Rising Markups and the Role of Consumer Preferences*

Hendrik Döpper
DICE[†]

Alexander MacKay
Harvard University[‡]

Nathan H. Miller
Georgetown University[§]

Joel Stiebale
DICE[¶]

June 5, 2024

Abstract

We characterize the evolution of markups for consumer products in the United States from 2006 to 2019. Using detailed data on prices and quantities for products in more than 100 distinct product categories, we estimate flexible demand systems and recover markups under an assumption that firms set prices to maximize profit. Our empirical strategy obtains a panel of consumer preferences and marginal costs based on the estimation of separate random coefficient models by category and year. We find that markups increased by about 30 percent on average over the sample period. The change is primarily attributable to decreases in marginal costs, as real prices only increased slightly from 2006 to 2019. Our estimates indicate that consumers have become less price sensitive over time.

JEL Codes: D2, D4, L1, L2, L6, L81

Keywords: Market Power, Markups, Demand Estimation, Consumer Products, Retailers

*We thank Chris Conlon, Charlie Murry, Ariel Pakes, and Gloria Sheu for helpful comments. We are grateful for the comments and suggestions from many seminar and conference participants. Researchers 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. Computational support and infrastructure was provided by the "Centre for Information and Media Technology" (ZIM) at Heinrich Heine University Düsseldorf. Hendrik Döpper gratefully acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG) (project 235577387/GRK1974).

[†]Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf. Email: doepper@dice.hhu.de.

[‡]Harvard University, Harvard Business School. Email: amackay@hbs.edu.

[§]Georgetown University, McDonough School of Business. Email: nathan.miller@georgetown.edu.

[¶]Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf. Email: stiebale@dice.hhu.de.

1 Introduction

Firms with market power set prices that reflect marginal costs, consumer preferences, and the prices of related products. Economic theory indicates that differences between prices and marginal costs—the markups—can have wide-ranging implications for market outcomes. All else equal, an increase in markups transfers wealth from consumers to producers and can cause consumers to change their purchase decisions. These effects lead to less efficient resource allocation and, through reduced production, affect the markets for inputs, such as labor. Changes in markups may also affect the long-run dynamics in an industry by distorting investment and innovation incentives (Aghion et al., 2005). Thus, empirical evidence that markups are rising in the United States and abroad (e.g., De Loecker et al., 2020; Ganapati, 2021; De Loecker and Eeckhout, 2021) raises important questions for economic policy.

In this paper, we study the markups that arise in the U.S. economy across a large number of firms and products, and examine how supply and demand conditions have influenced firms' pricing decisions over time. Our approach is to estimate models of differentiated-products demand and Nash-Bertrand pricing for more than 100 consumer product categories—such as cereals, shampoo, and over-the-counter cold medications—using data on prices, quantities, and consumer purchasing patterns. The models allow us to recover markups from estimated demand elasticities and first order conditions for profit maximization. We use the models to consider potential mechanisms that contribute to trends in markups, and to connect markups to other economic outcomes, such as consumer surplus and deadweight loss.

We contribute two main empirical findings to the literature. First, we estimate that average product-level markups increase by about 30 percent between 2006 and 2019. Second, we show that the increase in markups can be explained by a combination of declining marginal costs and a reduction in consumer price sensitivity. We also find that, despite rising markups, the standard measure of consumer surplus for the products in our sample increases over time, though the change in consumer surplus varies across the income distribution. These findings provide additional context for how to interpret changes in markups in the economy more broadly.

Our models feature the random coefficients logit demand system that is standard in industrial organization (e.g., Berry et al., 1995) and commonly applied in industrial organization to consumer product markets (e.g., Nevo, 2001). The demand system can accommodate flexible patterns of consumer substitution at the product-level. Given demand parameters and data on prices and quantities, one can recover marginal costs, and thus also markups, from the first order conditions that govern price-setting behavior. We do so under the assumption of profit maximization by competing firms. This *demand approach* to recovering markups allows us to use counterfactual simulations to evaluate how particular mechanisms affect equilibrium outcomes and welfare.[1]

[1] Our approach is distinct from the *production approach* (e.g., Hall, 1988; De Loecker and Warzynski, 2012) that uses a cost-minimization assumption to estimate markups without modeling demand or market competition.

We estimate the models using detailed sales information from the Kilts NielsenIQ Retail Scanner Data. We aggregate across UPCs to construct quantities and unitized prices of products (specific brands) for each retail chain, region, and quarter. Data from Capital IQ and Zephyr allows us to identify the parent company of each product in each year, which is important for the supply side of the model. We also use the Kilts NielsenIQ Consumer Panel Data, which contains household-level purchases and demographic information, to help capture differences in product preferences across households and to control for differential selection into the retail chains in our sample by consumers with different demographics.

Most previous applications that employ the random coefficients logit demand model focus on a single product category, such as ready-to-eat cereal (Nevo, 2001; Backus et al., 2021), beer (Miller and Weinberg, 2017), or yogurt (Villas-Boas, 2007; Hristakeva, 2022). We implement the methodology separately for each of 133 categories and across the 14 years in our data. We obtain a panel of preference parameters, marginal costs, and markups for individual products across retailers, geographic regions, and periods from 1,862 estimated models.

Our empirical strategy has two key components to support flexible estimation at this scale. First, we construct "micro-moments" from the Consumer Panel Data to identify the parameters that determine variation in preferences across consumers. Second, we apply the covariance restrictions approach of MacKay and Miller (2023) to resolve price endogeneity. Because the covariance restrictions approach requires the joint estimation of supply and demand, the supply-side assumptions can affect the demand estimates. However, these supply-side assumptions are also necessary to recover markups, so we view it as reasonable to employ them in estimation. Further, joint estimation of supply and demand is often used in the literature because it increases efficiency (e.g., Conlon and Gortmaker, 2020).

Our results indicate that the average Lerner index, our measure of markups, increases from approximately 0.45 to 0.60 from 2006 to 2019.[2] We find that the aggregate trend is driven by changes within products over time, rather than consumer substitution toward higher markup products. Larger absolute increases obtain for products with higher initial markups; however, in percentage terms, the changes that we estimate are similar for high- and low-markup products. Thus, the full distribution of product-level markups shifts upward over time.

Within-product markup changes reflect changes in price or marginal cost. We observe that real prices increase by seven percent over 2006 to 2012 but then decrease, such that real prices are only two percent higher in 2019 than in 2006. Thus, marginal cost reductions contribute to rising markups, especially in the later years of the sample. We find that marginal costs decrease at an average annual rate of 2.1 percent. On the demand side, product-specific demand elasticities decrease by 30 percent over the sample period. This is primarily driven by a reduction in consumer price sensitivity, which we define as a model primitive that captures the average

[2] The Lerner index is calculated as $\frac{p-c}{p}$, where p and c are price and marginal cost, respectively (Lerner, 1934). It typically falls between zero to one.

consumer's responsiveness to price.[3] This demand-side result points to a broad change that affects consumers across the income distribution.

We conduct three analyses that point to reductions in marginal cost and consumer price sensitivity as the primary mechanisms that generate rising markups for consumer products. First, in descriptive regressions that isolate cross-category variation in markup trends, we find that changes in marginal cost and consumer price sensitivity have substantially greater explanatory power than other factors, such as changes in consumer demographics and market concentration. Second, we use a counterfactual simulation to compute equilibrium markups that reflect only changes to marginal costs and consumer price sensitivity, while holding all other factors constant. This yields markup trends that closely track our baseline estimates. Finally, we develop a novel econometric decomposition that isolates the net contribution to markups of all observable factors, including prices, market shares, product ownership, market concentration, and consumer purchasing patterns. Although these observable factors explain most of the cross-category variation in markups, they do not contribute much to aggregate markup trends.

We take initial steps to explain declining consumer price sensitivity. In descriptive regressions that exploit cross-category variation in trends, we explore whether changes in price sensitivity reflect changes in retail shopping patterns (e.g., involving warehouse clubs, dollar stores, and online shopping) and differences in firm-level marketing and R&D investments. However, these factors do not not explain the changes in consumer price sensitivity that we estimate. Lower price sensitivity might arise instead from an increase in consumers' opportunity cost of time. In support of this possibility, we provide data that indicate that the use of coupons, which require a small effort by consumers to obtain lower prices, has been declining since the early 1990s and fell by 50 percent over our sample period. In addition, data from the American Time Use Survey indicate that the time that consumers spent shopping on consumer products fell by approximately 20 percent during our sample period.

With declining costs and changing preference parameters, rising markups could coincide with increases or decreases in consumer welfare. In our application, we find that consumer surplus per capita increased despite rising markups. The changes in consumer surplus vary across the income distribution, with more of the gains accruing to higher-income consumers. Even so, rising markups are consequential for consumer surplus. Using a counterfactual simulation, we find that consumer surplus would have been 18 percent higher in 2019 if markups were scaled down to 2006 levels. We use a standard measure of consumer surplus (Small and Rosen, 1981) that applies only to the products in our sample, limiting our ability to address overall consumer welfare. Nonetheless, these analyses indicate that market power can have important impacts on allocations across the income distribution and on aggregate welfare—subjects of longstanding interest (e.g., Harberger, 1954).

[3] Consumer price sensitivity can reflect both the strength of brand-specific preferences and the perceived value of lower prices; in the model, less price sensitive consumers require a greater difference in prices to switch to a less-preferred brand.

Our analysis is subject to limitations. Due to the scale of the empirical exercise, our demand specification is simpler than specifications employed by some studies that focus on a single product category. We conduct a series of comparisons with estimates derived from those studies and various validation checks to show that this limitation is unlikely to drive our estimated markup trends. As one example, for coffee, our estimates of marginal costs move one-for-one with the world commodity price index, and, like Nakamura and Zerom (2010), we find that the commodity price is roughly half of total marginal costs. Overall, our approach yields a reasonable distribution of demand elasticities. Another limitation is that we maintain the assumption of Nash-Bertrand competition throughout. Therefore, we do not identify changes in conduct or how they might affect inferences about markups.[4]

Our research contributes to a growing empirical literature on the evolution of market power. A seminal contribution is that of De Loecker et al. (2020), which applies the production approach to recover markups of firms that are publicly traded in the United States. A central finding is that average firm-level markups increase significantly from 1980 to 2016, primarily due to a reallocation of revenue toward higher-markup firms. Our research is complementary, as we use a different approach to estimate product-level markups from data on prices and quantities, rather than firm-level markups using data from financial reports. We also find that markups increase, albeit due to within-product changes rather than reallocation.

The papers in the literature that are closest to our own are Brand (2021) and Atalay et al. (2023). Both examine markup trends for consumer products using the demand approach under the assumption of Nash-Bertrand competition, as we do. Brand (2021) considers the hypothesis that increases in product variety lead to lower price sensitivity. He estimates demand in nine consumer product categories for the years 2006 and 2017 and finds less elastic demand and higher markups in the latter year. Key distinguishing factors in our analysis include the scope—we consider a broader set of product categories in every year—and our use of individual consumer data to link substitution patterns to variation in demographics. In addition, we address price endogeneity. Atalay et al. (2023) use a sample of 72 product categories from the NielsenIQ Retail Scanner Data and find similar trends in markups and consumer price sensitivity. They apply a different modeling framework that features nested logit demand, and they use Hausman instruments to address endogeneity in the price and nesting parameters.[5]

Our research also adds to a literature that applies the demand approach to recover markups in specific industries over long time horizons. Ganapati (2024) finds that the markups of wholesalers increased over 1992-2012 due to greater scale economies and the expansion of

[4]Although the Nash-Bertrand assumption is widely maintained in the literature, other forms of conduct may also be relevant in consumer product markets (e.g., Miller and Weinberg, 2017).

