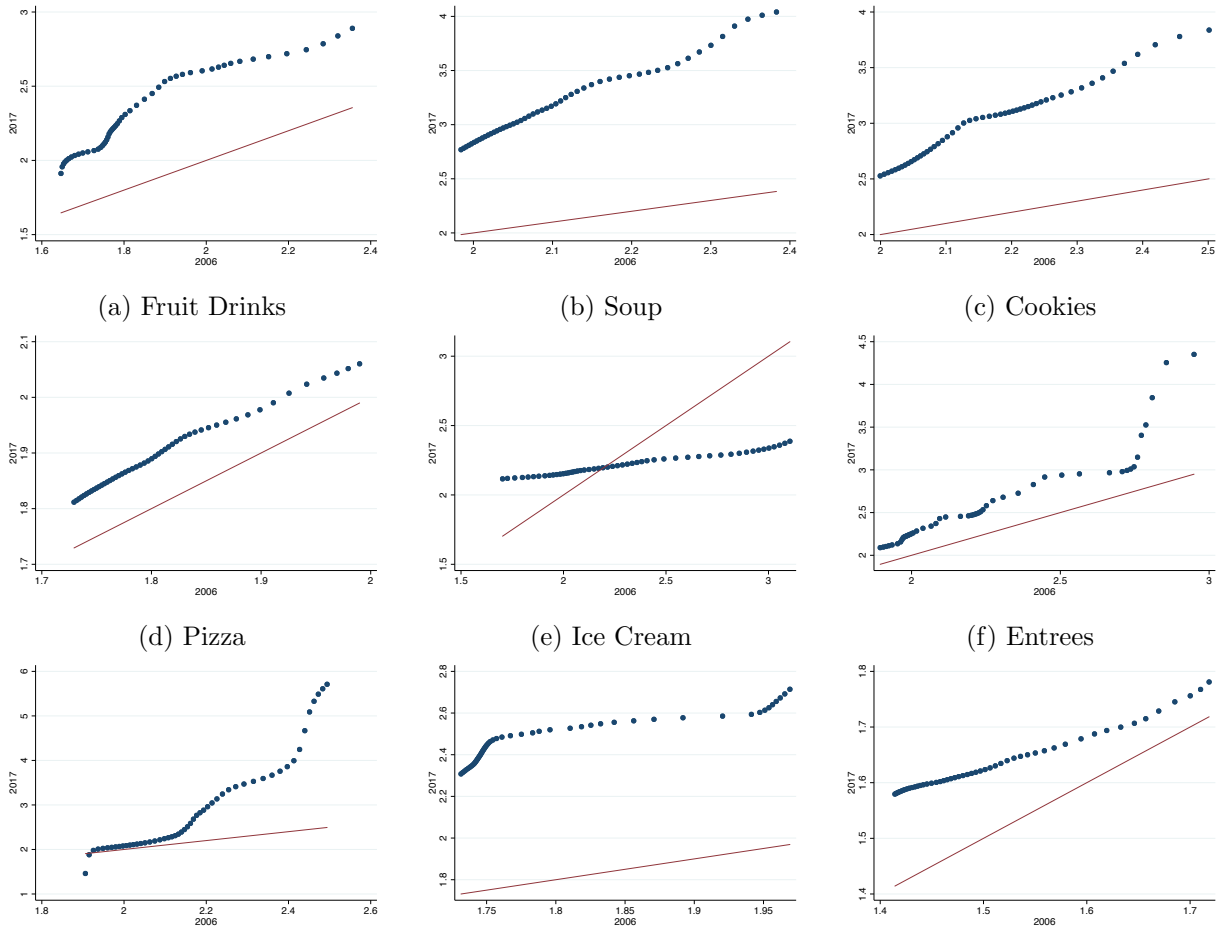
Figure 7: Interquartile Distribution of All Implied Markups (BLP)



Note: BLP estimates of implied markups across all nine modules in the first 16 weeks of 2006 and 2017. Markups were calculated under the assumption that firms price monopolistically within each store-module-week, taking into account cross-substitution between all goods within a module. Only the interquartile range of implied markups is shown here.

[40] This is similar to the results of the "qqplot" command in Stata, which plots quantiles of two distributions against each other. I focus on the mean in these equally sized bins because quantiles are often much more sensitive to small variations in the data.

Figure 8: Quantile-Quantile Comparisons of Markups in 2006 and 2017 (BLP), by Module



(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: I calculate 100 equally sized bins of implied markups in each year for each module, and calculate the mean within each bin. This figure plots the 2017 mean in a given bin (quantile) against the 2006 mean from the same bin, along the interquartile range for each module. The red line denotes the 45 degree line.

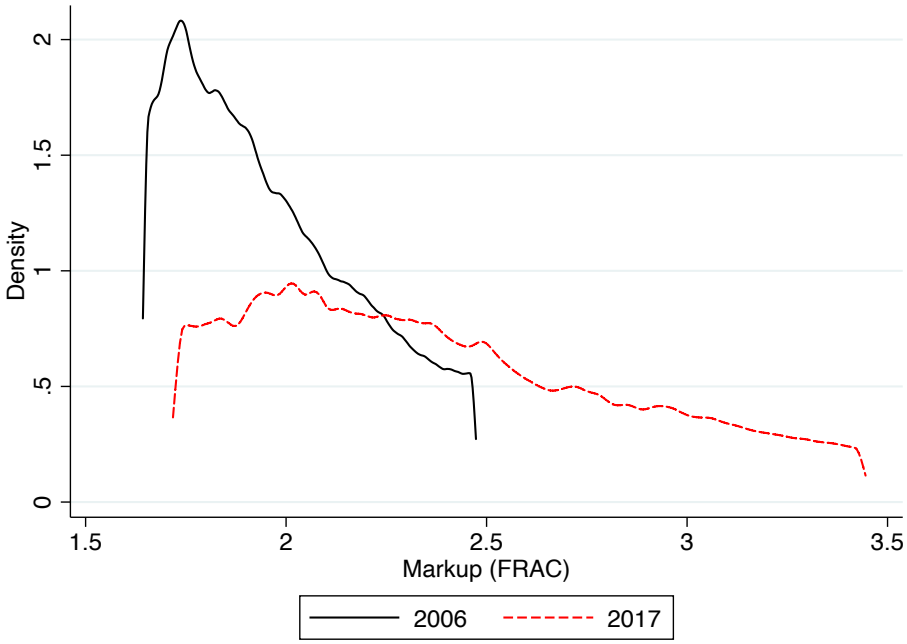
7.2.2 FRAC

Next, following the discussion in the previous section, we can convert FRAC demand estimates into estimates of markups under the same assumption, that each store prices monopolistically each week with in a module. Again, these markup estimates give us a ballpark figure to understand how a perfectly optimizing monopolist might have benefitted from the changes in demand documented in preceding sections. In Figure 9 I plot the interquartile range of monopolists' markups across all products, stores, and weeks in my sample according to FRAC demand estimates. Just as in Figure 7, the distribution of implied markups has shifted significantly to the right in 2017 relative to 2006. Consistent with my finding that the distribution

of price elasticities estimated by FRAC has a longer left tail than that of BLP estimates, the distribution of markups is somewhat wider in both years. As a result, some of the apparent markup growth here appears to come from the right tail of the distribution. Still, the general magnitude of implied markups are quite similar between the two Figures (7 and 9).

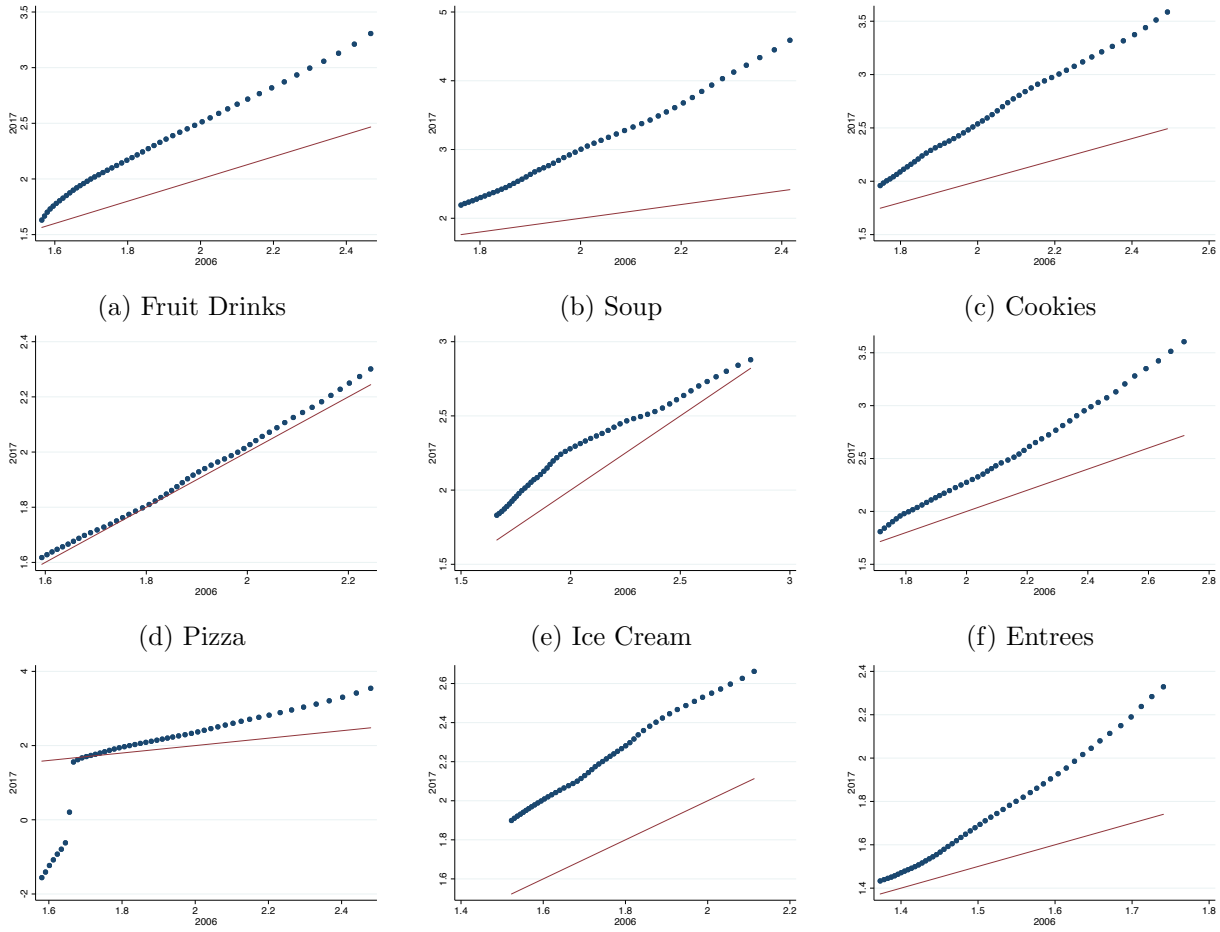
In Figure 10 I again plot quantiles of the implied markup distribution in 2006 and in 2017 as well as the 45-degree line. Just as is the case in my BLP results, I find that in the majority of categories the markup distribution has shifted noticeably to the right. The shape of these shifts differ by category, but the majority of modules demonstrate a significant shift particularly in the right tail of the distribution. Before moving on, we can note a few key differences. First, unlike in my BLP estimates, FRAC-implied markups in the Ice Cream category have increased over time, meaning that all modules see an increase at the median. However, at the bottom of the interquartile range, implied markups in Yogurt (g) are smaller in 2017 than in 2006, and in Pizza (d) the two distributions are quite similar across the entire range. Still, altogether, we see a general increase in implied markups across the vast majority of quantiles in most modules.

Figure 9: Interquartile Range of Markups (FRAC)



Note: FRAC estimates of implied markups across all nine modules in the first 16 weeks of 2006 and 2017. Markups were calculated under the assumption that firms price monopolistically within each store-module-week, taking into account cross-substitution between all goods within a module. Only the interquartile range of implied markups is shown here.

Figure 10: Quantile-Quantile Comparisons of Markups in 2006 and 2017 (FRAC), by Module



(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: I calculate 100 equally sized bins of markups implied by my FRAC estimates in each year in each module, and calculate the mean within each bin. This figure plots the 2017 mean in a given bin (quantile) against the 2006 mean from the same bin, along the interquartile range for each module. The red line denotes the 45 degree line.