[5]Berry and Jia (2010) also explore the connection between consumer preferences and markups, in the context of commercial airline markets, and find that an increase in consumer price sensitivity helps explain a (modest) decline in the markups over 1999–2006. This result suggests the caveat that the decreases in price sensitivity that we find for consumer products may not extend throughout the economy. As price sensitivity can reflect the strength of brand preferences, it may increase in some sectors even as it decreases in others.

distribution networks, and that consumers were benefiting from lower prices and access to higher quality goods. Grieco et al. (2023) find that the markups of automobile manufacturers decreased over 1980-2018 due to greater competition, despite dramatic increases in product quality and reductions in marginal costs. Miller et al. (2023) show that technology adoption in the cement industry over 1974-2019 increased markups and reduced marginal costs, with price levels changing only modestly. Consistent with our results, these studies highlight the role of cost savings as a determinant of long run economic outcomes.[6]

The paper proceeds as follows: In Section 2, we discuss the demand approach for recovering markups and specify the models of demand and supply that we use. Section 3 describes the data, the estimator, and our identifying assumptions. Section 4 summarizes results and explores trends in markups, prices, marginal costs, and consumer price sensitivity over time. It also contains descriptive regressions of markups on possible determinants and catalogs a series of robustness exercises that appear in appendices. Section 5 investigates mechanisms with counterfactual simulations, the econometric decomposition, and possible explanations for changing price sensitivity. The analysis of consumer surplus and welfare is in Section 6. Section 7 concludes. Our appendix materials contain a number of additional robustness analyses and validation exercises.

2 Modeling Framework

2.1 The Demand Approach to Recovering Markups

We follow the demand approach to recover markups. This approach is often used when data on prices and quantity are available, and it is a staple of the industrial organization literature. The approach invokes the assumption that firms maximize profits and then recovers an estimate for marginal costs that rationalizes observed prices. Take the case of a single-product firm that sets a price, P, given a residual demand schedule, $Q(P)$, and total costs, $C(Q)$. Differentiating its profit function with respect to price and rearranging yields a first order condition for profit maximization of the form:

$$\frac{P - C'}{P} = -\frac{1}{\varepsilon} \tag{1}$$

where $\varepsilon \equiv \frac{\partial Q(P)}{\partial P}\frac{P}{Q(P)}$ is the price elasticity of demand. The left-hand-side of the equation is the Lerner index, a measure of markups (Lerner, 1934; Elzinga and Mills, 2011). Knowledge of the demand elasticity identifies the Lerner index and the price-over-cost markup (i.e., P/C'). With price data, one also can recover marginal cost and the additive markup (i.e., $P - C'$).

The approach can be employed with more general demand systems that allow for multi-product firms (e.g., Berry, 1994; Berry et al., 1995) and with alternative assumptions about firm

[6] Miller (2024) reviews these industry studies in greater detail. Also related is Peltzman (2022), who analyzes accounting data on manufacturing firms over 1982-2012 and finds support for rising markups and productivity.

behavior, although the most typical equilibrium concepts are Nash-Bertrand and Nash-Cournot competition. The central idea is to use first-order conditions to infer the (unobserved) marginal costs that rationalize observed prices. With a demand system in hand, welfare statistics such as consumer surplus can be calculated, and it also becomes possible to conduct counterfactual simulations for policy evaluation or an exploration of causal mechanisms.

The main alternative is the so-called *production approach* that was pioneered in Hall (1986, 1988, 1990) and De Loecker and Warzynski (2012), and is applied to the evolution of markups in De Loecker et al. (2020) and De Loecker and Eeckhout (2021). Under an assumption of static cost minimization, the multiplicative markup (i.e., P/C') equals the product of (i) the elasticity of output with respect to a variable input and (ii) the ratio of revenue to expenditures on the variable input. Thus, firm-level markups can be recovered by estimating output elasticities and then scaling them with accounting data on revenues and expenditures. The production approach does not require prices or marginal costs to be observed.

As the demand approach and the production approach differ in their assumptions and data requirements, they may best be viewed as complements (e.g., De Loecker, 2011).[7] In applying the demand approach to consumer products, our research builds on the earlier research of De Loecker et al. (2020) and De Loecker and Eeckhout (2021), in that we construct markups at the (much more narrow) level of a product in a specific market. Our estimates are based on observed prices and quantities at this level, instead of firm-level revenue information that aggregates across many products and markets. Implementation comes with its own challenges and limitations. As inferences about markups are linked to the demand elasticities, an identification strategy is needed to obtain consistent estimates of the demand-side parameters in the presence of price endogeneity. Furthermore, the demand-side approach requires the researcher to specify the structure of the demand system and the nature of competition between firms.

2.2 Demand Model

For each product category and each year, we apply the random coefficients logit model of Berry et al. (1995). This demand system incorporates product differentiation and consumer heterogeneity, and it has been widely used in the literature to study consumer products. We work with data that are aggregated to the level of a retail chain, quarter, and geographic region. As in Backus et al. (2021), we assume that each consumer is affiliated with a single retail chain and geographic region, in the sense that they select among the products sold by one chain in their region. Let there be $j = 0, \ldots, J_{crt}$ products available for purchase in chain c, region r, and quarter t, including an outside good ($j = 0$). Affiliated consumers choose among these

[7] One working paper implements both approaches in the context of the U.S. brewing industry, and finds that they deliver similar results (De Loecker and Scott, 2022).

products. The indirect utility that consumer i receives from a purchase of product $j > 0$ is

$$u_{ijcrt} = \beta_i^* + \alpha_i^* p_{jcrt} + \xi_{jr} + \xi_{cr} + \xi_t + \Delta\xi_{jcrt} + \epsilon_{ijcrt} \tag{2}$$

where p_{jcrt} is the retail price, the terms $(\xi_{jr}, \xi_{cr}, \xi_t)$ are product×region, chain×region, and quarter fixed effects, respectively, $\Delta\xi_{jcrt}$ is a structural error term, and ϵ_{ijcrt} is a consumer-specific logit error term. A consumer that selects the outside good receives $u_{i0crt} = \epsilon_{i0crt}$.

We assume that the consumer-specific coefficients, β_i^* and α_i^*, depend on a set of observed and unobserved demographic variables according to

$$\alpha_i^* = \alpha + \Pi_1 D_i \tag{3}$$

$$\beta_i^* = \beta + \Pi_2 D_i + \sigma v_i \tag{4}$$

where D_i contains the observed demographics and $v_i \sim N(0,1)$ contains an unobserved consumer demographic. We demean the observed demographics so that α and β characterize average preferences. In doing so, we use a global mean to preserve variation across regions and time. We restrict the unobserved demographics to affect only the constant, rather than also prices, because we find that separately identifying both effects is difficult in practice. Allowing β to be absorbed by the product fixed effects, the structural parameters to be estimated are $\theta = (\alpha, \Pi_1, \Pi_2, \sigma)$.

Because we estimate the model separately for each category-year, all of the structural parameters and fixed effects are allowed to vary freely by product category and year. We omit subscripts for the year and category for notational brevity.

Quantity demanded is given by $q_{jcrt}(\boldsymbol{p}_{crt}; \theta) = s_{jcrt}(\boldsymbol{p}_{crt}; \theta) M_{crt}$, where $s_{jcrt}(\cdot)$ is the market share, $\boldsymbol{p}_{crt}$ is a vector of prices, and M_{crt} is the "market size" of the chain-region-period. Nevo (2000b) provides equations for market shares. We assume that market size—a measure of potential demand—is proportional to the population of the region and the number of stores operated by the chain within the region. Appendix A provides details on this calculation.

Our specification accommodates vertical differentiation among the inside goods because higher quality (more expensive) products may attract relatively price-insensitive consumers. This can be an important modeling feature in the context of markup trends, especially to the extent that prices or consumer incomes change over time. Relatedly, incorporating heterogeneity in price sensitivity allows the model to more reliably capture the curvature of demand (Miravete et al., 2022, 2023). Our specification also incorporates heterogeneity in the utility that consumers receive from the inside goods, which allows for flexible substitution patterns between the inside and outside goods. However, we do not incorporate product characteristics other than price and the fixed effects. Including more product characteristics would allow for the inclusion of random coefficients on these characteristics and yield more flexible cross-price elasticities. Doing so would be difficult to implement at scale because it would require category-by-category assessments about which characteristics are appropriate to include and whether or not relevant data are available.

2.3 Supply Model

Consumer products are produced by manufacturers and sold through retail chains. We assume that each manufacturer sets prices to maximize its profit, taking as given the prices of its competitors and passive cost-plus pricing on the part of retailers. Thus, the retail markup becomes part of the marginal cost that the manufacturer must pay to sell their products (Gandhi and Nevo, 2021). This assumption is maintained elsewhere (e.g., Miller and Weinberg, 2017; Backus et al., 2021) and is supported by evidence from the empirical literature.[8]

The first order conditions for profit maximization can be expressed in terms of the additive markup:

$$\boldsymbol{p}_{crt} - \boldsymbol{c}_{crt} = -\left(\Omega_{crt} \circ \left[\frac{\partial \boldsymbol{s}_{crt}(\boldsymbol{p}_{crt})}{\partial \boldsymbol{p}_{crt}}\right]'\right)^{-1} \boldsymbol{s}_{crt}(\boldsymbol{p}_{crt}) \tag{5}$$

where the vectors $\boldsymbol{p}_{crt}$, $\boldsymbol{s}_{crt}$, and $\boldsymbol{c}_{crt}$ collect the prices, market shares, and marginal costs of products $j = 1, \ldots, J_{crt}$, and Ω_{crt} is an "ownership matrix" in which each j^{th}, k^{th} element equals one if products j and k are produced by the same manufacturer, and zero otherwise.

We assume that marginal costs are constant in output, following the empirical literature on consumer products (Villas-Boas, 2007; Chevalier et al., 2003; Hendel and Nevo, 2013; Miller and Weinberg, 2017; Backus et al., 2021). We decompose marginal cost according to:

$$c_{jcrt} = \eta_{jr} + \eta_{cr} + \eta_t + \Delta\eta_{jcrt} \tag{6}$$

where $(\eta_{jr}, \eta_{cr}, \eta_t)$ are product×region, chain×region, and quarter fixed effects, and $\Delta\eta_{jcrt}$ is a supply-side structural error term. As in our demand specification, all fixed effects can vary freely by product category and year because we estimate separate models for each category-year combination. Thus, our model allows for changes in brand-specific technologies over time, and, on an annual frequency, these changes may be correlated with changes in demand (e.g., a plant closure). The supply-side structural error term reflects unobserved shocks to costs, including those due to "cost shifters" that have been used elsewhere in the literature as instruments in demand estimation, such as materials costs and distribution costs that affect products and chains differentially.

Given estimates of demand parameters, knowledge of product ownership, and observed values of prices and shares, we can construct the first-order conditions represented by equation

[8] De Loecker and Scott (2022) find evidence for perfect wholesale-retail pass-through indicating competitive retail markets. There is also evidence that retail prices respond to cost shocks (Butters et al., 2022) but not shocks to retailer demand (Arcidiacono et al., 2020). Finally, evidence suggests that retail markups have been relatively stable over the period 1980-2014, despite large changes in demand (Anderson et al., 2023). Our modeling assumptions are also consistent with nonlinear contracts that specify slotting fees or other fixed transfers.

(5) and recover unobserved marginal costs (c_{jcrt}). We use these values to obtain product-level markups, which we calculate as $\frac{p_{jcrt}-c_{jcrt}}{p_{jcrt}}$ (the Lerner index).

3 Data and Empirical Strategy

3.1 Data

Our primary sources of data are the Retail Scanner Data and Consumer Panel Data of Kilts NielsenIQ, which span the years 2006–2019. The scanner data contain unit sales and revenue at the level of the universal product code (UPC), store, and week from a sample of retail chains. The consumer panel data contain the purchases of a sample of panelists by UPC code, retailer, and day, along with demographic information on the panelists. We use aggregation and a number of screens to construct samples that are suitable for the empirical model.

We take as given that consumers choose between the products that are grouped by NielsenIQ into the same product category (or "module"). The categories contain UPCs that are likely substitutes. Our baseline sample comprises 133 product categories that cover 55 percent of revenues in the Retail Scanner Data. We obtain these categories by first identifying the top 200 categories by revenue. These categories account for 74 percent of revenues in the Retail Scanner Data. Within these 200 categories, we apply a screen to select those with relatively modest product differentiation—the ones that the model can reasonably be expected to fit. The screen removes categories for which the 99th percentile of unit prices is greater than five times the median unit price. An example of an omitted category is "Batteries," which has some products that are reasonably close substitutes, such as various brands of AAA batteries, along with other products that are functionally quite different, such as D batteries.

Within these categories, we define products at the brand level, which consolidates thousands of UPC codes into a more manageable set. Brands are defined by NielsenIQ and are fairly narrow. For example, in ready-to-eat cereals, "Cheerios," "Honey Nut Cheerios," and "Multigrain Cheerios" are three distinct brands. Within a brand, we aggregate sales across UPCs by unit of measurement, which characterizes volume (e.g., liters), mass (e.g., ounces), or count (e.g., six-pack), depending on the category.[9] We measure price using the ratio of revenue to equivalent unit sales, following the standard practice to adjust for differences in package size (e.g., Nevo, 2001; Miller and Weinberg, 2017; Backus et al., 2021). Within each category, we treat the top 20 brands by revenue as distinct products, and we collapse the remaining brands into a single composite "fringe" product that we assume is priced by an independent firm. The top 20 brands within each category account for approximately 84 percent of category revenues and typically include a private label product. The average market share of private label products

[9] In a handful of categories, UPCs differ in terms of whether units are reported in terms of volume, mass, or count. For those, we include only UPCs associated with the type of measurement that accounts for the greatest revenue.

Table 1: Sample of Product Categories

Rank	Product Category	Observations	Revenue ($ Millions)	Retailer-DMA Combinations	Brands Per Market	Share Top 20 Brands	Share Private Label
1	Cereal - Ready To Eat	231,178	22,557	333	19.3	0.58	0.08
2	Candy - Chocolate	229,065	16,162	335	18.9	0.54	0.03
3	Candy - Non-Chocolate	225,336	9,420	334	18.6	0.61	0.14
4	Deodorants - Personal	221,618	7,186	333	18.3	0.79	0.00
5	Soap - Specialty	214,153	5,563	355	17.5	0.68	0.05
6	Tooth Cleaners	212,056	7,343	333	17.6	0.71	0.00
7	Shampoo - Liquid/Powder	202,923	7,490	332	16.8	0.65	0.04
8	Cookies	202,880	17,191	334	16.8	0.64	0.18
9	Sanitary Napkins	201,864	5,128	333	16.7	0.79	0.18
10	Cold Remedies - Adult	201,134	9,111	332	16.6	0.85	0.40
20	Bottled Water	160,454	23,333	335	13.2	0.90	0.38
40	Baby Formula	133,082	10,616	323	12.1	0.76	0.05
60	Nuts - Bags	107,314	6,500	334	8.9	0.79	0.24
80	Fresh Muffins	85,228	3,899	332	7.6	0.85	0.17
100	Tuna - Shelf Stable	68,711	4,099	332	5.7	0.98	0.13
120	Cream - Refrigerated	52,297	3,402	330	4.6	0.70	0.30
130	Frozen Poultry	33,428	2,145	300	3.9	0.86	0.27
133	Fresh Mushrooms	25,510	2,772	246	3.4	0.95	0.28
	Mean Values	108,442	6,766	319	9.8	0.84	0.16

Notes: This table shows summary statistics for a selection of product categories. The chosen categories are sorted by the number observations in the estimation sample and are indexed by *rank*. *Revenue* provides total sales in millions of US $ (deflated to 2010 Q1) from 2006 to 2019. The two groups are separated by a horizontal rule. Statistics are calculated after the data cleaning steps described in the text. The last three columns report raw means across retailer-DMA-year-quarter markets. Shares in this table reflect inside shares (i.e., excluding the outside good).

across categories is about 16 percent and it has no meaningful trend over time.[10]

We use the designated market areas (DMAs) in the NielsenIQ data as the geographic regions. We restrict attention to the 22 DMAs for which there are at least 500 panelists in every year in the consumer panel data. These DMAs account for about half of the total revenue observed in the scanner data. Within each DMA, we aggregate the store-level data up to the level of the retail chain, as many retail chains set common prices among nearby stores (DellaVigna and Gentzkow, 2019). Finally, we aggregate the week-level data up to the level of quarters, following Miller and Weinberg (2017). The average number of retail chains per region is 9.3, and the average number of products per category, retail chain, and region is 10.3.[11]

Table 1 provides summary statistics for a selection of product categories in the estimation sample sorted by number of observations. We observe more products in the RTE Cereals category than in any other product category. Among the chains and DMAs in the baseline sample, the RTE Cereals category generated more than $22 billion in revenue over 2006-2019. The category for which we observe the fewest products is Fresh Mushrooms.

We use household demographic data to account for differences in the composition of consumers across markets and changes within markets over time. We generate consumer-specific demographic draws by sampling 5,000 consumers from the Consumer Panel Data for each re-

[10]See Figure A.1 for the distribution of private label market shares across categories and their evolution over time.
[11]For additional details about the data, see Appendix A.

gion and year.[12] We sample with replacement and use the projection weights provided by NielsenIQ. We focus on two demographics that we expect to influence demand for many of the consumer products in the data: household income and an indicator for the presence of children in the household. We assume that log of income is what enters demand through equations (3) and (4). We demean the demographics prior to estimation using a global mean, so that variation over time and across markets remains. We also divide the income measure by its standard deviation. The unobserved demographic is drawn from a standard normal distribution that is independent from the observed demographics.

We also use the Consumer Panel Data to construct "micro-moments" that identify the demographic parameters. Specifically, we calculate the average values of the observed demographics for consumers that purchase each product in a given region and year, again using the projection weights. In estimation, we search for preference parameters that generate corresponding patterns in the model. Thus, our approach attempts to ensure that, for example, the average income of households that purchased Honey Nut Cheerios in Chicago in 2015 matches the data. When constructing the micro-moments, we use only reported purchases from mass merchandisers, grocery stores, and drug stores, and we drop purchases from the largest chains that are not in the Retail Scanner Data. We make these restrictions to match better the distribution of retailers in the Retail Scanner Data, which is disproportionately selected from the aforementioned three channels. Note that we do not restrict our sample of demographic draws based on which retailers the consumers purchase from, so the micro-moments allow us to pin down differences in consumer tendencies to purchase from the retailers in our data (captured by β_i^*).

We account for multi-product ownership using auxiliary data, as ownership information is not provided in the NielsenIQ databases. We start with a manual search in which we identify the company that owns each product. Because multiple company names could be associated with the same manufacturer when a conglomerate has multiple subsidiary companies, we use data from Capital IQ to obtain the ultimate parent company for each product. This process provides a snapshot of product ownership at the end of our sample period. We backcast ownership for the preceding years using information on mergers and acquisitions (M&A) from the Zephyr database, compiled by Bureau van Dijk. Compared with most other M&A databases, Zephyr has the advantage that there is no minimum deal value for a transaction to be included. We assume that prices are chosen to maximize the profit of the ultimate parent company within each category. Finally, we match our sample with firm-level financial data from Compustat to obtain information on marketing expenditures and R&D. We use these variables to explain variation in price sensitivities across brands and time.[13] As Compustat covers only publicly-traded firms, this information is available for about half of the observations in our sample.

[12]We do not sample at the chain-region-year level because the decision to purchase from a chain is endogenous in our model. Our approach does not allow for, e.g., within-region geographic patterns that would shift the demographics of potential shoppers for each chain.

[13]The analyses appear in Appendix F.3.

We deflate prices and incomes using the Consumer Price Index such that they are in real dollars as of the first quarter of 2010.[14]

3.2 Estimation and Identification

We estimate the model using the generalized method of moments (GMM). We estimate separate models for each category and year, and allow the parameters for estimation, $\theta = (\alpha, \Pi_1, \Pi_2, \sigma)$, to vary arbitrarily across the models. The GMM estimator for θ is:

$$\hat{\theta} = \arg\min_{\theta} g(\theta)' W g(\theta), \qquad g(\theta) = \begin{bmatrix} g^{MM}(\theta) \\ g^{CR}(\theta) \end{bmatrix} \tag{7}$$

where W is a weighting matrix, $g^{MM}(\theta)$ collects micro-moments that summarizes how well the model matches the correlations between demographics and product purchases that we observe in the NielsenIQ Panelist dataset, and $g^{CR}(\theta)$ assesses a covariance restriction between demand-side and cost-side structural error terms. We take a two-step approach to estimation in which we first estimate $\theta_2 = (\Pi_1, \Pi_2, \sigma)$ then estimate the price parameter, α (see Appendix B). This reflects that micro-moments identify θ_2 but not α (Berry et al., 2004; Berry and Haile, 2022), and that the covariance restriction identifies α conditional on θ_2 (MacKay and Miller, 2023).

For the micro-moments, we use variation in purchase patterns across products and regions to capture heterogeneity in preferences. Each element corresponding to product j and demographic k is given by

$$g_{jk}^{MM}(\theta) = \frac{1}{T_j} \sum_{c,r,t} \left(\frac{\sum_i \omega_i s_{ijcrt}(\theta) D_{ik}}{\sum_i \omega_i s_{jcrt}(\theta)} - \mathcal{M}_{jrk} \right) \tag{8}$$

where T_j is the number of chain-region-quarter combinations in which product j is sold, ω_i is the weight that we place on consumer i, $s_{ijcrt}(\theta)$ is the consumer-specific choice probability implied by the candidate parameter vector, and $\mathcal{M}_{jrk}$ is the mean demographic observed in the data for product and region. That is, we match the implied average demographic of consumers for each product-chain-region-quarter to the average demographic observed in the data for the corresponding product-region (allowing for differences across years and categories). In our baseline specification, we use two observed demographic variables and at most 21 products, so there can be up to 42 micro-moments.

The covariance restriction is that the demand-side and supply-side structural error terms are uncorrelated in expectation: $\mathbb{E}[\Delta\xi_{jcrt}\Delta\eta_{jcrt}] = 0$. This provides a moment that we use in estimation to identify the price parameter. We construct the empirical analog of the moment as

[14]We deflate using the Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average. This CPI measure is predominantly constructed from products and services outside of the categories in our sample. The inflation data are monthly and seasonally adjusted.

follows:

$$g^{CR}(\theta) = \frac{1}{T} \sum_{c,r,t} \Delta \boldsymbol{\xi}_{crt}(\theta)' \Delta \boldsymbol{\eta}_{crt}(\theta) \tag{9}$$

where the $\Delta \boldsymbol{\xi}_{crt}(\theta)$ and $\Delta \boldsymbol{\eta}_{crt}(\theta)$ terms are recovered for each candidate θ, and T is the number of chain-region-quarter-product combinations for a given year.

Our identification strategy is well-suited for estimation at scale. The primary alternative would be an instrumental variables approach. However, many of the instruments used in the literature require data on non-price product characteristics to construct markup-shifters (e.g., Berry et al., 1995; Gandhi and Houde, 2020) or cost-shifters (e.g., Backus et al., 2021). Furthermore, instruments require a "first-stage" relevance conditions to be satisfied, and there is no guarantee that the instruments would be relevant for every category. By contrast, the micro-moments and covariance restriction do not require data on non-price product characteristics, and there is no equivalent relevance condition that must be satisfied.

As we have specified the model, the supply-side structural error term ($\Delta\eta_{jcrt}$) reflects all residual variation in costs, including factors that correspond to observed instruments. For example, variation in $\Delta\eta_{jcrt}$ will capture product-specific shipping costs (Miller and Weinberg, 2017) and prices of product-specific ingredients (Backus et al., 2021), even though these are not observed directly in our approach. When most of the variation in $\Delta\eta_{jcrt}$ is due to such factors, similar economic reasoning can be used to justify the assumption $\mathbb{E}[\Delta\xi_{jcrt}\Delta\eta_{jcrt}] = 0$ as one would use for $\mathbb{E}[\Delta\xi_{jcrt}z_{jcrt}] = 0$, where z_{jcrt} is a vector of cost-shifter instruments.[15]

We absorb potentially confounding sources of variation by incorporating fixed effects at the product×region, chain×region, and quarter levels (all of which are estimated separately by year and category). For example, the fixed effects control for the possibility that higher-quality products are more expensive to produce. Nonetheless, some threats to validity remain. As a leading example, if firms' marginal costs vary with output, then demand shocks would shift marginal costs, and the covariance restriction would be violated. Correlated shifts in quality and marginal cost that occur within a year also could be problematic. Alternatively, if retailers run non-price promotions when they have excess product inventory, this may boost demand for unobservable reasons, precisely when opportunity costs are low.