7.2.3 Constant Elasticity

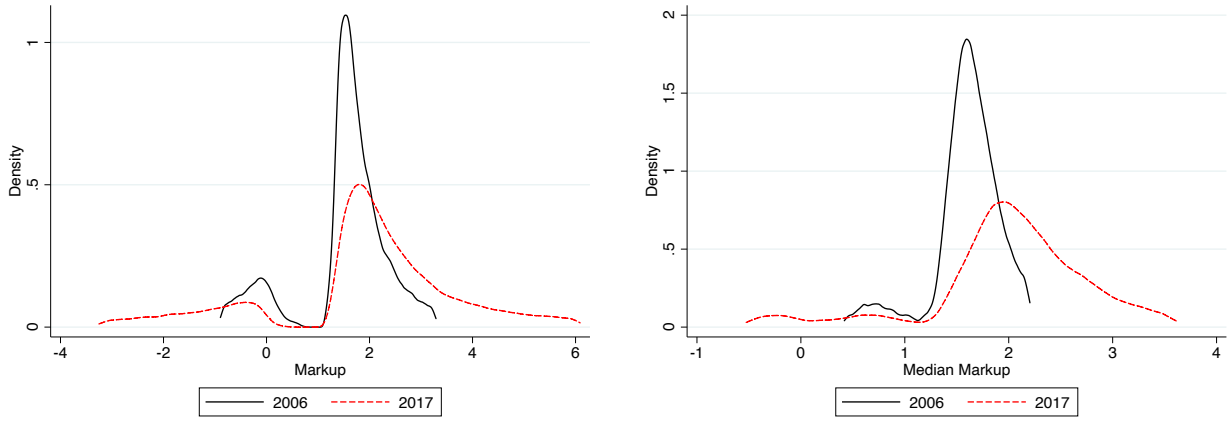
Under the constant elasticity specification of demand, optimal pricing follows a very simple rule. Assuming that firms price each good entirely independently, optimal markups in each store can be characterized by

$$\frac{p_{jsw}}{mc_{jsw}} = \frac{\eta_s}{1+\eta_s} \tag{13}$$

Thus, we can plug estimates of η_s (store-level price elasticities) into this equation to derive markups implied by this pricing assumption. In Figure 11(a) I present the distributions of these implied markups in each year, trimming the top and bottom 10% of each distribution. Note that, unlike in the preceding estimates, in this case optimal markups are calculated not for each product separately, but rather at the store-module level, because the estimated model assumes that price elasticities for different products in a module-store-year are identical. Further, these estimates are solely regarding *optimal* implied markups, and do not provide any information as to whether firms are actually charging higher markups. Preceding estimates, to the contrary, represented estimates of markups implied by observed prices under an optimal pricing rule.

The findings in Figure 11(a) may appear to be ambiguous, as many estimates in both years are negative. That is, the distribution in 2017 includes a longer right tail than in 2006, but also a longer left tail. This is to be expected, given the significant fraction of price elasticities which are less than 1 in Figure 6. To remove some of the impact of these categories, I also calculate the store-level median of markup estimates (i.e. across modules) in each year, and plot these distributions in Figure 11(b), again trimming the top and bottom 10% of estimates. This figure presents a story much more clearly consistent with the demand results, that optimal markups have increased across the vast majority of stores.

Figure 11: Distributions of All Optimal Markups (Constant Elasticity)



(a) (b)

Note: Figure (a) presents the distribution of all plug-in estimates of optimal markups from the constant elasticity demand model (Equation 13), and Figure (b) presents the distribution of store-level medians (across modules) of these estimates. The top and bottom 10% of each distribution are omitted.

8 Testing Model Predictions

Thus far, I have documented substantial changes in consumers' willingness to substitute between products. These changes, by definition, imply that perceived differentiation has increased. Still, in the simple model I presented I argued that the most natural way for differentiation to have changed over time was through the changes in the variance of preferences for horizontal attributes of goods, which is in part due to the introduction of more niche goods. This section aims to discuss the strength of the evidence for these claims in particular.

8.1 Nicheness vs. Price Sensitivity

As discussed in Section 6, in estimating the preceding models I normalize the variance of horizontal preferences (σ) because this variance is not separately identified from the disutility of price without additional structure. Thus, the results in Figure 4, which demonstrate the changes in the *estimated* disutility of price, confound the two. It is therefore possible that the results thus far represent a change in demographics, e.g. wealth, which have caused consumers' price sensitivities to decline absent any change in preference heterogeneity. It is also possible that newer products are more vertically differentiated (e.g. higher quality), which to this point I have generally left unmentioned. To explore the roles each of these mechanisms may have played, and to isolate the contribution of preference heterogeneity to the observed trends, consider the following logit model, where utility takes the form:

$$u_{ijsw} = \begin{cases} \alpha_{2006} p_{jsw} + \bar{\xi}_{js} + \Delta\xi_{jsw} + \epsilon_{ijsw}, \; Var(\epsilon_{ijsw}) = \sigma^2_{2006}\frac{\pi^2}{6} & \text{in 2006} \\ \alpha_{2017} p_{jsw} + \bar{\xi}_{js} + \Delta\xi_{jsw} + \epsilon_{ijsw}, \; Var(\epsilon_{ijsw}) = \sigma^2_{2017}\frac{\pi^2}{6} & \text{in 2017} \end{cases} \tag{14}$$

Note that I do not include store-week fixed effects here, as these would make the interpretation of product fixed effects $\bar{\xi}_{js}$ more complicated. Though I omit this notation, I allow α and σ to vary not only across years, but also *at every store*. Otherwise, my notation here is largely the same as in Section 5, except that I make explicit that both the coefficient on price in the utility function and the variance of ϵ have changed over time. Note that I assume, to the contrary, that the fixed effects $\bar{\xi}_{js}$ do not change over time for a given product. To the extent that $\bar{\xi}_{js}$ represents preferences for fixed, unobserved, product characteristics, this assumption is relatively weak. On the other hand, these fixed effects may indeed change over time if advertising has increased average preferences for some goods relative to their competitors.

To identify the relative importance of the disutility of price and horizontal differentiation

in determining changes in substitution patterns, I first apply the standard logit inversion in each year:

$$log\Big(s_{jsw}/s_{0sw}\Big) = \begin{cases} \frac{\alpha_{2006}}{\sigma_{2006}} p_{jsw} + \frac{\bar{\xi}_{js}}{\sigma_{2006}} + \Delta\xi_{jt} & \text{in 2006} \\ \frac{\alpha_{2017}}{\sigma_{2017}} p_{jsw} + \frac{\bar{\xi}_{js}}{\sigma_{2017}} + \Delta\xi_{jt} & \text{in 2017} \end{cases} \tag{15}$$

where I emphasize that the effective/estimated fixed effects and price coefficients in each year are dependent on the variance of ϵ, just as in Equation 3.[41] What is important to see here is that although the difference between price coefficients over time represents a combination of changes in price sensitivities (α) and differentiation (σ), changes in the magnitudes of estimated fixed effects offer an estimate of the latter alone. Thus, by comparing the magnitudes of estimates of product fixed effects in 2006 and 2017, we can derive an estimate of $\frac{\sigma_{2017}}{\sigma_{2006}}$, which will allow us to decompose the changes in demand in the previous section to show to what extent price elasticities would have changed in the absence of any rising preference heterogeneity. There are many choices for deriving an estimate of $\frac{\sigma_{2017}}{\sigma_{2006}}$. Perhaps the most direct route is to calculate an estimate of σ_{2006} for each product which is sold in a store in both 2006 and 2017. However, this is sensitive to estimation noise in individual fixed effects, which can be quite large. Instead, I take the standard deviation of product fixed effects (among products sold in a store in both years) as an estimate of $\frac{1}{\sigma_y}$ for each year y, and use the ratio of these quantities as a measure of the ratio of scale parameters of ϵ.[42]

From this, we can also derive the extent to which α (the true disutility of price) changed as well, as

$$\overbrace{\frac{\hat{\alpha}_{2006}}{\hat{\alpha}_{2017}}}^{\text{Ratio of estimated price coefficients}} \times \overbrace{\frac{\widehat{\sigma_{2006}}}{\sigma_{2017}}}^{\text{Estimated change in pref. het}} = \frac{\widehat{\alpha_{2006}}}{\sigma_{2006}} \times \frac{\widehat{\sigma_{2017}}}{\alpha_{2017}} \times \frac{\widehat{\sigma_{2006}}}{\sigma_{2017}}$$

$$= \underbrace{\frac{\widehat{\alpha_{2006}}}{\alpha_{2017}}}_{\text{Estimated change in price disutility}}$$

With estimates of the changes in σ and in α separately, we can now study the importance of each by estimating price elasticities in 2017 under 2006 levels of horizontal differentiation and/or price disutility. What this algebra shows is that if we find that estimated fixed effects

[41] The residual $\Delta\xi_{jt}$ is also rescaled, but I omit that notation for brevity.

[42] I estimate this model using the *reghdfe* command in Stata (Correia, 2016), to estimate many fixed effects and many store-level coefficients on price at once.

and estimated price coefficients have changed by the same multiplicative factor, we should conclude that most changes in price elasticities are due to rising horizontal differentiation. To the extent that the two differ, the disutility of price may have also changed over time.

I conduct this exercise by constructing two counterfactual distributions of price elasticities in 2017. Note that estimated own-price elasticities in a logit model are equal to

$$\frac{\partial s_{jt}}{\partial p_{jt}} \frac{p_{jt}}{s_{jt}} = \hat{\alpha} p_{jt} \Big(1 - \frac{exp(\hat{\alpha} p_{jt} + \hat{\xi}_{jt})}{1 + \sum_j exp(\hat{\alpha} p_{jt} + \hat{\xi}_{jt})} \Big)$$

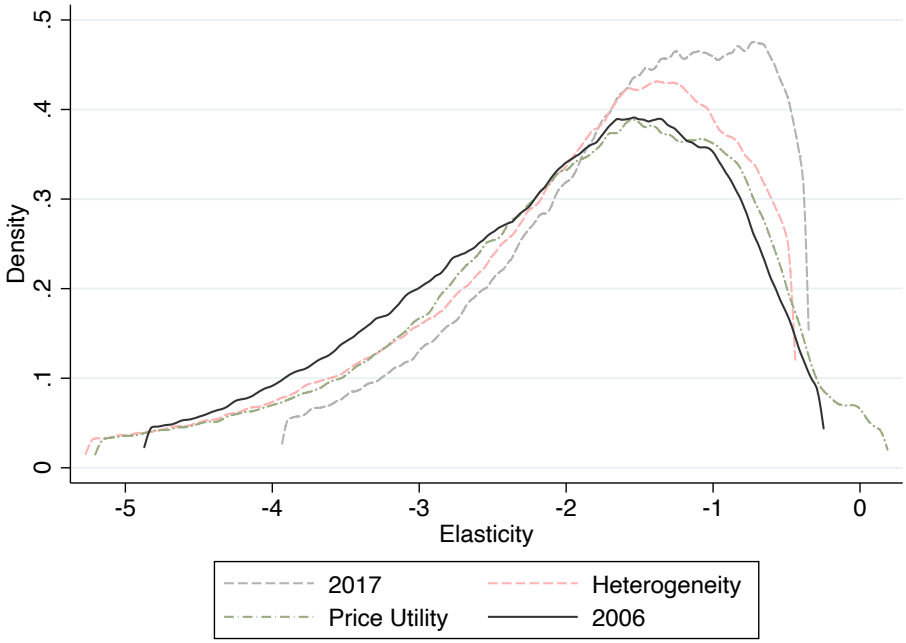
Thus, to calculate the desired counterfactuals I first rescale the mean utility term $\alpha_{2017} p_{jsw} + \xi_{jsw}$, as well as my estimates of α_{2017}, by $\frac{\sigma_{2017}}{\sigma_{2006}}$. This sets the effective variance of ϵ (i.e. preference heterogeneity) in each store to its 2006 levels. In a second counterfactual, I also rescale $\hat{\alpha}_{2017}$ by my estimate of $\frac{\alpha_{2006}}{\alpha_{2017}}$, thereby setting *both* the variance of ϵ and the disutility of price to 2006 levels.

In Figure 12 I plot the results of this exercise in the form of four distributions. In black (solid) is the distribution of estimated price elasticities in 2006. The estimated price elasticities in 2017 are in grey. In red (dashed) is the first counterfactual distribution, in which I recalculate all price elasticities 2017 after rescaling to set the variance of ϵ to 2006 levels for each store. The magnitude of the importance of horizontal differentiation is apparent from this plot. Removing the effects of preference heterogeneity over this decade shifts the distribution of price elasticities in 2017 much further to the left and much closer to the estimated distribution in 2006, meaning that explanations for the demand estimates herein that rely solely on the disutility of price or on changes in the price distribution leave much of the observed changes unexplained. In Figure A.12 of the Appendix I show that this result is true to varying degrees in most modules as well. Finally, in green (dot-dashed) I show the second counterfactual distribution, which calculates price elasticities in 2017 under 2006 levels of both preference heterogeneity and disutility of price. Parts of this distribution are significantly closer to my estimates in 2006 than even the first counterfactual, meaning that the disutility of price plays a role in declining price elasticities which is distinct from the rise of preference heterogeneity.

To make the similarity of the red (counterfactual preference heterogeneity) and black (2017 estimates) lines more concrete, I also present Table 5, in which I display the mean and selected quantiles of each of the three distributions. The proximity of the counterfactual 2017 distribution to the estimated distribution in 2006 is stark. Although they differ noticeably at most quantiles, the differences between these two distributions is in all cases smaller than the dif-

ference between the estimated distributions in 2006 and 2017. In this way, rising preference heterogeneity appears to have played a significant role in the trends in differentiation shown above.

Figure 12: Distributions of Estimated and Counterfactual Own-Price Elasticities



Note: In black and grey I show estimates of price elasticities from a logit model with store-specific preferences in 2006 and 2017 (respectively). In red (dotted) I plot the distribution of 2017 price elasticities under the counterfactual in which horizontal differentiation in each store were set to its 2006 level, and in green (dot-dashed) I plot this distribution after also adjusting the disutility of price to 2006 levels.

Table 5: Distribution of Estimated and Counterfactual Own-Price Elasticities

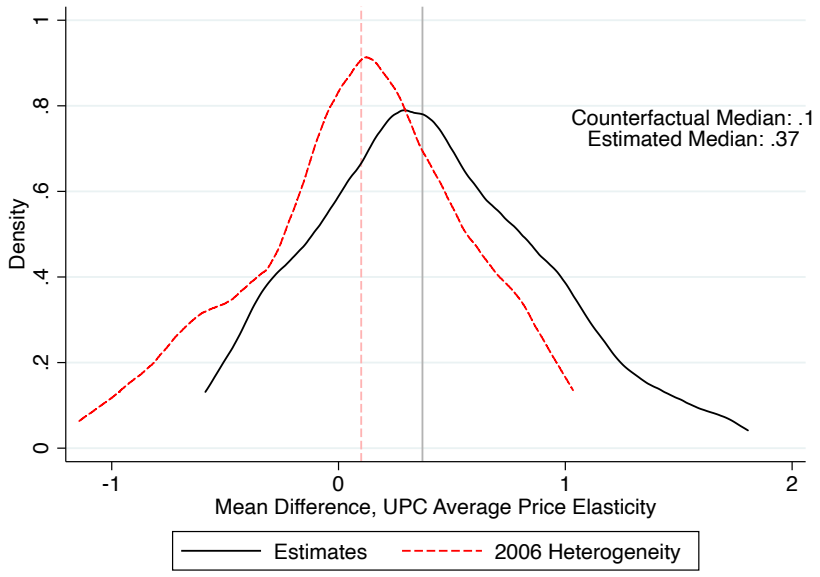
	2006 Estimates	2017 Counterfactuals	2017 Estimates
Mean	-2.13	-2.33	-1.7
10%	-4.00	-4.08	-3.21
25%	-2.89	-2.73	-2.25
Median	-1.89	-1.76	-1.47
75%	-1.16	-1.10	-.87
90%	-.63	-.64	-.50

Note: Mean and selected quantiles of estimated own-price elasticities in 2006 and 2017, and counterfactual 2017 estimates in which store-level preference heterogeneity is set to 2006 levels (dashed, red in Figure 12). Estimates derived by store-level logit models (Equation 15).

One concern with Figure 12 is that it could, in principle, mask significant heterogeneity

across products. That is, although the full distribution of counterfactual estimates is much closer to the 2006 distribution than are the 2017 estimates, it could be that average price elasticities at the UPC level have moved substantially within this distribution, meaning that Figure 12 overstates the role of preference heterogeneity in my demand estimates. To address this concern, I calculate average price elasticities for each product in each year. I calculate the difference in this average between years for each product which appears in both years, both for my estimates and under the desired counterfactual. The distributions of realized and counterfactual differences are plotted in Figure 13, where I show that while estimated elasticities increase by 0.38 at the median, UPC-average counterfactual elasticities increase by only 0.1. This makes clear that, in fact, returning differentiation to 2006 levels reduces the estimated average change in product-level price elasticities significantly.

Figure 13: Counterfactual Changes in Price Elasticities without Preference Heterogeneity



Note: Distributions of average changes (between 2006 and 2017) in own-price elasticities at the product level. Changes implied by store-level logit estimates are in black (solid), and changes implied by a counterfactual which sets preference heterogeneity to 2006 levels in each store are in red (dashed).

8.2 Are New Products More Niche?

The final prediction of the simple model in Section 2 is that newer products are likely to be more niche than older products. Although I take the preceding section as evidence that products are more niche in 2017 than in 2006 on average, it could be that *all* products have become more

horizontally differentiated over time due to, for example, effective advertising. This is difficult to test directly, in large part because it's unclear how best to define "new" given that only two years (a decade apart) are observed. Given the caveat that any such definition will be coarse, we can however permit consumers to have more heterogeneous preferences for some products than others. To do this, I estimate a model identical to that in Section 6.1 for each module in 2017 with two exceptions. First, I introduce an additional scaling parameter which rescales the utility (net of the logit error) for "new" products relative to existing products, such that

$$u_{ijsw} = \begin{cases} \alpha_i p_{jsw} + \xi_{jsw} + \epsilon_{ijsw} \text{ for old products} \\ \rho(\alpha_i p_{jsw} + \xi_{jsw}) + \epsilon_{ijsw} \text{ for new products} \end{cases} \tag{16}$$

where I define a product as "old" if it was ever in my sample in 2006 and "new" otherwise.[43] If stores only sold either all new goods or all old goods, then this form of scaling would be identical to letting the variance of ϵ differ across these two types of goods. Because stores sell a combination of both types of goods, these two models are not equivalent, but they are quite similar. For large values of ρ, the importance of the logit error in determining demand for a new product shrinks relative to the vertical characteristic and price. For small ρ, the logit error becomes increasingly important, meaning consumers are less likely to substitute away from that product in response to price increases. Thus, the hypothesis of the simple model is that $\rho < 1$. The second deviation from Section 6.1 is that I only include UPC and store fixed effects separately (rather than UPC-store fixed effects). I do this because I find in practice that including store-UPC fixed effects tends to make the search procedure behave much more poorly and make the results less robust to changes in sample size. I am forced by time and computational constraints to draw a random sample of 1,000 stores for each module and estimate the model using only these samples. Estimates with this reduced number of fixed effects appear to be robust to drawing alternative samples.

I present the results of this estimation procedure in Table 6. I find that in all modules except for Remaining Fruit the estimated value of ρ is statistically significantly less than 1, meaning that demand for newer products is more dependent on logit errors (i.e. preference heterogeneity) than is demand for older products. This is consistent with the model, which predicted that newer products would be more niche than older products and thus that consumers would be

[43] One may wonder why I introduce ρ instead of directly permitting the variance of ϵ to vary across products. This is because the standard logit form (which describes the maximum among *symmetric* logit draws) fails to hold when the scale of ϵ differs across products, making the estimation of such a model significantly more complicated.

less sensitive to price changes in the former. The extent to which these newer products are more niche varies substantially by category. In Light Beer, $\hat{\rho}$ is much smaller than 1, implying that consumers of new light beers are less than half as sensitive to price changes as consumers of older beers.[44]

Table 6: Estimates of Scaling Parameter for Newer Products

	ρ	$\bar{\alpha}$	σ_α
Fruit Drinks	0.92	-0.63	1.04
	(0.02)	(0.01)	(0.03)
Soup	0.92	-0.69	1.23
	(0.01)	(0)	(0.01)
Cookies	0.98	-0.5	1.24
	(0.01)	(0.01)	(0.03)
Pizza	0.93	-0.48	3.63
	(0.03)	(0.01)	(0.21)
Ice Cream	0.67	-0.53	2.02
	(0.02)	(0.01)	(0.05)
Entrees	0.82	-0.18	1.02
	(0.02)	(0.01)	(0.05)
Yogurt	0.79	0.08	4.12
	(0.02)	(0.02)	(0.12)
Remaining Fruit	1.08	-0.36	1.17
	(0.04)	(0.01)	(0.05)
Light Beer	0.34	-0.25	0.45
	(0.01)	(0)	(0.01)

Note: Estimates of utility parameters and rescaling parameter for each module. The parameter ρ is an estimate of the rescaling parameter which multiplies "new" products, defined as products (i.e. UPCs) which were not sold in 2006 in my data.

8.3 Vertical Differentiation

Finally, one may wonder whether vertical differentiation can explain the changing substitution patterns presented herein. If increasingly efficient supply chains have reduced firms' marginal costs significantly, some firms may have invested in higher quality goods which were

[44]One concern with these estimates is that σ_α is very large relative to $\bar{\alpha}$, implying that a large number of price coefficients are positive. I have also estimated models for some categories in which α_i is distributed as a log normal random variable. Though the search procedure in this model is more likely to generate numerical issues, estimates are similar when models converge. Estimates from such a model are presented in Table A.4.

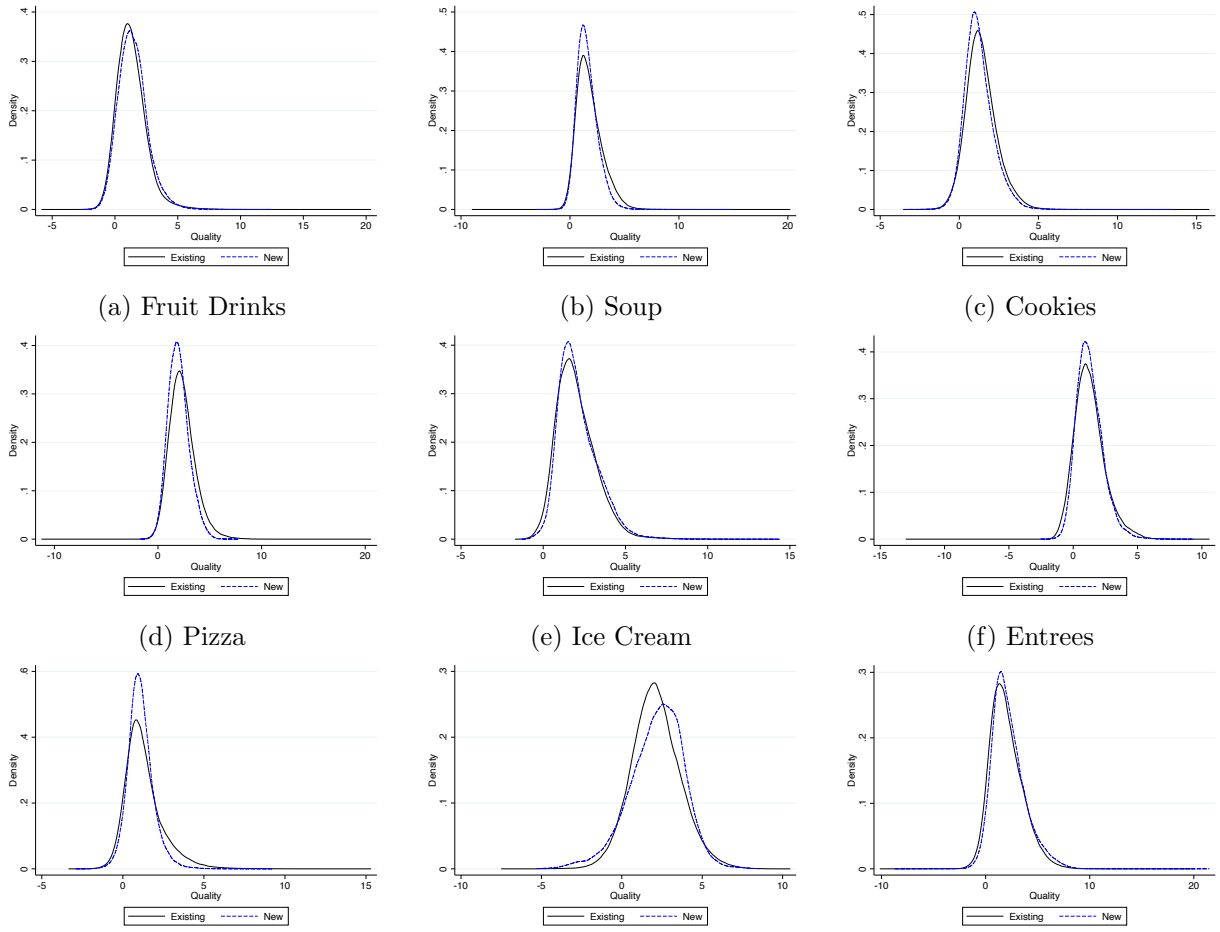
too expensive in the past. To address this possibility, I briefly summarize the extent to which the quality of newer goods differs from that of existing goods in Figure 14. In this figure, I plot the distribution of product fixed effects (which are estimated in the store-level logit regressions in Equation 14) for new and old/existing products in 2017. I define "new" slightly differently here. Rather than specifying a UPC as being "new" depending on whether it was ever sold in 2006, I call a product "new" if a that product was not sold *in a given store* in 2006. This definition allows us to quantify the extent to which newer goods in each store are preferred on average by the consumers in that store, without constraining average preferences (or the variance of logit errors) to be identical across stores.