We find that our model and identification strategy deliver a reasonable distribution of own-price elasticities, which we report in Appendix Figure D.2. This is an indicator of the potential of our empirical strategy, as demand elasticities are theoretically linked to markups and marginal costs in our model. We also find that our demand estimates are consistent with the literature, for a few categories where comparisons are possible (see Appendix C).

We consider the robustness of our results to model specification along a few important dimensions. First, we consider the role of the consumer heterogeneity pararameters in our model, which are identified by micro-moments constructed from the Consumer Panel data. In

[15] See MacKay and Miller (2023) for a more detailed discussion and additional examples.

Appendix D.1, we estimate an alternative standard logit demand model ($\Pi_1 = 0$, $\Pi_2 = 0$, $\sigma = 0$) for all categories and years. Relative to the logit specification, our baseline estimates obtain more elastic demand estimates and smaller markups.

Next, we consider the omission of product characteristics from the model. First, as noted above, we obtain elasticities and margins that are similar to the literature in other categories. This provides some evidence that product characteristics may not be critical for pinning down markups.[16] A particular example is Backus et al. (2021), who use similar data and assumptions in their analysis of ready-to-eat cereal but include product characteristics. To address this further, we construct product characteristics for ready-to-eat cereals following the approach of Backus et al. (2021). We include these characteristics in an alternative set of estimates as a specification check, and we report the results in Appendix D.2. Our estimates for this category are not materially affected by the inclusion of additional characteristics.

Finally, we consider alternative identification strategies for the price parameter. Relative to the covariance restriction approach, using Hausman instruments or assuming that prices are exogenous yield results that are less plausible. With the covariance restriction, demand is always downward-sloping, and the median own-price elasticities are larger than one in magnitude for 95 percent of the category-year combinations. Thus, in our baseline model, demand is elastic at most observed prices, consistent with profit maximization by firms with positive marginal costs. By contrast, assuming prices are exogenous or using Hausman instruments produces median own-price elasticities greater than one for only 71 and 66 percent of the category-year combinations, respectively, and demand is upward-sloping in several cases.

4 The Evolution of Markups in Consumer Products

In this section, we document the evolution of markups over the sample period. We first summarize median category-level markups and the overall distribution of markups. We then disentangle within-product variation from across-product variation, and decompose markup changes into price and marginal cost changes. We also consider how demand for consumer products changes over time, and examine the economic forces that correlate with markup changes.

4.1 Aggregate Markup Trends

Our estimation procedure yields a panel of 14.4 million product-level observations across 133 categories and 14 years. To evaluate aggregate trends, we first consider changes in the category-level markups in the 1,862 category-year combinations in our data. We take the median markup within each category-year, and we then calculate the mean across categories in each year. Figure

[16] Incorporating product characteristics can allow for richer patterns of horizontal differentiation, which can yield more accurate counterfactual simulations that depend on specific cross-price elasticities, such as merger simulation (e.g., Nevo, 2000a) or studies of entry and exit (e.g., Ciliberto et al., 2021).

Figure 1: Markups Over Time Across Product Categories

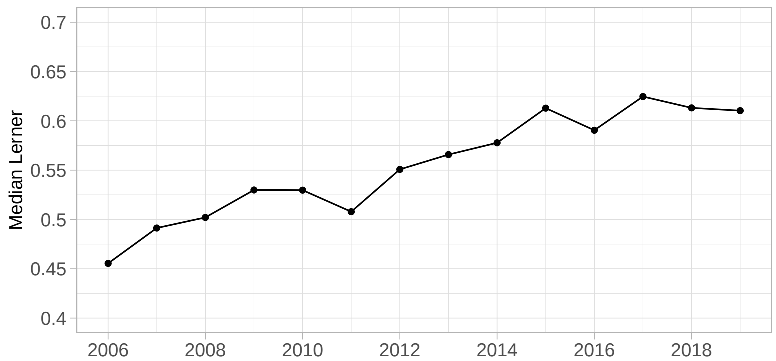

Notes: This figure plots the mean of within-category median markups over time. Markups are defined by the Lerner index, $(p - mc)/p$, and are estimated separately by product category and year. When calculating the mean, we winsorize the upper and lower 2.5 percent of observations across all categories and years.

1 plots this statistic over time. Averaging across categories, we find an increase in the median Lerner index from approximately 0.45 in 2006 to over 0.60 towards the end of our sample period. This corresponds to an average annual growth rate in markups of 2.3 percent.

In addition to the median, we find that the full distribution of within-category markups is shifting upward over time. Appendix Figure G.1 reports the trends for different percentiles of the markup distribution. We regress different percentiles of the within-category markup distribution on year dummies and report the resulting coefficients in panel (a). These can be interpreted as changes in markups relative to 2006, the first year of our sample. In panel (b), we repeat the exercise with the log of the Lerner index, $\ln(\frac{p-c}{p})$.[17] Panel (a) shows that the level increase in markups is slightly higher for the upper half of the distribution, while panel (b) indicates that the relative increase in markups is in fact quite similar across the distribution. These results suggest that the aggregate trends in markups in our sample are not isolated to a subset of products or firms, but rather reflect a broader trend affecting overall categories.

4.2 Within-Product Changes in Markups, Prices, and Marginal Costs

Different economic effects could account for an upward shift in the distribution of markups. Firms could be charging higher markups on existing products, or replacing lower-markup products with newer, higher-markup products. Further, an increasing markup distribution could reflect higher prices, lower marginal costs, or both, or it could reflect a reallocation of market share from lower-markup products to higher-markup products.

To evaluate these possibilities, we analyze the change in markups, prices, and marginal costs at the product level. For markups, we regress the log of the Lerner index on quarter, year, and

[17] We find similar changes in the distribution of firm-level markups which we calculate as quantity-weighted averages over brands owned by each parent company.

Figure 2: Product-Level Changes in Markups, Prices, and Marginal Costs

(a) Markups

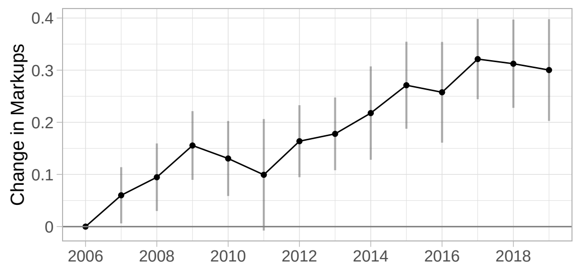

(b) Prices

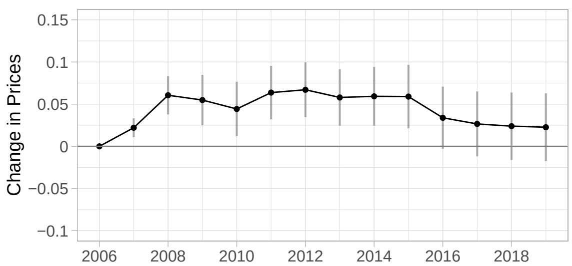

(c) Marginal Costs

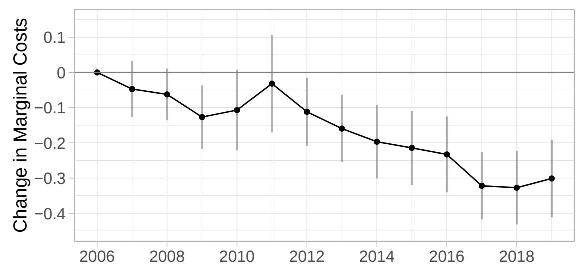

Notes: This figure shows coefficients and 95 percent confidence intervals of a regressions of the log of the Lerner index, real prices, and real marginal costs at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.

product-chain-region fixed effects, using revenues as weights.[18] The unit of observation is a unique product-chain-region-quarter-year combination. The results of this regression are documented in panel (a) of Figure 2. The figure displays point estimates and 95 percent confidence intervals for the year fixed effects. The results indicate a within-product increase in markups of about 30 percent between 2006 and 2019. The estimated annual growth rate of within-product markups is 2.2 percent per year. Thus, we find that the shift in the distribution of markups can be explained by changes within products over time.

We find similar results if we use unweighted regressions or if we use a balanced panel of products. Together, these results indicate that the estimated trend is not driven by product entry and exit or a reallocation of market shares across products. We report the full regression results and these alternative specifications in Tables G.1, G.2, and G.3 in the Appendix. We also consider alternative specifications in which we replace year fixed effects with a linear time trend, drop product-chain-DMA fixed effects, or use category fixed effects. We obtain qualitatively similar results across these specifications, and estimate average yearly increases in markups between 1.7 and 2.2 percent. We estimate larger changes when controlling for product-level fixed effects, indicating that within-product changes are greater than the aggre-

[18]We weight by revenues instead of quantities to assign higher weights to products with higher initial prices. Revenue-weighted relative changes, which we measure by changes in log markups, are consistent with quantity-weighted absolute changes in a consumption basket.

gate (revenue-weighted) changes in markups. Though these differences are not substantial, they suggest that some of the product-level increase in markups may be offset by the introduction of lower-markup products over time.

Mechanically, within-product markup changes must be driven by price changes, marginal cost changes, or both. To explore, we regress log prices and log marginal costs on product-DMA-retailer, quarter, and year fixed effects. Prices and marginal costs are deflated by core CPI and indexed to Q1 of 2010. The yearly coefficients are documented in panels (b) and (c) of Figure 2. Panel (b) shows that prices increase at the beginning of our sample period, but decline in later years. The average price for products in our sample increases by 7 percent over 2006 to 2012, but prices are only 2 percent higher in 2019 than in 2006.

Panel (c) of the figure reports the yearly coefficients for log marginal costs. The results indicate a within-product decrease in marginal costs of about 30 percent between 2006 and 2019. This corresponds to a decline of 2.1 percent per year on average. Our point estimates suggest that marginal costs increased between 2010–2011, though there is some noise and a linear trend cannot be rejected. Thus, higher real prices account for part of the increase in markups during the first half of our sample, while the higher markups we observe at the end of our sample arise primarily from lower real marginal costs. In nominal (i.e., non-deflated) terms, marginal costs are roughly constant over time. See Appendix Figure G.2.

The extent to which marginal cost reductions lead to higher markups, in a causal sense, depends on equilibrium pass-through relationships and cannot be ascertained solely from the trends shown in Figure 2. For example, perhaps prices would have fallen with marginal costs, if not for changes in demand or other factors such as concentration. This motivates a number of empirical exercises later in the paper that investigate causal mechanisms.

However, the marginal cost reductions that we estimate are economically significant on their own. In many markets, costs may decline over time due to innovations in production/distribution technology and operational efficiencies.[19] We observe that, for consumer products, manufacturers sought ways to reduce costs over the sample period. Procter & Gamble, one of the largest companies in our data, began a "productivity and cost savings plan" in 2012 that was estimated to reduce annual costs by $3.6 billion in 2019. Similarly, Unilever reports realizing €6 billion in cost savings over 2017-2019.[20] Overall, our finding of declines in marginal costs is consistent with secular increases in productivity across the economy.

4.3 Changes in Demand

We now examine the changes in demand that are captured by our empirical model. We start with own-price elasticities of demand because these are linked to markups theoretically. We

[19] For one comparison, Grieco et al. (2023) estimate that marginal costs of automobile manufacturers decrease by 1.4 percent per year on average over 1980-2018, conditional on vehicle attributes.

[20] See the 2019 Annual Reports of Procter & Gamble Company and Unilever.

regress the logarithm of the absolute value of own-price elasticity at the product level on quarter, year, and product-chain-region fixed effects, as before. We present the results in panel (a) of Figure 3. The coefficients show that price elasticities decline in magnitude, indicating that demand becomes less responsive to prices over time. Price elasticities can be influenced by the level of prices, several underlying aspects of consumer preferences, and supply-side factors such as quality and competition that contribute to the shape of residual demand curves.