Using this definition, I plot the full distribution of estimated fixed effects for "existing" (black, solid) and "new" (blue, dashed) products. Unlike in many of the preceding figures, the differences between the distributions in 2006 and 2017 are generally small. Consistent with Bronnenberg and Ellickson (2015), who comment that "the supply and availability of fresh products and the diversity and quality of products on offer" have improved, I find the most significant increase in this vertical characteristic to be in the subcategory of fruits in my sample (subfigure (h)). This also helps make sense of the fact that price elasticities have declined in this category even while newer products are less niche (according to Table 6). The other most notable differences are in the Soup, Yogurt, and Frozen Pizza categories, all of which appear to indicate that newer products rank slightly *lower* on vertical dimensions than newer. In general, these changes are small, bolstering the notion that horizontal differentiation has played the dominant role in rising differentiation.

9 Concluding Remarks

In this paper I have shown that, for many popular goods in food stores and mass merchandisers, consumers have become significantly less price sensitive over the past decade. I offer one explanation for this trend, that retailers have changed the assortment of products they offer in order to sustain larger markups. This finding relates to a growing literature which finds that profit rates, and markups in particular, have risen among many firms in the United States in recent decades (De Loecker, Eeckhout and Unger, 2020). What is perhaps most novel here is that the explanation I offer is entirely within firm. Even in the absence of changes in market or monopsony power, retailers may have been able to sustain large markups because the types of technological advances they have experienced have allowed them to offer more differentiated goods. Moreover, markups induced by this mechanism can in principle be welfare enhancing.

Figure 14: Distributions of Quality of New and Existing Products, by Module



(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Plotted are the distributions of product fixed effects in 2017 from the store-level logit model in Equation 14. Estimates of fixed effects of existing products are in black (solid) and those from new products are in blue (dashed).

Prices have moved little while consumers have become able to find products for which they have a high willingness to pay, meaning both producer and consumer surplus may have risen. My results are also consistent with recent evidence that the number of products households purchase in a given category has declined over time (Neiman and Vavra, 2019).

The strongest evidence I present here demonstrates that own- and cross-price elasticities have declined significantly between 2006 and 2017. These results are robust across nine product modules, which cover thousands of products sold in 5,000-10,000 stores (differing by module) across the country. In models with geographically heterogeneous preferences I show that the

disutility of price has declined significantly in many markets, meaning that the observed changes in substitution patterns are in part due to changing preferences (broadly defined). The results also hold when I estimate more than 100,000 store-level constant elasticity demand curves, which make strong assumptions about cross-product substitution but which are robust to misspecification of market sizes.

I then show that these changes are broadly consistent with the simple model in which the decline of stocking costs incentivizes firms to offer goods for which consumers have more heterogeneous preferences, which I call "niche" products. I do this by providing evidence for two predictions of the model, namely that the average variance of preferences for products in 2017 is larger than that in 2006, and in particular that newer products in 2017 are more niche than older products. I demonstrate the former via a counterfactual in which I eliminate the role of rising preference heterogeneity, and the latter by estimating a model which permits the effective "nicheness" to vary across products by rescaling the importance of idiosyncratic (horizontal) preferences relative to the utility derived from price and quality.

References

Abaluck, Jason, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc. 2020. "Mortality Effects and Choice Across Private Health Insurance Plans." *Working Paper*.

Adams, Brian, and Kevin R Williams. 2019. "Zone pricing in retail oligopoly." *American Economic Journal: Microeconomics*, 11(1): 124–56.

Arcidiacono, Peter, Paul B Ellickson, Carl F Mela, and John D Singleton. 2019. "The competitive effects of entry: Evidence from supercenter expansion." *Available at SSRN 3045492*.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The fall of the labor share and the rise of superstar firms." *The Quarterly Journal of Economics*, 135(2): 645–709.

Barkai, Simcha. 2020. "Declining labor and capital shares." *The Journal of Finance*, 75(5): 2421–2463.

Baumol, William J, and Edward A Ide. 1956. "Variety in retailing." *Management Science*, 3(1): 93–101.

Berry, Steven, Amit Gandhi, and Philip Haile. 2013. "Connected substitutes and invertibility of demand." *Econometrica*, 81(5): 2087–2111.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Bhardwaj, Vertica, and Ann Fairhurst. 2010. "Fast fashion: response to changes in the fashion industry." *The international review of retail, distribution and consumer research*, 20(1): 165–173.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. 2020. "Some Unpleasant Markup Arithmetic: Production Function Elasticities and their Estimation from Production Data." National Bureau of Economic Research.

Brand, James. 2019. "Estimating Productivity and Markups Under Imperfect Competition."

Bronnenberg, Bart J, and Paul B Ellickson. 2015. "Adolescence and the path to maturity in global retail." *Journal of Economic Perspectives*, 29(4): 113–34.

Butters, Andrew, Daniel Sacks, and Boyoung Seo. 2020. "How do national firms respond to local shocks? Evidence from excise taxes." *Working Paper*.

Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson. 2016. "Health care exceptionalism? Performance and allocation in the US health care sector." *American Economic Review*, 106(8): 2110–44.

Conlon, Christopher, and Jeff Gortmaker. 2019. "Best practices for differentiated products demand estimation with pyblp." Working paper.

Conlon, Christopher T, and Nirupama Rao. 2019. "Discrete prices and the incidence and efficiency of excise taxes." *Available at SSRN 2813016*.

Consumer Reports. 2014. "What to Do When There Are Too Many Product Choices on the Store Shelves?"

Correia, Sergio. 2016. "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator." Working Paper.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. "Uniform pricing in us retail chains." *The Quarterly Journal of Economics*, 134(4): 2011–2084.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The rise of market power and the macroeconomic implications." *The Quarterly Journal of Economics*, 135(2): 561–644.

De Loecker, Jan, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik. 2016. "Prices, markups, and trade reform." *Econometrica*, 84(2): 445–510.

Draganska, Michaela, Michael Mazzeo, and Katja Seim. 2009. "Beyond plain vanilla: Modeling joint product assortment and pricing decisions." *QME*, 7(2): 105–146.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference prices, costs, and nominal rigidities." *American Economic Review*, 101(1): 234–62.

Fan, Ying. 2013. "Ownership consolidation and product characteristics: A study of the US daily newspaper market." *American Economic Review*, 103(5): 1598–1628.

Fernie, John, and Nobukazu Azuma. 2004. "The changing nature of Japanese fashion." *European Journal of Marketing*, 38(7): 790–808.

Fernie, John, Leigh Sparks, and Alan C McKinnon. 2010. "Retail logistics in the UK: past, present and future." *International Journal of Retail & Distribution Management*, 38(11/12): 894–914.

Flynn, Zach, Amit Gandhi, and James Traina. 2019. "Identifying market power from production data." Working paper.

Foster, Lucia, John Haltiwanger, and Cornell 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.

Ganapati, Sharat. 2021. "Growing Oligopolies, Prices, Output, and Productivity." *Forthcoming, American Economic Journal: Microeconomics*.

Gandhi, Amit, and Jean-François Houde. 2019. "Measuring substitution patterns in differentiated products industries." National Bureau of Economic Research.

Grieco, Paul, Charles Murry, and Ali Yurukoglu. 2020. "The Evolution of Market Power in the US Auto Industry." *Working Paper*.

Hendel, Igal, and Aviv Nevo. 2006. "Measuring the implications of sales and consumer inventory behavior." *Econometrica*, 74(6): 1637–1673.

Hitsch, Günter J, Ali Hortacsu, and Xiliang Lin. 2019. "Prices and promotions in us retail markets: Evidence from big data." National Bureau of Economic Research.

Hortaçsu, Ali, and Chad Syverson. 2015. "The ongoing evolution of US retail: A format tug-of-war." *Journal of Economic Perspectives*, 29(4): 89–112.

Hsieh, Chang-Tai, and Esteban Rossi-Hansberg. 2019. "The industrial revolution in services." National Bureau of Economic Research.

Hull, Peter. 2018. "Estimating hospital quality with quasi-experimental data." *Available at SSRN 3118358*.

Jaumandreu, Jordi. 2018. "Dangerous Shortcuts: Ignoring Marginal Cost Determinants in Markup Estimation." mimeo, Boston University.

Kiedaisch, Christian, Andreas Chai, and Nicholas Rohde. 2018. "The saturation of spending diversity and the truth about Mr Brown and Mrs Jones." *Entropy*, 5: 6.

Kuhn, Heinrich, and Michael G Sternbeck. 2013. "Integrative retail logistics: an exploratory study." *Operations Management Research*, 6(1-2): 2–18.

Miller, Nathan H, and Matthew C Weinberg. 2017. "Understanding the price effects of the MillerCoors joint venture." *Econometrica*, 85(6): 1763–1791.

Morris, Carl N. 1983. "Parametric empirical Bayes inference: theory and applications." *Journal of the American statistical Association*, 78(381): 47–55.

Neiman, Brent, and Joseph S Vavra. 2019. "The rise of niche consumption." National Bureau of Economic Research.

Reynaerts, Jo, R Varadha, and John C Nash. 2012. "Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping." *Working Paper*.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter. 2018. "Diverging trends in national and local concentration." National Bureau of Economic Research.

Salanié, Bernard, and Frank A Wolak. 2019. "Fast, "robust", and approximately correct: estimating mixed demand systems." National Bureau of Economic Research.

Shin, Seungjae, and Burak Eksioglu. 2014. "Effects of RFID technology on efficiency and profitability in retail supply chains." *Journal of Applied Business Research (JABR)*, 30(3): 633–646.

Smith, Dominic, and Sergio Ocampo Díaz. 2020. "The Evolution of US Retail Concentration."

Soliman, Khalid S, Brian D Janz, Edmund Prater, Gregory V Frazier, and Pedro M Reyes. 2005. "Future impacts of RFID on e-supply chains in grocery retailing." *Supply Chain Management: An International Journal*.

Somaini, Paulo, and Frank A Wolak. 2016. "An algorithm to estimate the two-way fixed effects model." *Journal of Econometric Methods*, 5(1): 143–152.

Sparks, Leigh. 2010. "Supply chain management and retailing." Vol. 11, 4–12, Taylor & Francis.

Stamatopoulos, Ioannis, Achal Bassamboo, and Antonio Moreno. 2020. "The Effects of Menu Costs on Retail Performance: Evidence from Adoption of the Electronic Shelf Label Technology." *Management Science*.

Stansbury, Anna, and Lawrence H Summers. 2020. "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy." National Bureau of Economic Research.

Sullivan, Christopher. 2017. "The ice cream split: Empirically distinguishing price and product space collusion." *Available at SSRN 3321948*.

Varadhan, Ravi, and Christophe Roland. 2008. "Simple and globally convergent methods for accelerating the convergence of any EM algorithm." *Scandinavian Journal of Statistics*, 35(2): 335–353.

Waldfogel, Joel. 2003. "Preference externalities: an empirical study of who benefits whom in differentiated-product markets." *RAND Journal of Economics*, 34(3): 557–569.

A Omitted Figures and Tables

Table A.1: Changing Brand Structure and Characteristics

	2006	2017
Flavored Products	239885	188499
Unique Flavors	14378	29744
Organic	3518	16311

Note: Constructed from products files provided by Nielsen. Flavors are defined by Nielsen's unique flavor codes, and products are counted as "organic" if they are marked as having a USDA seal indicating the product is organic.

Table A.2: Correlation Between FRAC Estimates Across Modules, 2006 and 2017

2006	Fruit Drinks	Soup	Cookies	Pizza	Ice Cream	Entrees	Yogurt	Fruit
Fruit Drinks								
Soup	0.26							
Cookies	0.55	0.14						
Pizza	0.53	0.19	0.49					
Ice Cream	0.35	0.33	0.34	0.45				
Entrees	0.38	0.49	0.42	0.5	0.5			
Yogurt	0.29	0.41	0.32	0.35	0.41	0.41		
Remaining Fruit	0.28	0.39	0.49	0.48	0.11	0.51	0.27	
Light Beer	0.34	0.46	0.35	0.36	0.08	0.28	0.18	0.53
2017								
Fruit Drinks								
Soup	0.39							
Cookies	0.24	0.31						
Pizza	0.42	0.57	0.4					
Ice Cream	0.27	0.05	0.35	0.24				
Entrees	0.26	0.39	0.47	0.53	0.44			
Yogurt	0.05	0.03	0.37	0.14	0.05	0.12		
Remaining Fruit	0.23	-0.02	0.17	-0.06	0.13	0.03	0.23	
Light Beer	0.34	0.23	0.25	0.33	0.3	0.27	-0.04	0.08

Note: Correlation matrix of all estimates of the mean coefficient on price in consumer utility functions ($\bar{\alpha}$) in 2006 and 2017.

Table A.3: Estimated Price Elasticities by Year, Constant Elasticity

	2006		2017	
	Estimates	Post-Shrinkage	Estimates	Post-Shrinkage
Mean	-1.87	-1.94	-1.54	-1.56
5%	-3.99	-3.74	-3.11	-3.02
25%	-2.74	-2.66	-2.05	-1.99
Median	-1.97	-1.97	-1.49	-1.48
75%	-0.92	-1.15	-0.93	-1.03
95%	0.24	-0.12	-0.22	-0.42

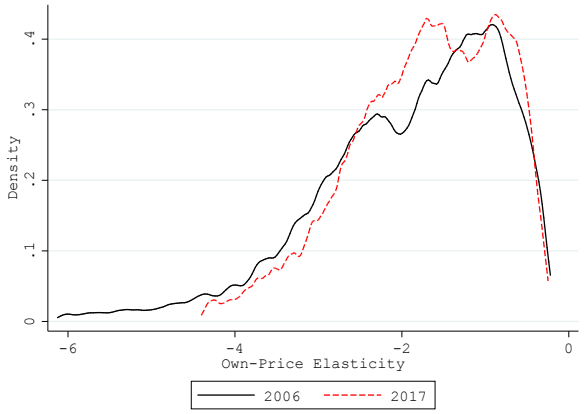
Note: Distribution of estimated and post-shrinkage (empirical Bayes) estimates of price elasticities η_s in 2006 and 2017, estimated in the constant elasticity model in Equation 11.

Table A.4: Estimates of Scaling Parameter with Log-Normal Preferences

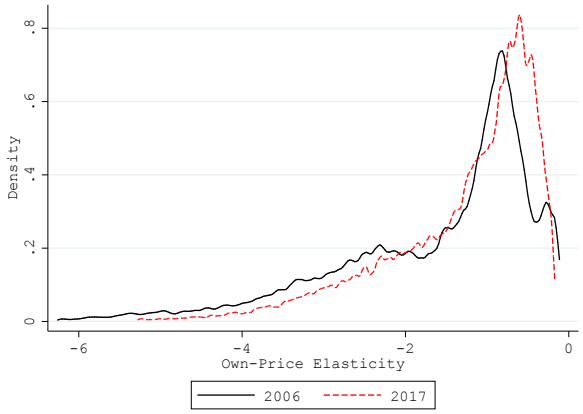
	ρ
Fruit Drinks	
Soup	0.88
	(0.04)
Cookies*	0.93
	(0.14)
Pizza*	0.77
	(0.04)
Ice Cream	0.90
	(0.10)
Entrees*	0.75
	(0.06)
Yogurt	
Remaining Fruit	1.07
	(0.04)
Light Beer	0.61
	(0.02)

Note: Estimates of scaling parameter ρ from Equation 16, with α_i distributed according to a log normal distribution which is constrained to be negative. Modules with asterisks were estimated using 200 randomly sampled stores; otherwise, 400 stores were used. The estimation procedures for Fruit Drinks and Yogurt do not converge and often yield numerical issues.

Figure A.1: Own Elasticity Distribution, Additional Modules



(a) 4100- Eggs



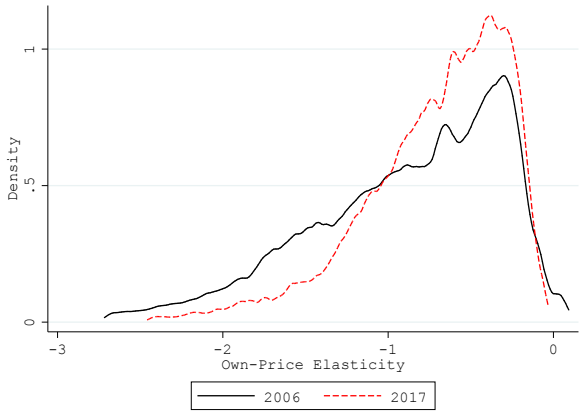
(b) 1484- Sodas



(c) 1553- Diet Sodas



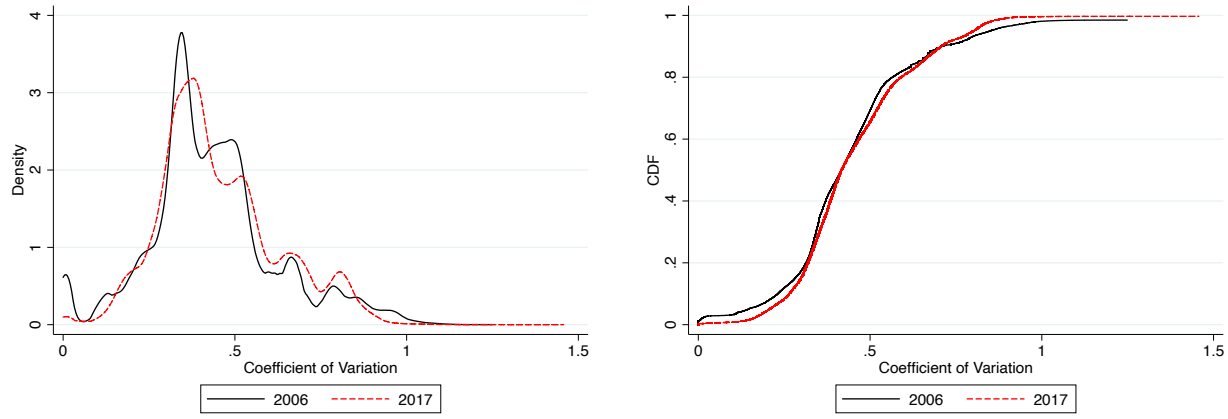
(d) 3574- Sliced Lunchmeat



(e) 1272- Baby Milk

Note: Distribution of estimated own-price elasticities. Estimates come from store-level logit models in each module in each year.

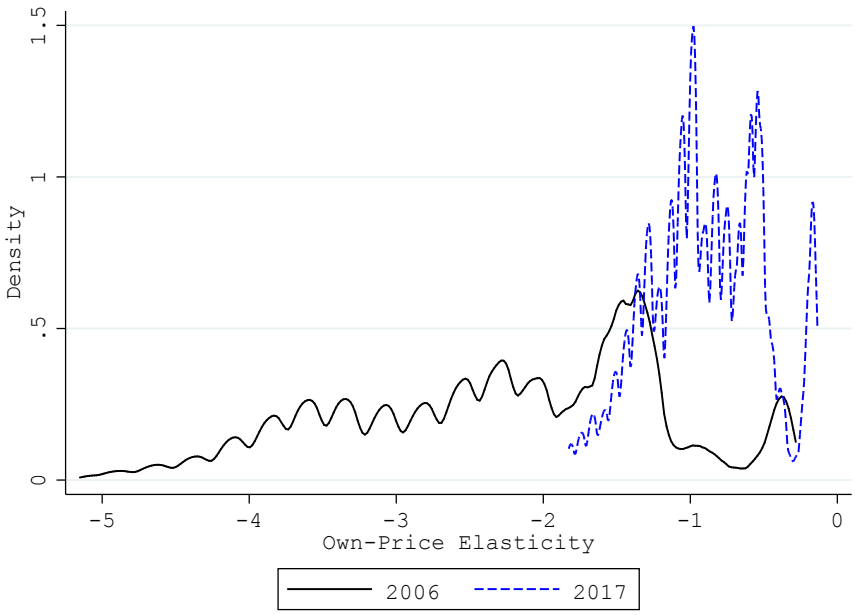
Figure A.2: Coefficients of Variation Within Module, 2006 and 2017



(a) (b)

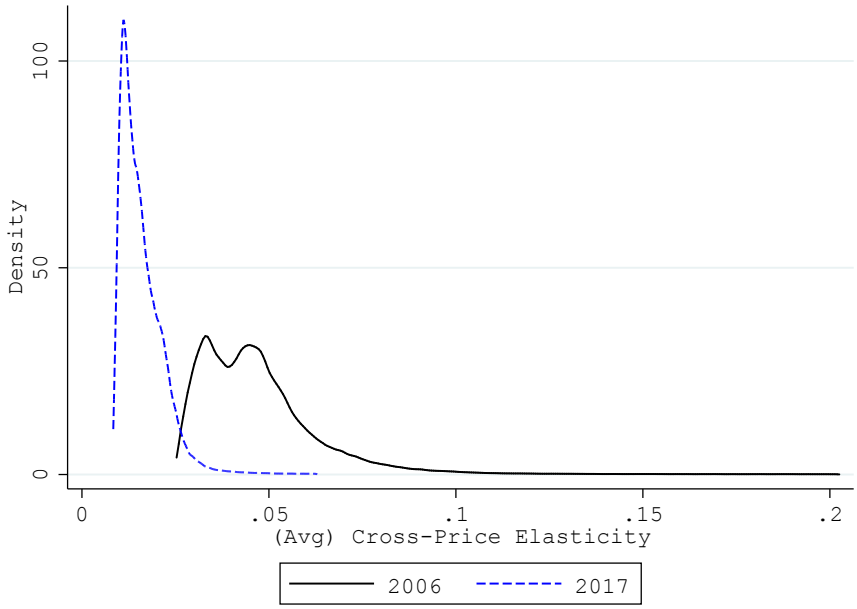
Note: Constructed from scanner files in Nielsen. Figure (a) presents the estimated kernel density of coefficients of variation within a store-module in the first week of 2006 (black, solid) and 2017 (red, dashed). Figure (b) presents empirical cumulative distribution functions.

Figure A.3: Own-Price Elasticity Distribution, Nested Logit, Light Beer



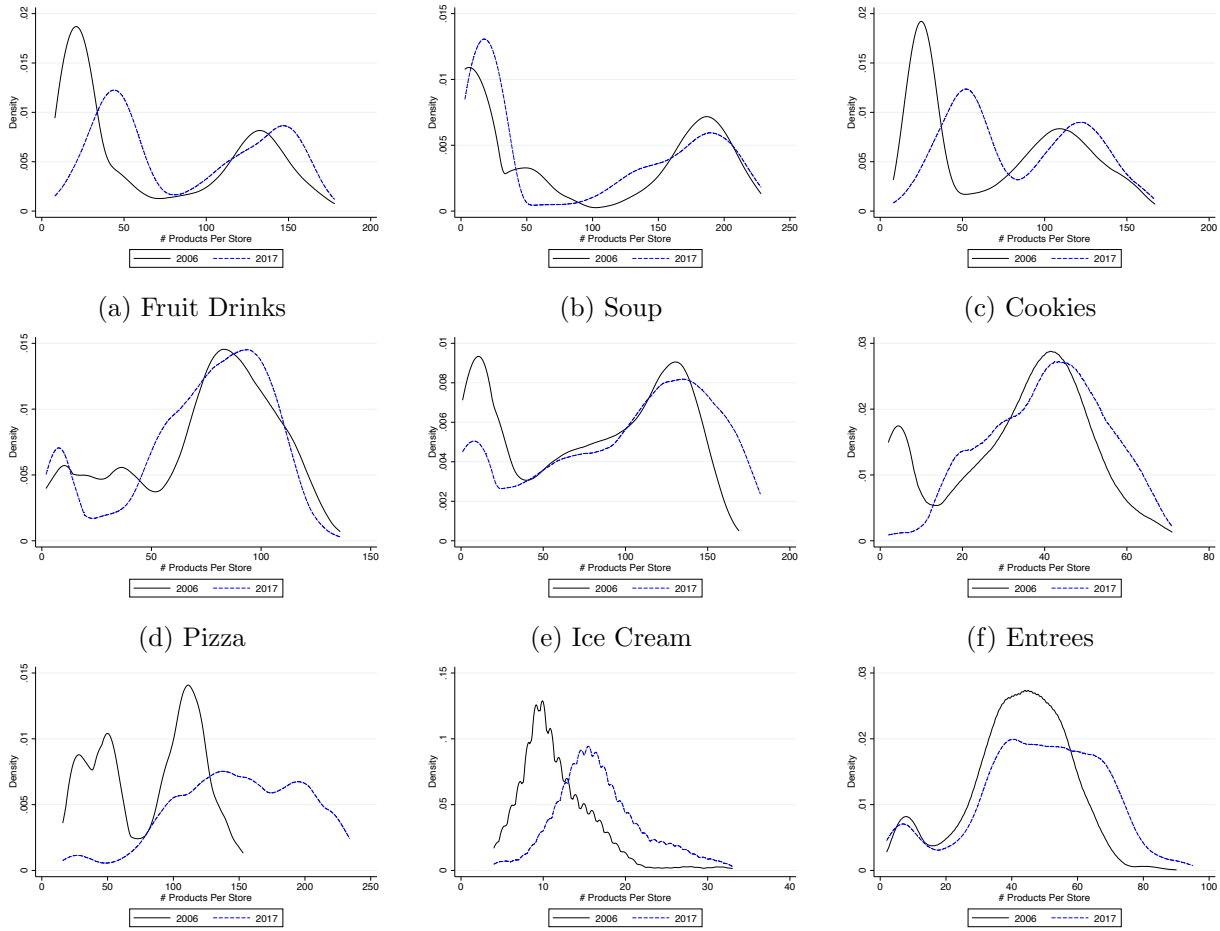
Note: Estimated own-price elasticities in 2006 and 2017 for the Light Beer module. Estimates are from a nested logit model estimated separately in each year with UPC and week fixed effects.

Figure A.4: Average Cross-Price Elasticity Distribution, Nested Logit, Light Beer



Note: Estimated average cross-price elasticities in 2006 and 2017 for the Light Beer module. Estimates are from a nested logit model estimated separately in each year with UPC and week fixed effects.