We find that the time-series variation in price elasticities is highly correlated with the mean price coefficient, α, which we estimate separately for each category and year. We repeat the regression exercise using price sensitivity, which we define as the log absolute value of the mean price coefficient (i.e., $\log(-\alpha)$), as the dependent variable. Panel (b) shows a decline in price sensitivity. The decreases are large through 2012, corresponding with the initial increase in prices that we observe. In the model, changes to α reflect changes that are common to all consumers, as the model controls for factors such as changing consumer demographics and selection by consumers into retailers and products.

While changes in the elasticity may be driven by vertical or horizontal shifts in demand or movement along the demand curve, the changes that are captured by α reflect a rotation in the demand curve. A demand rotation that decreases price sensitivity manifests in the data as a reduced within-market ratio of the variance in quantities over the variance in prices. We observe that this ratio of variances is falling, driven primarily by a decrease in within-market share dispersion (while within-market price dispersion increases slightly). Intuitively, this pattern indicates that demand is becoming less sensitive to price variation. The covariance restriction approach to estimation exploits this empirical pattern to identify changes in price sensitivity. See MacKay and Miller (2023) for additional details.[21]

Our estimates allow us to examine other changes in demand as well. For instance, the fixed effects allow us to characterize changes in perceived product quality over time, relative to that provided by the outside good. We measure product quality as the value that an average consumer obtains from the product (relative to the outside good); to improve comparability across categories we standardize values using the category-level means and standard deviations. Figure G.3 in the Appendix shows that perceived product quality declines over the sample period.[22] Improvements in the outside good—which includes shopping through online retailers for example—could contribute to this trend. The same appendix figure plots changes in the coefficients that characterize how observed consumer demographics affect the consumer-specific price coefficient and category-level constant (Π_1, Π_2). Changes in these parameters can also affect substitution to the outside good, relative to substitution among products within a

[21]This change can be also captured by alternative approaches, even those that are known to deliver biased elasticities. A rotation in demand affects the shape of the price-quantity cloud, which will, for example, affect even a linear ordinary least squares coefficient. In Appendix D.4, we examine trends under the alternative identification assumption that prices are exogenous.

[22]A similar but slightly smaller trend is obtained when we scale quality by the estimated price coefficients.

Figure 3: Changes in Demand

(a) Log Absolute Elasticity (b) Price Sensitivity

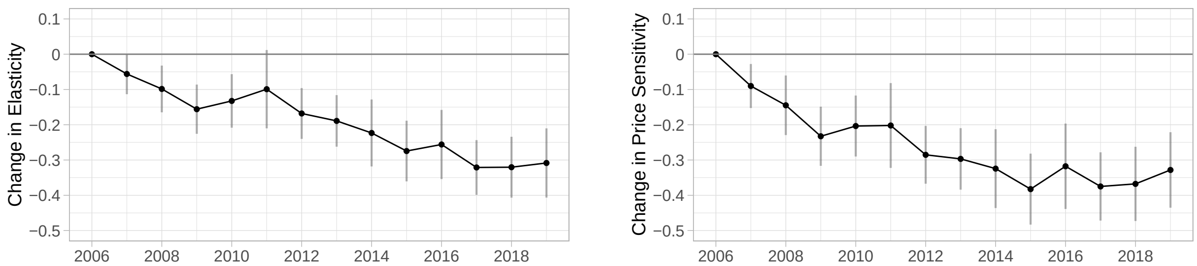

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of log absolute elasticity and price sensitivity at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.

market. In Appendix Figure G.4, we report trends in market elasticities and diversion ratios.

4.4 Descriptive Analysis

Thus far, we have shown that within-product markups increase over the sample period and that, from a mechanical perspective, the change is primarily driven by marginal cost decreases, at least by the end of the sample. We also have shown that consumer price sensitivity and perceived product quality (relative to the outside good) decrease over the sample period. In our model, each of these factors—marginal cost, price sensitivity, and product quality—affect equilibrium markups. Consumer demographics and competitive events, such as mergers and product entry, also matter for equilibrium markups.

In this section, we take initial steps to identify why markups increase, exploiting the unique panel structure of our data. Our approach uses descriptive regressions to identify how changes in different market factors correlate with changes in markups. We use fixed effects to absorb the straight-forward time-series variation that our figures have highlighted previously.[23] Thus, the question we ask is a relative one: Do products with greater markup increases also have larger marginal cost reductions, bigger decreases in consumer price sensitivity, and so on? One might expect that the factors that explain markup growth may also have important causal effects, although causality cannot be established from descriptive regressions alone. We use the analysis to inform the more formal exploration of mechanisms that we conduct in the next section.

Specifically, we regress product-level log markups on marginal costs, consumer preference parameters, consumer demographics, and measures of market concentration. We incorporate

[23]These figures indicate that, in the time-series, rising markups correlate with lower marginal costs, lower consumer price sensitivity, and lower product quality. We do not observe similar correlations between markups, income, and market concentration. In our sample period, real income first decreases, then increases. Similarly, the trends in the concentration measures shown in Appendix Figure G.5 do not cleanly track the markup trends.

Table 2: Factors Predicting Cross-Category Variation in Markup Trends

	(1)	(2)	(3)	(4)	(5)	(6)
Marginal Cost (Standardized)	-0.585^{***}					-0.461^{***}
	(0.020)					(0.021)
Price Sensitivity		-0.728^{***}				-0.397^{***}
		(0.025)				(0.022)
Quality (Standardized)			-0.137^{***}			0.002
			(0.021)			(0.006)
Income (Log)				0.101^{***}		0.061^{***}
				(0.029)		(0.013)
Children at Home				-0.114^{*}		-0.017
				(0.065)		(0.049)
Parent HHI					0.323	0.221^{***}
					(0.206)	(0.060)
Brand HHI					-0.004	-0.102^{**}
					(0.186)	(0.051)
Retailer HHI					0.192^{***}	0.069^{**}
					(0.068)	(0.027)
Brand-Category-DMA-Retailer FEs	X	X	X	X	X	X
Time Period FEs	X	X	X	X	X	X
Observations	14,406,731	14,406,731	14,406,731	14,406,731	14,406,674	14,406,674
R^2 (Within)	0.715	0.476	0.045	0.000	0.003	0.825

Notes: This table reports regression results where the dependent variable is log markups. Observations are at the brand-category-DMA-retailer-year-quarter level, and brand-category-DMA-retailer and year-quarter fixed effects are included in each specification. Standard errors are clustered at the category level and are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

fixed effects at the brand-category-DMA-retailer level. We also include year fixed effects, as explained above. Thus, the regression coefficients capture deviations from aggregate trends. Although the regression coefficients are of interest, our main focus is on the ability of the regressors to explain changes in markups, as reflected by their contribution to the R^2.

For the consumer preference parameters, we include price sensitivity and perceived product quality, as defined in the previous section. We standardize the product qualities and marginal costs, separately for each category, so that they have a variance of one.[24] For consumer demographics, we use log income and the presence of young children at home. Finally, for market concentration, we examine three constructions of the Herfindahl-Hirschman Index (HHI). Parent HHI is calculated for the upstream parent companies of the products (i.e. for the brand manufacturers). Brand HHI is calculated under the counterfactual of single-product firms, and serves to isolate changes in market concentration that are unrelated to product ownership. Finally, Retailer HHI is calculated for the retailers, separately for each category and region. We measure the HHIs on a zero-to-one scale.[25]

[24] Standardization improves comparability across categories and also eases interpretation of the coefficients. We choose this approach to standardization, rather than logs, to include observations with negative values.

[25] We use the consumer panel data to construct HHI measures. Our results are qualitatively similar if we instead use the retail scanner data.

Table 2 summarizes the results. Column (1) indicates that marginal cost reductions alone can explain 72 percent of the within-product variation in markups (within $R^2 = 0.715$). The coefficient implies that a one standard deviation reduction in marginal costs is associated with a 59 percent increase in markups. Similarly, column (2) indicates that declines in price sensitivity alone can explain 48 percent of the within-product variation in markups (within $R^2 = 0.476$). The coefficient indicates that a 10 percent decrease in price sensitivity is associated with a 7.3 percent increase in markups.[26]

Columns (3), (4), and (5) examine perceived quality, consumer demographics, and concentration. Column (4), for example, indicates that markups increase more in regions with more income growth. Although some of the coefficients are statistically significant, these measures explain little of the variation in log markups. The within R^2 values are all less than 0.05.

In column (6), we include all of the regressors. The coefficients on price sensitivity and marginal costs decline modestly, but remain large in magnitude and statistically significant. The coefficient on quality becomes effectively zero. Thus, though declines in relative perceived quality are correlated with increasing markups in the time series, products with greater increases in quality do not realize differential changes in markups. The coefficients on the demographic variables also shrink, though income remains statistically significant.

Turning to concentration, we find that parent HHI and retailer HHI are are positively correlated with changes in markups. Yet, these coefficients remain modest. The parent-HHI coefficient of 0.221 in column (6) indicates that a change in parent company HHI of 0.02—i.e., a 200-point change on a 0 to 10,000 scale—is associated with a 0.4 percent increase in markups. The relationship between markups and changes in concentration at the brand level (which ignores multi-product ownership) is weaker.

The within R^2 for the multivariate regression is 0.83. Thus, the regressors explain the bulk of within-product markup deviations from the aggregate trends. Among these regressors, the two that contribute the most are marginal costs and price sensitivity; a specification with these two together yields a within R^2 of 0.82. This helps focus our remaining efforts to understand underlying mechanisms, to which we turn in Section 5.

4.5 Sensitivity Analyses and Robustness Checks

In this section, we describe several additional robustness checks and alternative specifications to examine the sensitivity of our main results. In each of the analyses, the estimated trends in markups and price sensitivity track our baseline estimates. We provide additional details and present the results in Appendix E.

[26]Note that price sensitivity is measured at the category-year level, whereas markups and marginal costs may vary across brands, DMAs, and retailers within each category-year. If we run regressions at the product category level, we find similar coefficients and a higher within R^2 for price sensitivity. We report these results in Table G.4 in the Appendix.

(1) Markup measure: Our baseline measure of markups is the Lerner index, $\frac{p-c}{p}$. We find very similar results if we instead use price-over-cost markups, p/c, as used by, e.g., De Loecker et al. (2020). We report the alternative markup trend in Appendix E.1.

(2) Product and retailer composition: To show that our trends are not due to changes in the composition of products and retailers within the sample, we re-estimate the model using a balanced panel of brand-chain-DMA combinations. The results are presented in Appendix E.2.

(3) Sample of categories: To obtain our baseline estimates, we use a sample of 133 categories that feature modest degrees of price dispersion. We present our estimated markup and price sensitivity trends using a broader sample of the top 200 categories by revenue. The results are in Appendix E.3.

(4) Sample of retailers: Our baseline sample is constructed from the Retail Scanner Data, which comes from a sample of retail chains. One possible concern is that consumers select into the chains in our sample and that the selection has changed over time. To address this concern, we construct price and quantity data from the Consumer Panel Data for large retailers that are not included in the Retail Scanner Data, and we estimate the model including these additional retailers. We report the results in Appendix E.4.

(5) Sample of geographic markets: Our baseline sample includes the 22 DMAs with at least 500 households in the Consumer Panel data for each year over 2006–2019. We re-estimate the model using an expanded sample of 30 DMAs and a restricted sample of the 15 largest DMAs. The set of 30 DMAs includes all DMAs with at least 500 households in the Consumer Panel data for 2007–2019. We present the results in Appendix E.5.

(6) Number of brands: Our baseline analysis includes the top 20 brands in each category, and it aggregates the rest to a synthetic "fringe" brand. One question is whether our results depend on the number of brands aggregated into the fringe. As a robustness check, we estimate our model instead using the top 15 brands, and we present the results in Appendix E.6.

(7) Time aggregation: Our baseline analysis aggregates the data to the quarterly level. Time aggregation involves a trade-off between the number of observations that are used to identify demand parameters and the sensitivity to short-run fluctuations induced, for instance, by temporary sales. To check the robustness of our results to the level of time aggregation, we include an alternative specification that aggregates to the semiannual level, instead of the quarterly level. We present these results in Appendix E.7.