Figure A.5: Distributions of the Number of Products Sold, by Module



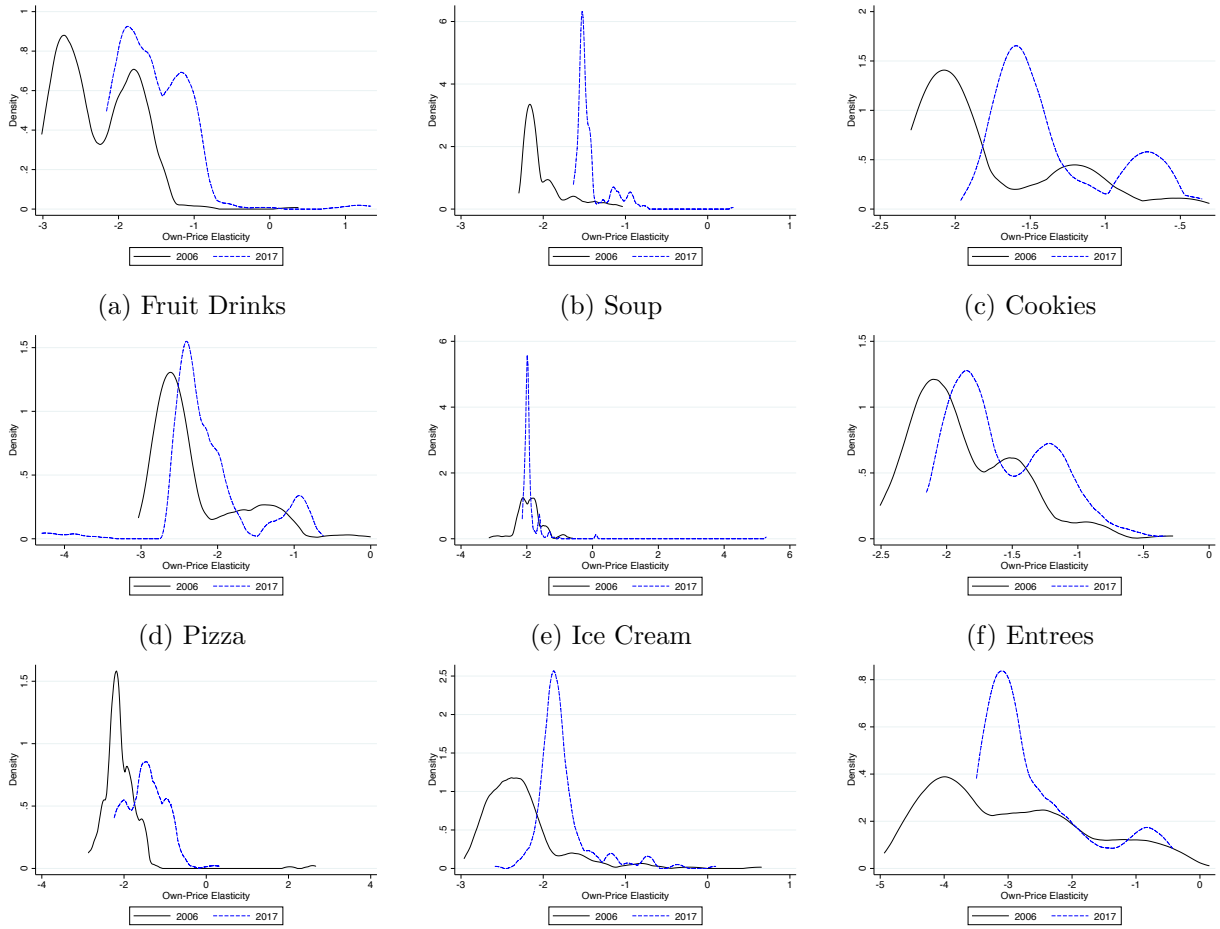
(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Plotted are the distributions (across stores) of the number of products in my sample sold in the first week of 2006 (black, solid) and 2017 (blue, dashed).

Figure A.6: Distributions of Mean Own-Price Elasticities by UPC, BLP



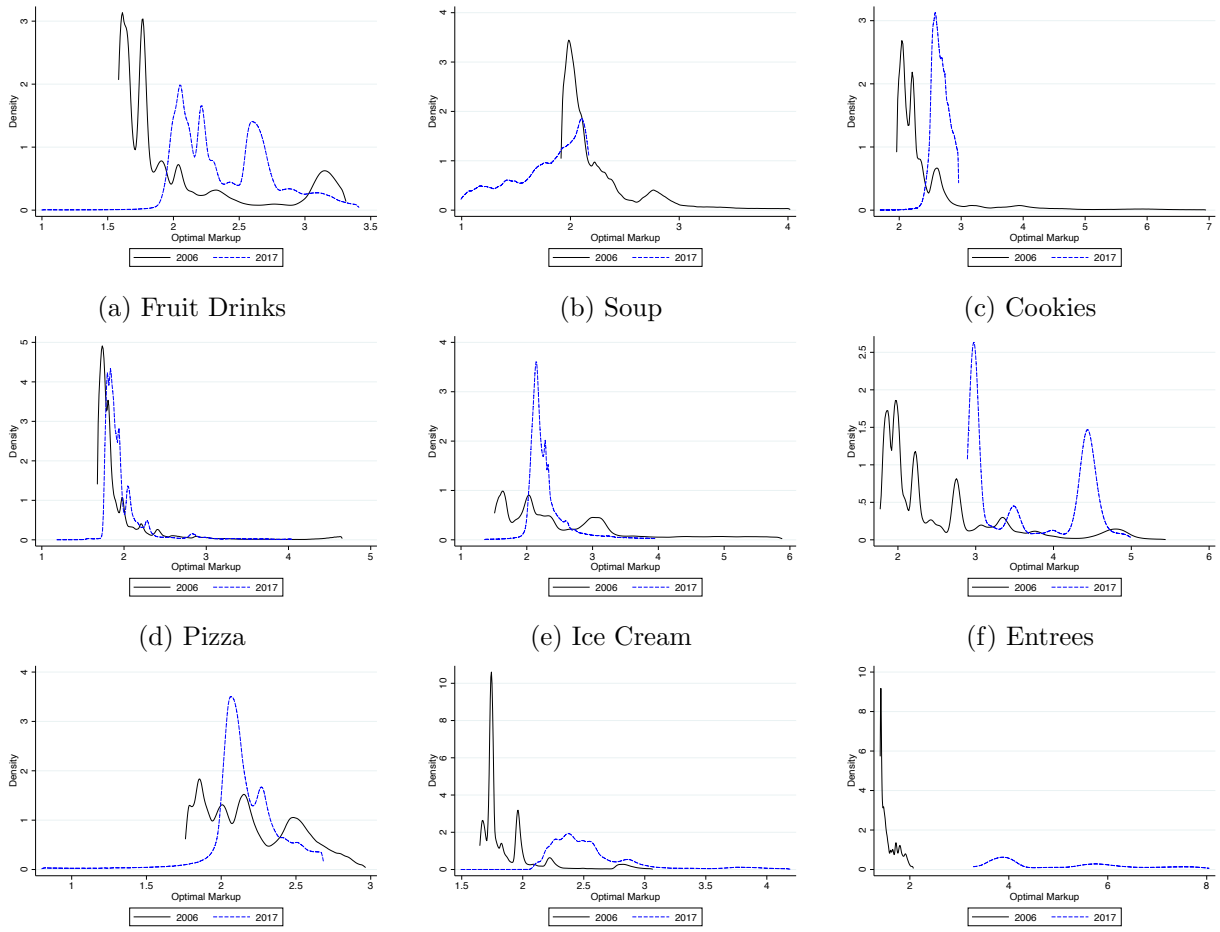
(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Plotted are the distributions (across UPCs) of average own-price elasticities in 2006 (black, solid) and 2017 (blue, dashed) implied by estimates of the BLP-style model.

Figure A.7: Distributions of Optimal Markups by Module, BLP Estimates



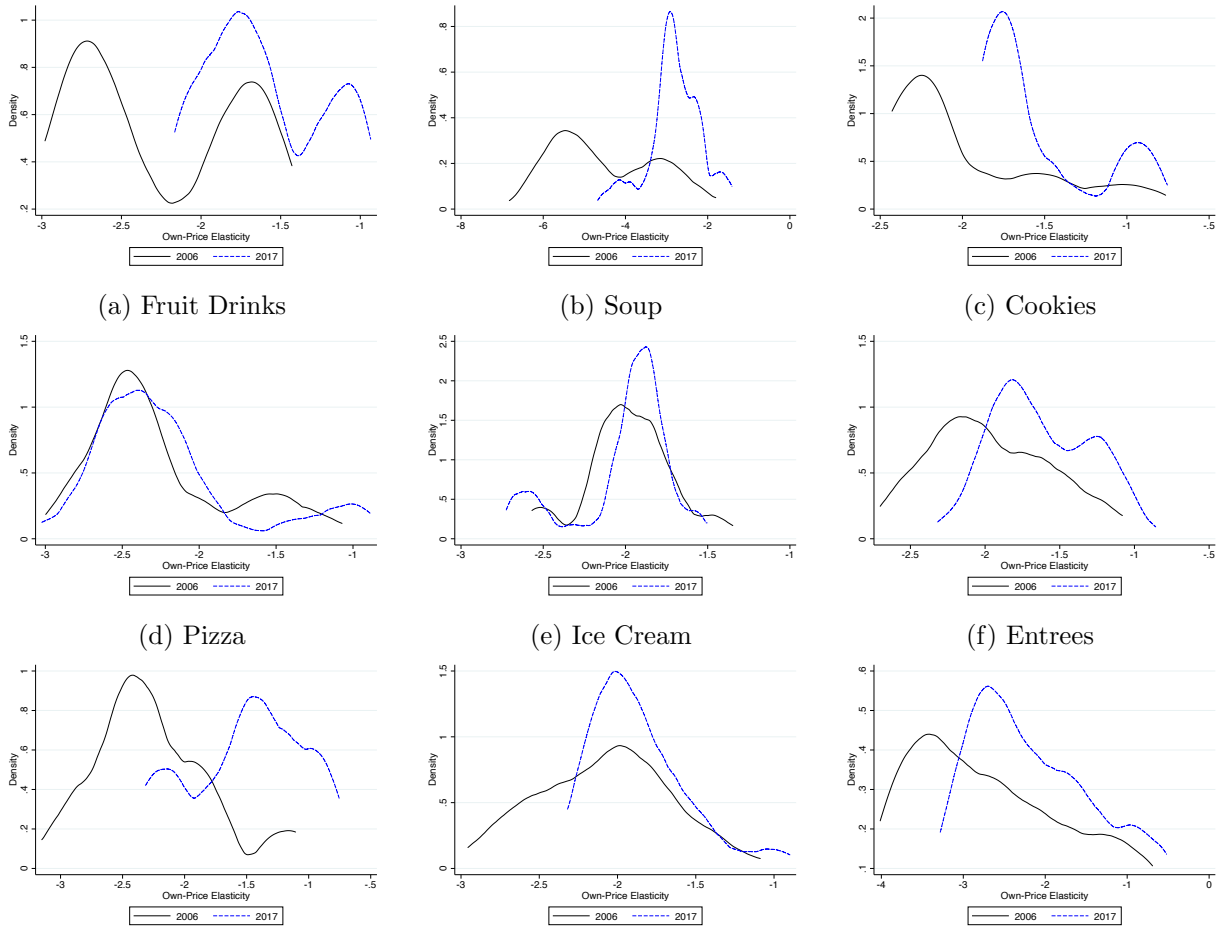
(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Plotted are the distributions of markups in 2006 (black, solid) and 2017 (blue, dashed) according to BLP estimates, conditioning on estimates with positive implied marginal costs and trimming the top and bottom five percent of the distributions.

Figure A.8: Distributions of Mean Own-Price Elasticities by UPC, FRAC



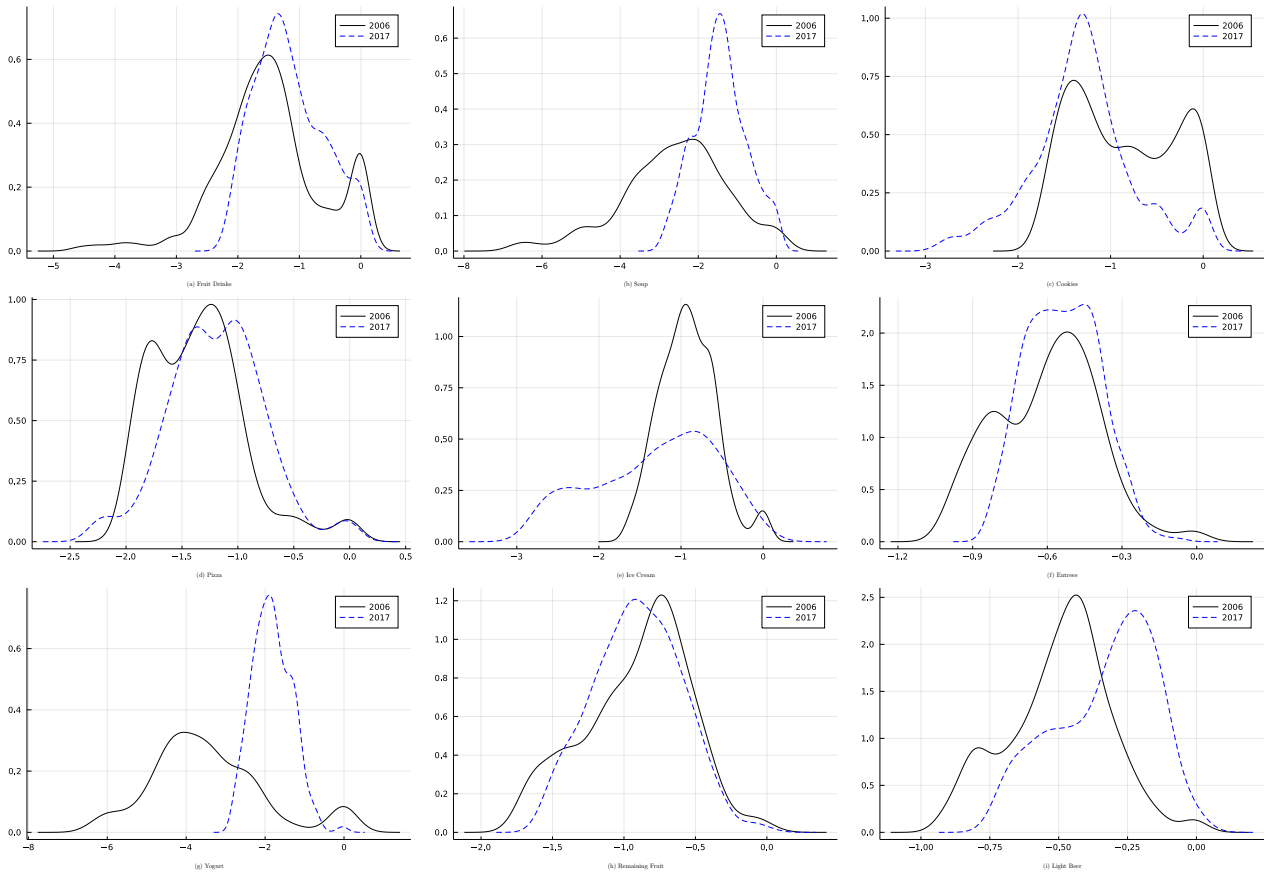
(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

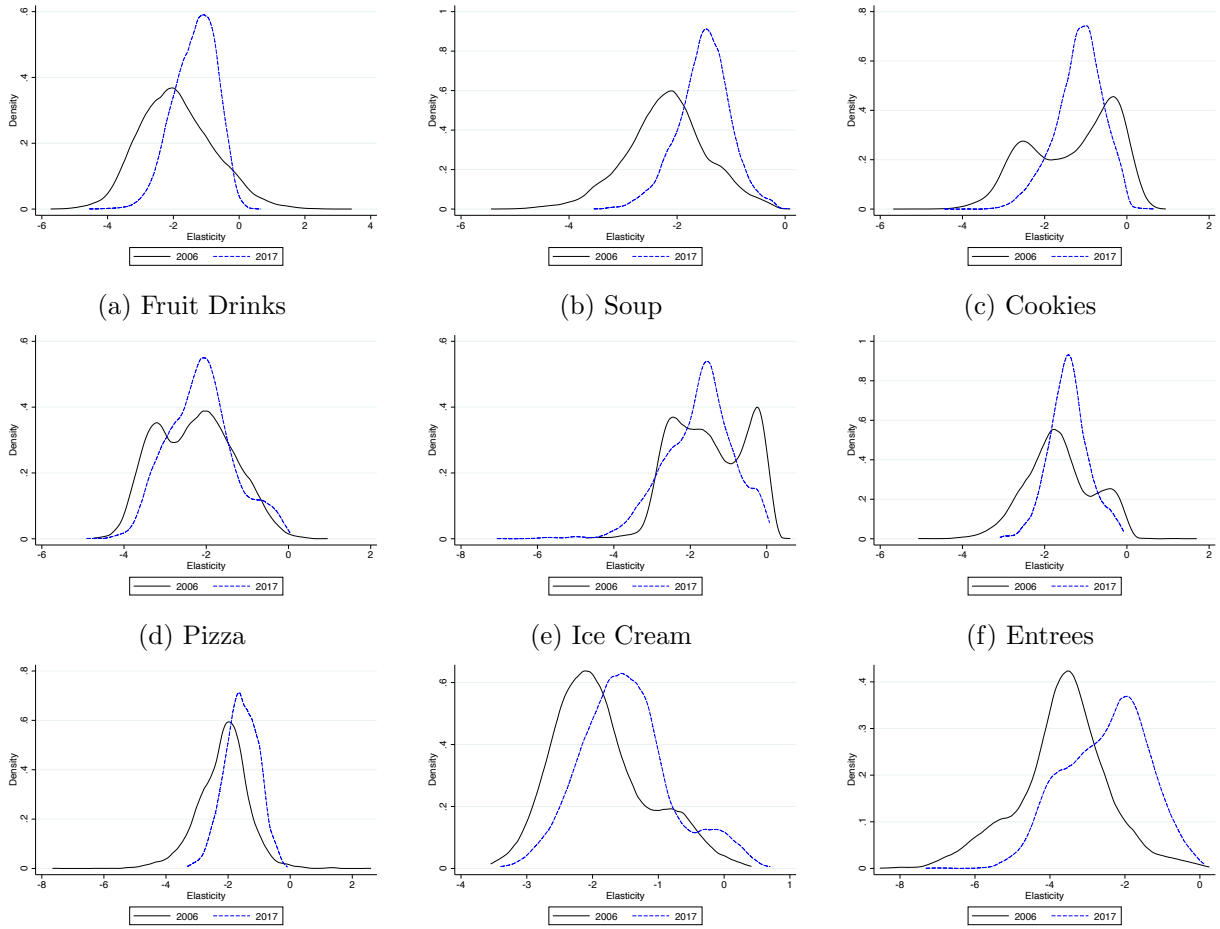
Note: Plotted are the full distributions (across UPCs) of average own-price elasticities in 2006 (black, solid) and 2017 (blue, dashed) implied by FRAC estimates of demand.

Figure A.9: Distribution of Estimates of ZIP Code-Level Price Coefficients, Instrumenting for Price



Note: Plotted are the distributions of mean preferences for price in 2006 (black, solid) and 2017 (blue, dashed) in each module, estimated by FRAC. The bottom 10% of estimates in each module-year are excluded from these figures and only ZIP codes with at least 500 observations are included. In contrast to the main estimates in the paper, these estimates instrument for p_{jt} at a given store using the average price of good j in week t at all other stores within a three-digit ZIP code. This instrument serves as cost-shifter under the assumptions that the marginal costs associated with a given UPC are similar in nearby stores and that demand shocks (net of absorbed fixed effects) are uncorrelated across stores.

Figure A.10: Distributions of the Store-Level Price Elasticities, by Module



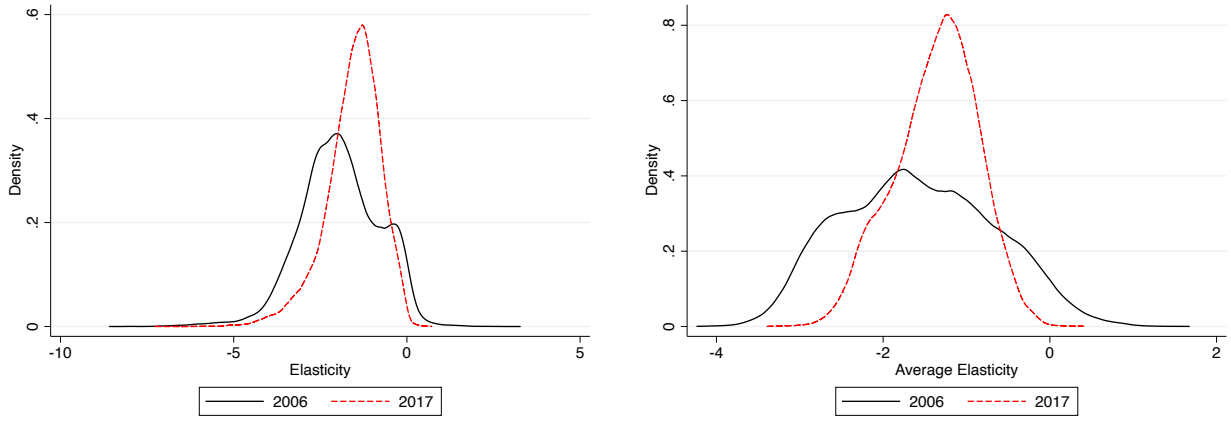
(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Plotted are the distributions (across stores) of estimated price elasticities (η_s) in the constant elasticity demand model. Estimates from 2006 are in black (solid) and those from 2017 are in blue (dashed).

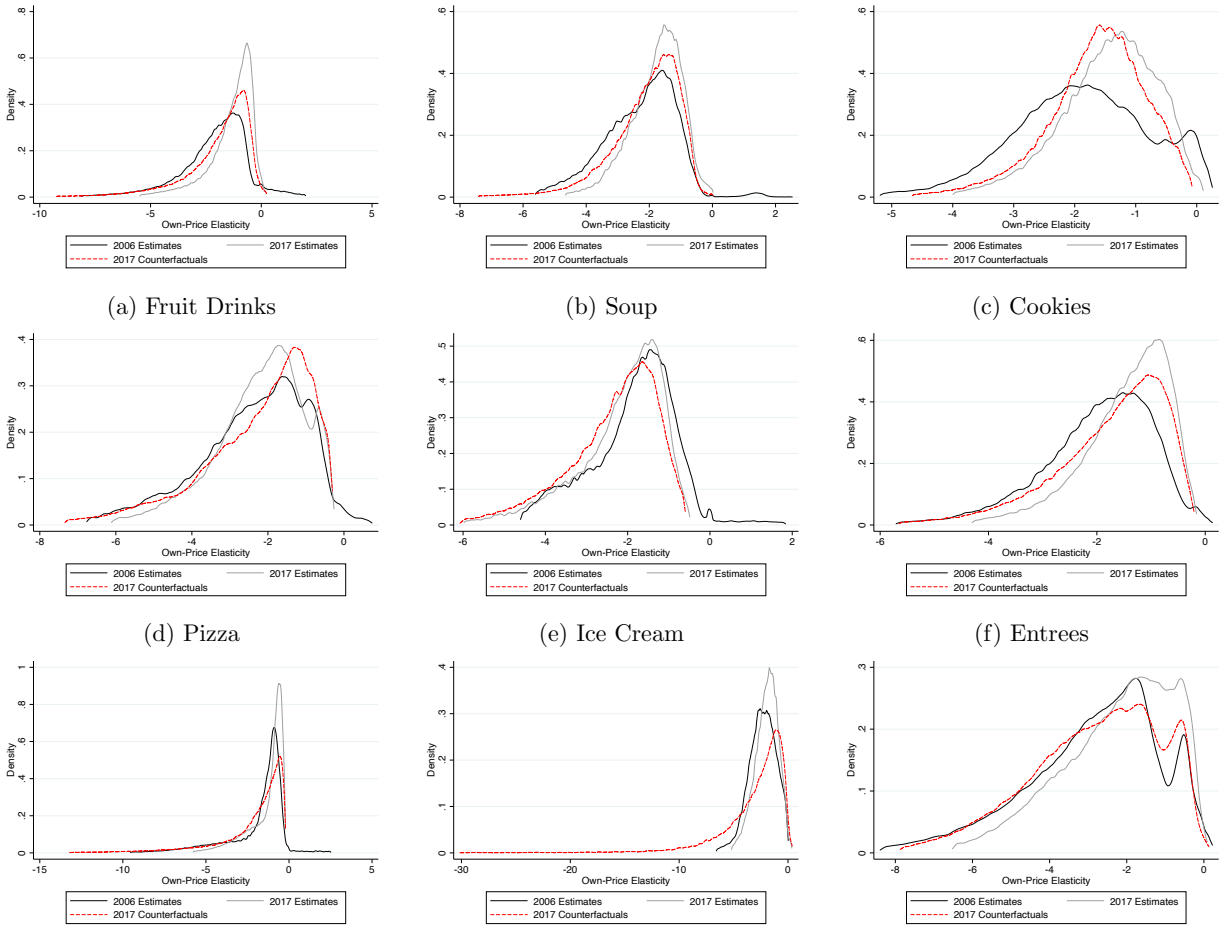
Figure A.11: Distributions of Price Elasticities (Constant Elasticity), Controlling for Competing Prices



(a) (b)

Note: Figure (a) presents the distribution of estimates of η_s, the store-module-level price elasticity estimated from regressions of the form $log(s_{jsw}) = \eta_s log(p_{jsw}) + \beta_s log(\bar{p}_{-jt}) + \bar{\xi}_{js} + \Delta\xi_{jsw}$, where the term $\bar{p}_{-jt}$ is the average price of other goods within a store-module-week and each regression includes UPC fixed effects. Figure (b) presents the distribution of store-level averages of these estimates (i.e. averaging across the modules offered by a store).