(8) Market definition: We assume that each consumer shops within a retailer-DMA-year-quarter-category. Thus, our model does not allow consumers to shop across categories or across retailers within a DMA. Accounting for cross-category substitution by consumers goes beyond the scope of our analysis. As an indicator for the potential impact of cross-retailer substitution on our estimates, we consider an alternative specification where the market is defined at the level of a DMA-year-quarter-category. We estimate this specification by first aggregating the data across retailers within a DMA. We report the resulting estimates in Appendix E.8.

(9) Market size: A larger market size implies higher diversion to the outside good and lower diversion among the inside goods. Some demand systems, such as logit, place restrictions on diversion such that the estimates of consumer substitution are sensitive to the assumption on market size. Our random coefficients demand model can flexibly capture substitution between the inside goods and the outside good, so the market size assumption may be less consequential. Still, we re-estimate the model using different assumptions on market size. We present the results of these robustness analyses in Appendix E.9.

5 Mechanisms

The previous section presented our findings of rising markups and provided descriptive analyses that indicate a high degree of correlation between rising markups, marginal cost reductions, and decreasing consumer price sensitivity. Here, we use counterfactual simulations and a novel econometric decomposition to better understand whether these relationships are causal. We also provide additional descriptive evidence about the declines in price sensitivity.

5.1 Counterfactual Simulations

We use counterfactual simulations to assess the causal effects of aggregate trends in marginal costs and price sensitivity on markups. We start with the estimated model and data from 2006. We then compute equilibrium prices over 2007-2019 under simulated changes in marginal costs and price sensitivity, holding fixed product assortments, consumer demographics, and other demand parameters. First, we scale marginal costs uniformly based on the estimated trend from Figure 2, holding price sensitivity fixed at 2006 levels. Second, we scale price sensitivity uniformly based on the estimated trend from Figure 3, holding marginal costs fixed at 2006 levels. Third, we scale both marginal costs and price sensitivity. We obtain hypothetical time series that inform the extent to which aggregate trends contribute to rising markups.

For reporting purposes, we regress the log of the Lerner indices that we obtain under each of these counterfactual scenarios on quarter, year, and product-chain-region fixed effects, weighting by revenue. This follows the approach we use in Sections 4.2 and 4.3. The year fixed effects then summarize within-product markup changes over time.

Figure 4 depicts the results. The dash-dotted line plots changes in markups under the assumption that marginal costs track the aggregate trend but that demand and other factors are fixed. We find that marginal cost reductions alone yield an increase in markups of about 13 percent over the sample period. The dashed line plots changes in markups under the assumption that price sensitivity tracks the aggregate trend but that other factors are fixed. Declines in

Figure 4: Simulated Markup Changes

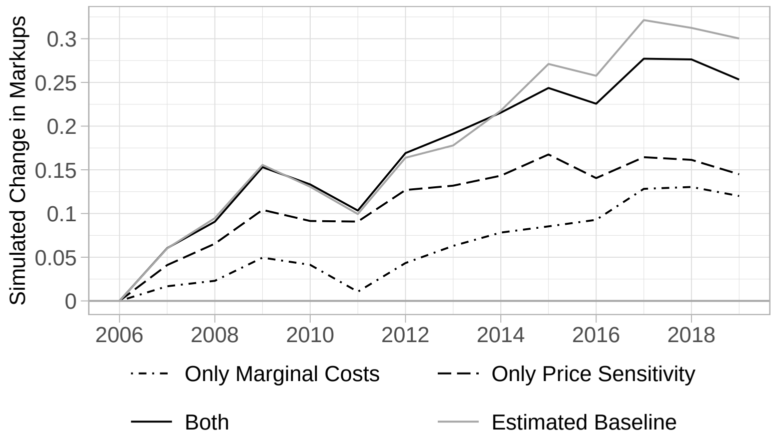

Notes: This figure plots counterfactual log changes in markups from simulations that scale marginal costs (dash-dotted line), price sensitivities (dashed line), or both (solid line) according to the average realized changes that are reported in Figures 2 and 3. Markups are defined by the Lerner index, $(p - mc)/p$, and changes are reported relative to 2006. Product assortments, consumer demographics, and other demand parameters are held fixed at 2006 values in each simulated year. The solid gray line plots the estimated change in log markups in the realized data for comparison.

price sensitivity alone yield in an increase in markups by about 15 percent. The solid black line plots markup changes when both marginal costs and price sensitivity track aggregate trends, while everything is held fixed. In that scenario, markups increase by about 29 percent (roughly 0.25 log points), and the overall trend is similar to the estimated changes in markups reported in Figure 2, which we plot here with the solid gray line.

These counterfactual simulations indicate that aggregate changes in marginal costs and price sensitivity can account for the vast majority of the time-series variation in markups. Marginal cost reductions lead to higher markups because, in our estimated model, there is incomplete pass-through of costs to equilibrium prices. A reduction in price sensitivity leads to firms choosing higher prices and markups, all else equal. These factors have different relative impacts over time. For the first half of our sample, increases in markups can be mainly attributed to changes in price sensitivity, while falling marginal costs are the primary driver of rising markups in the second half. Together, the two trends explain approximately 85 percent of the aggregate increase in markups by 2019.

The black line in Figure 4 nearly overlaps the estimated markup trend from 2006 through 2014. From 2014 to 2019, both the counterfactual and the estimated markups continue to increase, but the lines diverge slightly. Several factors could explain these differences, including factors we hold fixed (e.g., product assortment, demographics), as well as the fact that changes in marginal costs and price sensitivity were not uniform across products. A counterfactual that instead imposes category-level mean changes in marginal costs and price sensitivity yields larger changes in markups that can fully explain the estimated change by 2019.

5.2 Econometric Decomposition

We now develop a novel econometric decomposition that provides an alternative approach to assessing the impact of price sensitivity. The thought experiment involves holding fixed the data, and considering the implications of trends in the data for trends in markups, when interpreted through the lens of the model. Thus, it is conceptually distinct from the counterfactual simulations of the previous section, which take as primitives the demand and cost parameters.

Our starting point is the result of MacKay and Miller (2023) that product-level additive markups can be expressed in terms of the mean price parameter α and a known function of data, $\lambda(\cdot)$, for a broad class of oligopoly models. With random coefficients logit demand and Bertrand pricing, this takes the form:

$$p_{jcrt} - c_{jcrt} = -\frac{1}{\alpha}\lambda_{jcrt}(\boldsymbol{s}_{crt}, \boldsymbol{p}_{crt}, \boldsymbol{x}_{crt}, \boldsymbol{D}_{rt}, \boldsymbol{\nu}_{rt}, \boldsymbol{\Omega}_{rt}; \Pi_1, \Pi_2, \sigma), \tag{10}$$

The arguments in $\lambda(\cdot)$ include vectors of market shares and prices ($\boldsymbol{s}_{crt}$ and $\boldsymbol{p}_{crt}$). It also includes a matrix, $\boldsymbol{x}_{crt}$, of non-price data that is interacted with the parameters Π_1, Π_2, and σ; in our implementation, $\boldsymbol{x}_{crt}$ is a vector of ones. Finally, $\boldsymbol{D}_{rt}$ and $\boldsymbol{\nu}_{rt}$ contain consumer demographics, and $\boldsymbol{\Omega}_{rt}$ summarizes product ownership. Though $\lambda(\cdot)$ depends on the parameters Π_1, Π_2, and σ, we identify these parameters by matching micro-moments and therefore treat them as data for this exercise. Importantly, α has no effect on $\lambda(\cdot)$ conditional on the data.

Taking the quantity-weighted average within each category and year and dividing by average price, we obtain an expression for the aggregate Lerner index,

$$\overline{L} = \frac{\overline{p} - \overline{c}}{\overline{p}} = -\frac{1}{\alpha}\frac{\overline{\lambda}}{\overline{p}} \tag{11}$$

In logs, we obtain:

$$\ln \overline{L} = \underbrace{-\ln(-\alpha)}_{-1\times\text{Price Sensitivity}} + \underbrace{\ln\left(\frac{\overline{\lambda}}{\overline{p}}\right)}_{\text{Observable Factors}} \tag{12}$$

This decomposes (log) category markups into price sensitivity (i.e., $\ln(-\alpha)$) and a term that captures the net effect of observable factors: prices and market shares, product characteristics, consumer demographics, consumer purchasing patterns, and the ownership of products (e.g., market concentration). By contrast, α is not observed and requires an additional identifying assumption. Thus, the econometric decomposition allows us to assess the role of α in determining markups, holding fixed the model specification and the observable factors in the data.

Figure 5 plots changes in the aggregate log Lerner index ($\ln \overline{L}$) and the contribution of observable factors ($\ln\left(\overline{\lambda}/\overline{p}\right)$) over the sample period, relative to 2006. Consistent with our earlier findings, the average log Lerner index increases by 0.25 from 2006 to 2019. However, the structural component decreases from 2006 to 2011 and remains below 2006 levels thereafter.

Figure 5: Econometric Decomposition of Markup Trends

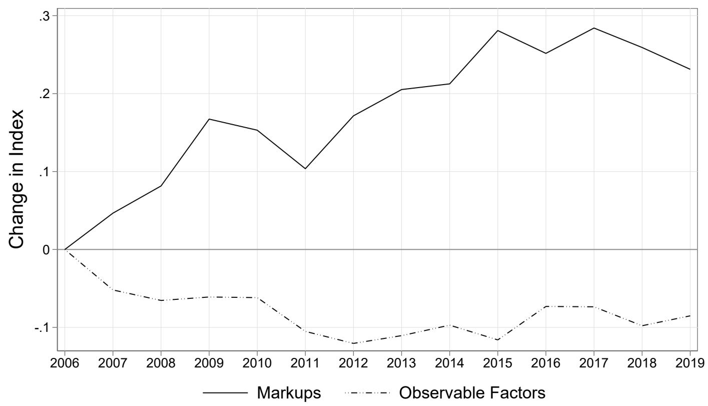

Notes: This figure shows the changes to the aggregate log Lerner Index (black line) and the empirical factors (dash-dotted line) specified by equation (12). The difference between the lines is due to changes in price sensitivity.

In 2019, it is 5 percent lower than in 2006. Thus, without a change in price sensitivity, the trends in the data are consistent with a modest decline in markups. If we were to estimate a restricted model in which α were held fixed over time, then we would estimate that the log Lerner index exactly tracks the contribution of the observable factors.

Thus, to understand rising markups among the consumer products that we examine, it is necessary to have an understanding of consumer price sensitivity and how it has changed over time. An econometrician with data on product ownership, market shares, prices, and consumer purchasing patterns—which are sufficient to recover $\lambda(\cdot)$ within a specific modeling context—could make incorrect inferences about markup trends *unless* the model also allows for changes in price sensitivity. This points to a strength of our modeling approach: as we estimate demand separately for each category and each year, our estimates of price sensitivity can adjust flexibility over time with the shifts in the empirical variation in the data.

We also explore the cross-category variation in the components of equation (12). Within individual years, we find that variation in price sensitivity explains a modest fraction of the cross-sectional variation in markups: 16 percent in 2006 and 10 percent in 2019. Thus, it must be that other features, such as consumer purchasing patterns and multi-product ownership, explain most of the variation in markups across categories. As these features are included in the observable factors component of the econometric decomposition, observable factors are informative of cross-sectional variation in markups even without identification of α.

We also consider a panel regression with observations at the category$\times$year level in which the dependent variable is the year-over-year change in the (log) aggregate Lerner Index and the independent variable is the year-over-year change in price sensitivity. We find that changes in price sensitivity within categories over time explain 59 percent of the within-category variation in markups over time. Though changes in price sensitivity explain the majority of the changes in markups at the category level, a meaningful portion of the category-level variation over time is explained by observable factors. We report these results via regression in Appendix Table G.5.

This finding, combined with the cross-sectional results, points to the flexibility of our model to incorporate observable factors that are important for determining markups. Further, this highlights that our demand specification is sufficiently rich to attribute much of the variation in markups across categories to observable factors that are uncorrelated with mean price sensitivity. This need not be the case with less flexible demand systems. For example, with constant elasticity demand, the Lerner index only varies due to differences in price sensitivity (i.e., $\lambda_t = p_t$ and $\ln(\lambda_t/p_t) = 0$).

5.3 Explanations for the Decline in Price Sensitivity

Why does consumer price sensitivity decline? We consider three sets of possible explanations. First, we consider whether our results might reflect an increasing selection by price-sensitive consumers away from the retailers in our sample. Second, we consider whether the change may reflect investments by firms to reduce price sensitivity directly, e.g., through marketing expenditures. Third, we consider whether there is corroborating evidence for broad secular changes in price sensitivity. We caveat that our analyses below are suggestive and may be better addressed in future research.

To address selection, we analyze how aggregate consumer spending patterns change over time, and we consider how our estimates change with an extended sample of retailers. In Appendix F, we document that revenue shares for the main retail channels that our products are purchased—mass merchandisers, grocery stores, drug stores, warehouse clubs, and dollar stores—are stable over our sample period. We validate this within the NielsenIQ data and using auxiliary revenue data from large retailers. These data indicate that there are no broad shifts in consumer spending among these channels. In Appendix E.4, we address selection across retailers within these channels by providing estimates from an extended sample that incorporates additional large retailers that are observed in the Consumer Panel Data, but not the Retail Scanner Data. Our estimated markup and price sensitivity trends are very similar to our baseline estimates. These analyses mitigate concerns that shifts across channels by price-sensitive consumers might be driving our results.

The second set of possibilities is that firms make investment decisions that lower consumer price sensitivity. Such decisions might be reflected in marketing expenditures, R&D expenditures, or the variety of products that they offer for a particular brand. In Appendix F.3, we also show that changes in these variables do not explain changes in price sensitivity. Therefore, we do not find support for the hypotheses that declining price sensitivity is due to firm-level investment decisions.

Finally, we consider whether changes in price sensitivity may reflect exogenous shifts in preferences that are not the result of changes to supply. To explore this possibility, we exam-

Figure 6: Coupon Use Over Time

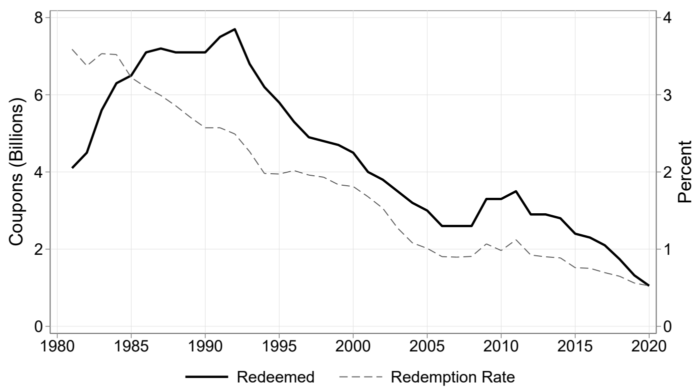

Notes: This figure shows the annual number of coupons redeemed (left axis) and the redemption rate out of all issued coupons (right axis). From 2006 to 2019, coupon redemptions fell from 2.6 billion to 1.3 billion, and the redemption rate fell from 0.90 percent to 0.56 percent. Annual estimates reflect total coupon usage for consumer products in the United States across all channels, including free standing inserts and electronic coupons.

ine other information about consumer shopping patterns. In particular, we look at the use of coupons and estimates of time spent shopping for consumer products. Coupon redemptions are a plausible proxy for price sensitivity because they typically involve a small amount of effort in order to obtain a discount on price. To evaluate coupon use, we collect statistics on the number of coupons distributed and redeemed for consumer packaged goods from 1981 through 2020. These statistics reflect industry estimates of coupon use across all channels, including free standing inserts and electronic coupons.[27]

Figure 6 plots the aggregate coupon usage over time. The black line reports the number of coupons redeemed each year (left axis). From 1981 to 1992, the number of coupons redeemed roughly doubled, from 4.1 billion to 7.7 billion. Since that year, there has been a steady decline in the number of coupons redeemed, with the exception of a brief bump due to the recession starting in 2009. Over our sample period, the number of coupons redeemed has fallen in half, from 2.6 billion in 2006 to 1.3 billion in 2019.

This trend reflects a decreasing propensity of consumers to use coupons, rather than coupon availability. To highlight this, the dashed line plots the percent of coupons that are redeemed out of all the coupons that were distributed (right axis). Redemption rates are declining over the entire sample period. From 1981 to 1992, the decline reflects the fact that the growth in the distribution of coupons outpaced the growth in coupon redemption rates. From 1992 to 2015, the annual number of coupons issued remained high while redemption rates fell. In 2015, 316 billion coupons were distributed, compared to 309 billion in 1992. From 2016 to 2020, fewer coupons were distributed each year, but redemption fell even faster. The redemption rate fell from 0.90 in 2006 to 0.56 in 2019.

[27]Industry estimates were obtained from reports by two companies, NCH Marketing from 1981 through 2002, and Inmar Intelligence from 2003 through 2020.

Concurrently, adults in the U.S. spent less time shopping for consumer products. Data from the American Time Use Survey indicate that both the frequency and duration of shopping trips declined over our sample period. For adults between the ages of 25 and 54, time spent on consumer goods purchases fell by 21 percent, from 3.01 to 2.38 hours per week.[28] We also find that, in the consumer panel data, households visit approximately 10 percent fewer unique retailers each week on average in 2019 compared to 2006. This seems to reflect a long-run trend. Aguiar and Hurst (2007b) and Aguiar et al. (2013) provide evidence for reduced time spent shopping between the 1960s and 2003 and between 2003 and 2010, respectively.

Overall, the declining use of coupons and the reduced time spent purchasing consumer goods suggest a fundamental shift in consumer shopping behavior that is consistent with lower price sensitivity arising from exogenous factors. Both trends indicate that consumers are less willing to exert effort to obtain lower prices. For instance, Aguiar and Hurst (2007a) provide evidence that shopping intensity is negatively correlated with consumer prices. Notably, the decline in coupon use began in the early 1990s, before the rise of online retail. We view this as additional evidence that declining price sensitivity reflects a longer-run secular trend. A potential explanation for this trend is an increase in the opportunity costs of time spent shopping, possibly due to changes in preferences for leisure, or changes to labor supply and the within-household distribution of wages. Consistent with the latter, Griffith et al. (2022) provide evidence that the opportunity cost of time for households in the United Kingdom has increased since the 1980s, and that this change is correlated with an increase in labor force participation and earnings among secondary earners.[29]

6 Markups, Welfare, and Consumer Surplus

In this section, we analyze how consumer surplus, producer surplus, and total welfare for consumer products have changed over time. We also examine various counterfactual scenarios in order to estimate the deadweight loss from (changes in) market power and to explore the consequences of rising markups for consumers and firms.

We follow Small and Rosen (1981) and calculate consumer surplus in dollar terms.[30] This yields the additional consumer surplus provided by the products in our sample, relative to a counterfactual in which only the outside option is available to consumers. Thus, it can be

[28] The American Time Use Survey reports both the frequency of adults participating in an activity in a given day, which declined by 5 percent, and the daily time spent conditional on participation, which declined by 16 percent.

[29] An alternative potential explanation, following results in the marketing literature, is that consumers are responding to broad shifts in the pricing behavior of firms. For example, Mela et al. (1997) argues that price-oriented promotions increase consumer price sensitivity in the long run. Therefore, a decline in price sensitivity could potentially be a response to a large-scale decline in price-oriented promotional activity.

[30] With the observed set of products, consumer surplus is given by $CS = -\frac{1}{N}\sum_i \frac{1}{\alpha_i} \ln\left(\sum_j \exp\left(w_{ij}\right)\right)$ where $w_{ij} = \beta_i^* + \alpha_i^* p_{jcrt} + \xi_{jr} + \xi_{cr} + \xi_t + \Delta\xi_{jcrt}$ for the inside products ($j > 0$), $w_{0j} = 0$ for the outside good ($j = 0$), and N denotes the number of consumers.

Table 3: Annual Surplus and Welfare Per Capita

(a) 2006 Preferences and Costs

Specification	CS	PS	W	% change CS	% change W
Baseline	678	264	942	0.0	0.0
Prices Scaled to 2019 Price Levels	653	268	920	-3.7	-2.3
Markups Scaled to 2019 Markup Levels	600	293	893	-11.5	-5.2
Prices Equal to Marginal Costs	999	0	999	47.3	6.0

(b) 2019 Preferences and Costs

Specification	CS	PS	W	% change CS	% change W
Baseline	942	369	1311	0.0	0.0
Prices Scaled to 2006 Price Levels	975	348	1323	3.5	0.9
Markups Scaled to 2006 Markup Levels	1111	236	1347	17.9	2.7
Prices Equal to Marginal Costs	1379	0	1379	46.4	5.2

Notes: This table reports consumer surplus (CS), producer surplus (PS), and welfare (W) per capita based on estimated demand parameters ("Baseline") and for counterfactual scenarios that hold fixed preferences and marginal costs and vary the price levels.

interpreted as the equivalent variation that would compensate consumers for the loss of the product-retailer combinations in our data.[31] Our measure of producer surplus reflects variable profits and is measured as price less marginal costs multiplied with quantities: $PS = \sum_{j>0}(p_j - c_j)q_j$. Our measure of producer surplus does not reflect fixed costs, thus this analysis does not inform whether brand manufacturers earn economic profit. We measure welfare (W) as the sum of producer and consumer surplus. The deadweight loss that exists in an observed equilibrium can be calculated by comparing the welfare that obtains with the equilibrium to the welfare that obtains under a counterfactual with prices set equal to marginal costs.

These exercises have an important limitation when considering overall welfare. We do not take a stance on utility received outside of the products in our sample, and our consumer surplus calculations are relative to the outside utility value. Thus, overall consumer welfare may decline even if we find increasing consumer surplus for the products in our sample. For example, if housing became more expensive, consumers may get greater relative utility from consumer products but be worse off. For these reasons, we focus on exploring the relationship between markups and welfare within the products and markets of our sample.

Table 3 shows per capita consumer, producer surplus, and welfare for 2006 and 2019 using observed prices ("Baseline") and prices under different counterfactual scenarios. To compute counterfactual values, we hold fixed estimated preference parameters and marginal costs, and we simulate consumer choices using different prices. We consider three counterfactual scenarios. First, we scale all prices by the average realized price change for all products in the same category from one year to another (e.g., from 2006 to 2019). Second, we scale all markups

[31] The model assumes complete information, while, in reality, consumers may face search costs or other frictions that are not explicitly captured by this representation. Consumers may engage in search both across and within retailers. If search costs are increasing over time, that may be reflected in our estimates as a reduction in price sensitivity.

by the average realized markup change for all products in that category from one year to another. Because we hold marginal costs fixed, scaling 2006 prices to match 2019 markups results in higher prices than what we observe in the data. Third, we consider a counterfactual where prices equal marginal costs (i.e., no markups). The last two columns in each panel show changes in consumer surplus and welfare relative to the baseline scenario.[32]

Comparing the baseline scenarios, the results indicate that per capita consumer surplus increased by about 39 percent (i.e., about 2.6 percent annually) between 2006 and 2019, from \$678 to \$942. As average prices did not decline and perceived relative quality did not increase, the increase in consumer surplus is likely due to lower price sensitivity, i.e., that consumers receive lower disutility from any given price in 2019. Along with higher markups, producer surplus increased over the period, from \$264 to \$369 per capita. Thus, approximately 70% of the increase in welfare have accrued to consumers.

Markups are costly for consumers. With marginal cost pricing, consumer surplus would be substantially higher in both 2006 and 2019, as shown by the final specification in each panel. Our estimates suggest that markups in 2006 reduced per capita welfare from \$999 to \$942 (6 percent). In 2019, markups reduced welfare by about 5 percent.

The changes in markups over this period are economically meaningful. Holding fixed the 2006 preferences, marginal costs, and product assortments, increasing markups to 2019 levels would reduce consumer surplus by 11.5 percent. However, markups trends do not occur in isolation. Changes in markups are often concurrent with and in response to other factors. For example, declining marginal costs mitigate the impact of rising markups on prices and consumer welfare. When scaling up prices—which are the relevant demand variables—to match 2019 levels, the decrease in consumer surplus is much smaller (3.7 percent). Analogous results obtain if 2019 markups and prices are scaled down to 2006 levels.