Figure A.12: Effects of Increasing Preference Heterogeneity



(a) Fruit Drinks (b) Soup (c) Cookies

(d) Pizza (e) Ice Cream (f) Entrees

(g) Yogurt (h) Remaining Fruit (i) Light Beer

Note: Distributions of estimated own-price elasticities in 2006 (black) and 2017 (grey), as well as counterfactual price elasticities which set preference heterogeneity to 2006 levels (red, dashed). Estimates come from logit models estimated at the store-module-year level (Equation 14).

B Selecting Product Categories

As described in the text, in order to exclude any product categories in which consumers exhibit sales targeting behavior, I conduct a test for many product modules which draws on the arguments in Hendel and Nevo (2006): if consumers are unable to store goods, then the amount purchased on one trip to the store should not predict the time until the next trip. The goal here is to test the null hypothesis that consumers are not storing goods, and to restrict my sample to modules for which I cannot reject the null. As I observe nearly every trip taken by the households in the Homescan data, we can do so by testing whether the quantity purchased on a given date predicts the amount of time which passes before a consumer purchases again. To be precise, I run the following regression separately for 40 modules in 2006:

$$Time\ Since_{it} = f(p_{it}) + \beta q_{it-1} + \theta_h + \delta_j + \eta_{it} \tag{17}$$

where *Time Since* is the number of days since the most recent purchase before trip t by consumer i in a given module.[45] In this regression p_{it} and q_{it-1} are current price paid and previous quantity purchased, and θ_h and δ_j are household and product fixed effects respectively.[46] The coefficient β is the focus of the test, and the null hypothesis is that $\beta = 0$. Household fixed effects are necessary here because some households may consistently purchase higher quantities. Product fixed effects control for product size and whether or not the product is a pack of multiple smaller units. Without these fixed effects, consumers switching between 6- and 12-packs of sodas (for example) could bias estimates of β. Finally, I instrument for q_{it-1} with p_{it-1}, which links the regression to the thought experiment more directly: if additional units purchased previously *as a result of a sale* increase the time until a household purchases again, then consumers may be storing goods and targeting sales.

Consistent with intuition, the modules selected in this paper tend to be either perishable (e.g. Yogurt) and/or not easily stored in large quantities as other modules (e.g. Frozen Pizza). However, two points should be emphasized. First, the test above is for storage and sales targeting *behavior*, not for whether or not a good is perishable. Soup, which is included in my sample, can be stored for months in a pantry. What this test implies is that consumers are unlikely to purchase additional units of soup in response to a sale, or to avoid purchasing

[45] Before running these regressions, I restrict the sample to trips (1) which were the only trip taken by that household on a given day, and (2) in which the household bought only one product within the module of interest. These are necessary to make q_{it-1} well-defined and to make *Time Since* measurable (as I only observe the day of any given trip).

[46] I specify $f(\cdot)$ as a cubic function of price.

by consuming from an inventory at home. Second, although I argue that this test is one reasonable way to reduce the number of products to be studied, I do not claim that this is necessarily crucial to my results which follow. Even if consumers maintain inventories, most of the questions of interest here about *changes* in substitution patterns over time. As long as consumers target sales and store inventories at the same rate in 2006 and 2017, my conclusions will be unaffected.[47] The full set of p-values for all candidate modules are in Table B.1.

C Panel of Yogurt Data

In this section I restrict my focus to Refrigerated Yogurt, in order to demonstrate that the results in the paper have been happening gradually over time. If data from 2017 or 2006 were anomalous, we might worry that the results in the paper cannot be related to growing selection within stores, which has been gradual over time. As a separate issue, one may also worry that although I use deflated 2017 (i.e. real) prices throughout the paper, nominal prices are more relevant for demand estimation.

Including data from intermediate years will shed light on these concerns. For the results in this section, I use data from the first 6 weeks of 2006, 2008, 2010, 2013, 2015, and 2017. I begin by estimating a single logit model in each year, under four specifications: (1) all stores, and no controls, (2) only stores which are present in all years of the panel, (3) the same restricted sample, plus store-UPC and store-week fixed effects, and (4) the same as (3) but using nominal prices only. I plot the estimates of the coefficient on price over this panel in Figure C.1. Standard errors (clustered at the store-year level) on these estimates are also included in the graph, though they are generally small enough to be hidden by the marker. Each of these estimates indicates a trend in the price coefficient over time. Adding fixed effects naturally changes the price coefficient significantly, but restricting to a balanced panel of stores and using nominal prices has minimal impact.

Next, I estimate a logit model for every store in every year with more than 250 observations in a given year (most stores with fewer observations are dropped anyway when absorbing fixed effects). I include UPC and week fixed effects in every regression. Unlike in the main text, I do not restrict these estimates to stores which generate a balanced panel. This amounts to nearly 60,000 separate logit estimates in total over the panel, from which I present two results. First, in Figure C.2 I plot the mean, 25^{th}, and 75^{th} percentiles of the distribution of own-price elasticities

[47] Although the story I am most concerned about predicts that $\beta > 0$, in practice many estimates are negative. As this still implies that past purchases predict future behavior (albeit in a more complicated story), I treat these categories identically to those with $\hat{\beta} > 0$.

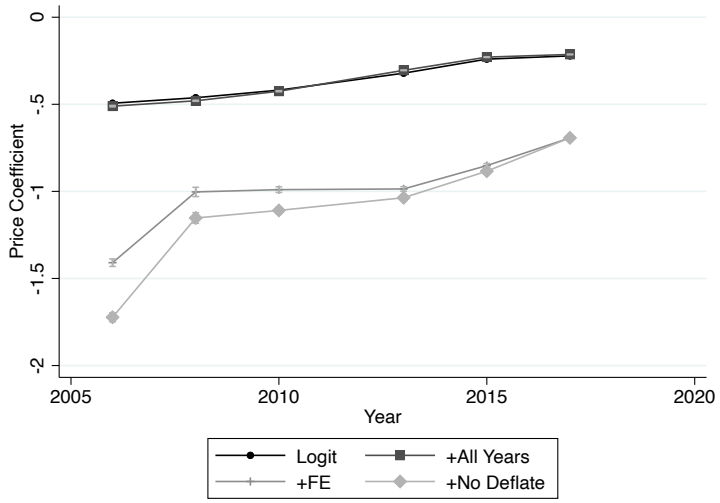
Table B.1: Results of All Tests for Inventory Behavior

Module Description	p-value	N
FRUIT- ORANGE- OTHER CONTAINER	< 0.001	189210
FRUIT DRINKS - OTHER CONTAINER	**0.991**	**84271**
BABY MILK AND MILK FLAVORING	0.217	4223
SOUP-CANNED	**0.644**	**98075**
CAT FOOD- WET TYPE	0.015	26209
DOG FOOD - DRY TYPE	< 0.001	58268
SNACKS - TORTILLA CHIPS	0.107	107960
CEREAL - READY TO EAT	< 0.001	188801
COOKIES	**0.993**	**197003**
GROUND AND WHOLE BEAN COFFEE	< 0.001	101399
SOFT DRINKS - CARBONATED	< 0.001	132651
WATER- BOTTLED	0.110	68437
CANDY- CHOCOLATE	0.015	158994
CANDY - NON-CHOCOLATE	0.010	107560
SOFT DRINKS - LOW CALORIE	0.001	119141
ENTREES- ITALIAN-1 FROOD - FROZEN	0.010	44622
PIZZA - FROZEN	**0.972**	**74221**
ICE CREAM - BULK	**0.675**	**142015**
FROZEN NOVELTIES	0.177	79076
LUNCHMEAT - SLICED-REFRIGERATED	0.164	105761
FRANKFURTERS - REFRIGERATED	< 0.001	83761
BACON- REFRIGERATED	< 0.001	108521
ENTREES - REFRIGERATED	**0.576**	**43736**
CHEESE- SHREDDED	< 0.001	100283
YOGURT - REFRIGERATED	**0.808**	**92011**
LUNCHMEAT - DELI POUCHES - REFRIGERATED	0.045	53413
DAIRY-MILK-REFRIGERATED	0.002	764194
BAKERY-CAKES-FRESH	0.056	82168
EGGS-FRESH	< 0.001	365468
FRESH FRUIT - REMAINING	**0.515**	**63883**
BEER	0.052	15945
LIGHT BEER (LOW CALORIE/ALCOHOL)	**0.742**	**13344**
DETERGENTS -HEAVY DUTY- LIQUID	< 0.001	97923
BLEACH-LIQUID/GEL	0.005	32798
TOILET TISSUE	< 0.001	184531
PAPER TOWELS	< 0.001	126729
BATTERIES	< 0.001	56657
ANTACIDS	< 0.001	22113
PAIN REMEDIES - HEADACHE	< 0.001	52689
COLD REMEDIES - ADULT	< 0.001	33085
DISPOSABLE DIAPERS	< 0.001	18024

Note: Reported are p-values for the test of the null hypothesis that $\beta = 0$ and sample sizes in estimating Equation 17 for each module. Categories come from Table 2b of the Online Appendix of DellaVigna and Gentzkow (2019), which lists many of the highest revenue products in the Nielsen data.

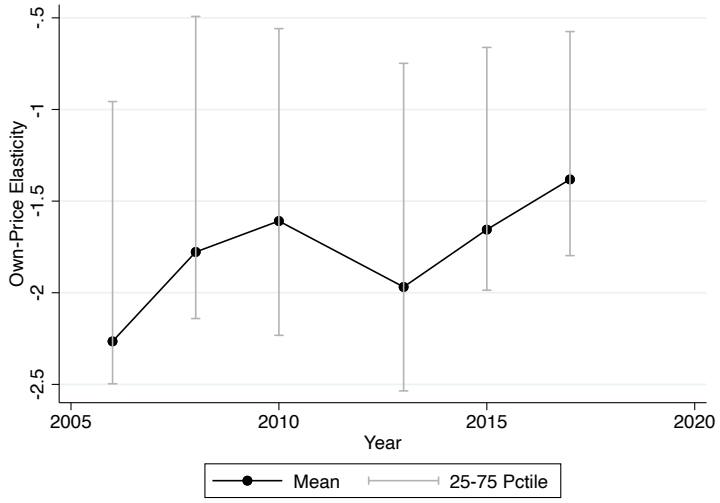
for each year of the panel. Although the trend here is not perfect (e.g. the decline between 2010 and 2013), the hypothesis that either of the end years is anomalous seems unlikely. The mean in 2008 is substantially higher than that in 2006, the mean in 2010 is higher still, and the mean in 2017 is the most positive of the sample. Given potentially many changes in this sample, e.g. the rise of greek yogurt and the complicated price responses which may have occurred, this figure is strong evidence of a trend. Second, in Figure C.3 I plot the same moments of the distribution of estimated utility coefficients on price over time. Here the trend is even more clear. Except for the slight decline at the mean in 2013, the average estimated coefficient has increased steadily over time.

Figure C.1: Utility Coefficients on Price Over Time, Refrigerated Yogurt



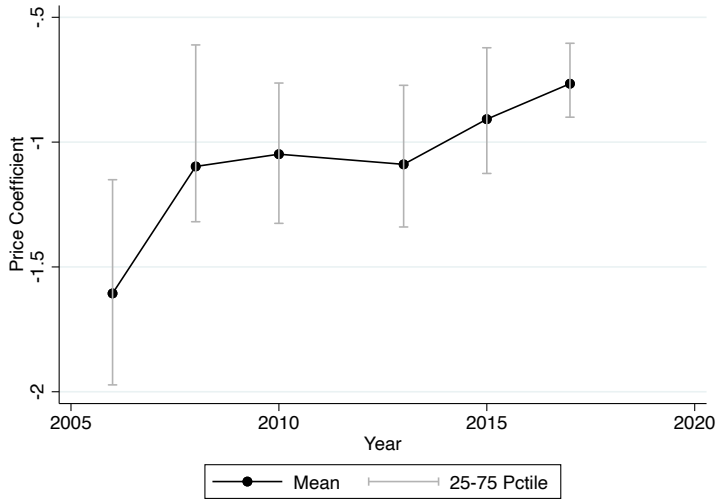
Note: Estimated coefficients on price from four simple logit models: (circle) all stores in each year, and no controls, (square) only stores which are present in all years of the panel, (plus) the same restricted sample, plus store-UPC and store-week fixed effects, and (diamond) the same as (plus) but using nominal prices only.

Figure C.2: Own Price Elasticity Distribution by Year, Refrigerated Yogurt



Note: Mean and interquartile range of own-price elasticities for Refrigerated Yogurt for selected years between 2006 and 2017.

Figure C.3: Store-Level Utility Coefficients Over Time, Refrigerated Yogurt



Note: Mean and interquartile range of estimated price coefficients for Refrigerated Yogurt in store-level logit models.