Thus, to interpret the impacts of changing markups on welfare, it is necessary to take a stand on what other factors are changing at the same time. Markups are equilibrium objects that are determined by supply and demand. If marginal costs and price sensitivity had not changed, the aggregate trends in markups would have likely looked quite different. This is an important consideration for potential policy responses to markup trends.

In our final analysis, we analyze how the change in consumer surplus varies by income. For this purpose, we calculate the log of consumer surplus per purchasing decision separately by each quartile of the income distribution and for each category-year. We relate these values to category and year fixed effects and document the coefficients across years in Figure 7. The results indicate that the increase in per capita consumer surplus between 2006 and 2019 is mainly driven by consumers with relatively high income and takes place during the second half of the sample period. In contrast, the lowest quartile of the income distribution has lower

[32]To mitigate the potential impact of outliers, we winsorize the consumer surplus estimates at the 1st and 99th percentiles from the pooled set of 133 categories across 14 years and 6 scenarios: baseline, scaling prices or markups to 2006 levels, scaling prices or markups to 2019 levels, and prices equal to marginal costs (11,166 observations).

Figure 7: Consumer Surplus Over Time By Income Group

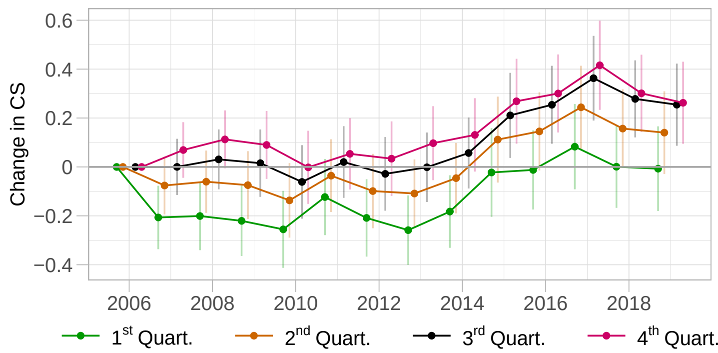

Notes: This figure reports coefficients and 95 percent confidence intervals of a regression of the log of consumer surplus by purchase on year dummies, controlling for category fixed effects, separately for different quartiles of the income distribution.

consumer surplus through 2016. The reduction in consumer surplus for the lowest-income households coincides with the increase in real prices in the first half of our sample. After this point, real prices fall and consumer surplus for this quartile increases, recovering to 2006 levels at the end of the sample period. In Figure G.6 in the Appendix, we repeat the analysis dividing the sample into deciles. The results confirm that changes in consumer surplus are strongly associated with the income distribution. Consumers in the highest income group see increases in consumer surplus over time, while lower income households have, on average, lower consumer surplus over our sample period. These findings suggest that changes in market power and consumer preferences over time have important distributional consequences.

7 Conclusion

This paper analyzes the evolution of markups in consumer products in the United States between 2006 and 2019. We estimate demand with flexible consumer preferences and recover time-varying markups for individual products under the assumption of profit maximization. Our results indicate that markups increased by about 30 percent during our sample period. This reflects within-product changes and is primarily due to reductions in marginal costs, rather than increases in (real) prices. Changes in marginal costs, along with declining consumer price sensitivity, account for the vast majority of the time series variation in aggregate markup changes between 2006 and 2019.

Our model indicates that consumer surplus has increased despite rising markups, though the increases are concentrated among higher-income consumers. The standard welfare measure in our analysis does not account for changes to outside options or why price sensitivity has declined, and it imposes simplifying assumptions, such as full information. We welcome future work that explores welfare impacts in greater detail while addressing some of these limitations.

References

AGHION, P., N. BLOOM, R. BLUNDELL, R. GRIFFITH, AND P. HOWITT (2005): "Competition and Innovation: An Inverted-U Relationship," Quarterly Journal of Economics, 120, 701–728.

AGUIAR, M. AND E. HURST (2007a): "Life-cycle prices and production," American Economic Review, 97, 1533–1559.

——— (2007b): "Measuring trends in leisure: The allocation of time over five decades," Quarterly Journal of Economics, 122, 969–1006.

AGUIAR, M., E. HURST, AND L. KARABARBOUNIS (2013): "Time use during the great recession," American Economic Review, 103, 1664–1696.

ANDERSON, E., S. REBELO, AND A. WONG (2023): "Markups Across Space and Time," Working paper.

ARCIDIACONO, P., P. B. ELLICKSON, C. F. MELA, AND J. D. SINGLETON (2020): "The Competitive Effects of Entry: Evidence from Supercenter Expansion," American Economic Journal: Applied Economics, 12, 175–206.

ATALAY, E., E. FROST, A. T. SORENSEN, C. J. SULLIVAN, AND W. ZHU (2023): "Scalable Demand and Markups," Tech. rep., National Bureau of Economic Research.

BACKUS, M., C. CONLON, AND M. SINKINSON (2021): "Common Ownership and Competition in the Ready-to-Eat Cereal Industry," Working paper.

BERRY, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," RAND Journal of Economics, 25, 242–262.

BERRY, S. AND P. JIA (2010): "An Empirical Analysis of the Airline Industry," American Economic Journal: Microeconomics, 2, 1–43.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," Econometrica, 841–890.

——— (2004): "Differentiated products demand systems from a combination of micro and macro data: The new car market," Journal of political Economy, 112, 68–105.

BERRY, S. T. AND P. A. HAILE (2021): "Foundations of Demand Estimation," *Handbook of Industrial Organization, Volume 4*, edited by Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, Elsevier.

——— (2022): "Nonparametric Identification of Differentiated Products Demand Using Micro Data," Working paper.

BRAND, J. (2021): "Differences in Differentiation: Rising Variety and Markups in Retail Food Stores," Working paper.

BRUNNER, D., F. HEISS, A. ROMAHN, AND C. WEISER (2020): "Reliable Estimation of Random Coefficient Logit Demand Models," Working paper.

BUTTERS, R. A., D. W. SACKS, AND B. SEO (2022): "How Do National Firms Respond to Local Shocks? Evidence from Excise Taxes," American Economic Review, 112, 1737–72.

CAOUI, E. H., B. HOLLENBECK, AND M. OSBORNE (2023): "The Impact of Dollar Store Expansion on Local Market Structure and Food Access," Working Paper.

CHEVALIER, J. A., A. K. KASHYAP, AND P. E. ROSSI (2003): "Why Don't Prices Rise during Periods of Peak Demand? Evidence from Scanner Data," American Economic Review, 93, 15–37.

CILIBERTO, F., C. MURRY, AND E. TAMER (2021): "Market Structure and Competition in Airline Markets," Journal of Political Economy, 129, 2995–3038.

CONLON, C. AND J. GORTMAKER (2020): "Best Practices for Differentiated Products Demand Estimation with PyBLP," RAND Journal of Economics, 51, 1108–1161.

DE LOECKER, J. (2011): "Recovering Markups from Production Data," International Journal of Industrial Organization, 29, 350–355.

DE LOECKER, J. AND J. EECKHOUT (2021): "Global Market Power," Working paper.

DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): "The Rise of Market Power and the Macroeconomic Implications," Quarterly Journal of Economics, 135, 561–644.

DE LOECKER, J. AND P. SCOTT (2022): "Markup Estimation using Production and Demand Data. An Application to the US Brewing Industry," Working paper.

DE LOECKER, J. AND F. WARZYNSKI (2012): "Markups and Firm-Level Export Status," American Economic Review, 102, 2437–2471.

DELLAVIGNA, S. AND M. GENTZKOW (2019): "Uniform Pricing in US Retail Chains," Quarterly Journal of Economics, 134, 2011–2084.

ELZINGA, K. G. AND D. E. MILLS (2011): "The Lerner Index of Monopoly Power: Origins and Uses," American Economic Review, 101, 558–564.

GANAPATI, S. (2021): "Growing Oligopolies, Prices, Output, and Productivity," Americal Economic Journal: Microeconomics, 13, 309–327.

——— (2024): "The Modern Wholesaler: Global Sourcing, Domestic Distribution, and Scale Economies," Working paper.

GANDHI, A. AND J.-F. HOUDE (2020): "Measuring Substitution Patterns in Differentiated Products Industries," Working paper.

GANDHI, A. AND A. NEVO (2021): "Empirical Models of Demand and Supply in Differentiated Products Industries," *Handbook of Industrial Organization, Volume 4*, edited by Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, Elsevier.

GRIECO, P. L. E., C. MURRY, AND A. YURUKOGLU (2023): "The Evolution of Market Power in the US Automobile Industry," Quarterly Journal of Economics, forthcoming.

GRIFFITH, R., W. M. JIN, AND V. LECHENE (2022): "The Decline of Home-cooked Food," Fiscal Studies, 43, 105–120.

HALL, R. E. (1986): "Market Structure and Macroeconomic Fluctuations," Brookings Papers on Economic Activity, 2, 285–322.

——— (1988): "The Relation Between Price and Marginal Cost in U.S. Industry," Journal of Political Economy, 96, 921–947.

——— (1990): "Invariance Properties of Solow's Productivity Residual," In Peter A. Diamond (ed), *Growth, Productivity, Unemployment: Essays to Celebrate Bob Solow's Birthday*, MIT Press.

HARBERGER, A. C. (1954): "Monopoly and Resource Allocation," American Economic Review (Papers and Proceedings), 44, 77–87.

HAUSMAN, J. A. (1996): "Valuation of New Goods Under Perfect and Imperfect Competition," in The Economics of New Goods, ed. by T. Bresnahan and R. Gordon, Chicago: University of Chicago Press, 209–248.

HENDEL, I. AND A. NEVO (2013): "Intertemporal Price Discrimination in Storable Goods Markets," American Economic Review, 103, 2722–2751.

HRISTAKEVA, S. (2022): "Vertical Contracts with Endogenous Product Selection: An Empirical Analysis of Vendor-Allowance Contracts," Journal of Political Economy, 130, 3202–3252.

LERNER, A. P. (1934): "The Concept of Monopoly and the Measurement of Monopoly Power," Review of Economic Studies, 1, 157–175.

MACKAY, A. AND N. H. MILLER (2023): "Estimating Models of Supply and Demand: Instruments and Covariance Restrictions," Working paper.

Mela, C. F., S. Gupta, and D. R. Lehmann (1997): "The Long-Term Impact of Promotion and Advertising on Consumer Brand Choice," Journal of Marketing Research, 34, 248–261.

Miller, N. H. (2024): "Industrial Organization and The Rise of Market Power," Working Paper.

Miller, N. H., M. Osborne, G. Sheu, and G. Sileo (2023): "Technology and Market Power: The United States Cement Industry, 1974-2019," Working paper.

Miller, N. H. and M. C. Weinberg (2017): "Understanding the Price Effects of the Miller-Coors Joint Venture," Econometrica, 85, 1763–1791.

Miravete, E. J., K. Seim, and J. Thurk (2022): "Robust pass-through estimation in discrete choice models," Tech. rep., Working paper, Department of Economics, University of Pennsylvania.

——— (2023): "Pass-through and tax incidence in differentiated product markets," International Journal of Industrial Organization, 90, 102985.

Nakamura, E. and D. Zerom (2010): "Accounting for Incomplete Pass-through," Review of Economic Studies, 77, 1192–1230.

Nevo, A. (2000a): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," RAND Journal of Economics, 31, 395–421.

——— (2000b): "A Practitioner's Guide to Estimation of Random Coefficients Logit Models of Demand," Journal of Economics and Management Strategy, 9, 513–548.

——— (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," Econometrica, 69, 307–342.

Peltzman, S. (2022): "Productivity, Prices and Productivity in Manufacturing: a Demsetzian Perspective," Journal of Law and Economics, 65.

Small, K. A. and H. S. Rosen (1981): "Applied Welfare Economics with Discrete Choice Models," Econometrica, 105–130.

Villas-Boas, S. B. (2007): "Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data," Review of Economic Studies, 74, 625–652.

Appendix

A Additional Details About the Data

A.1 Data Construction

The NielsenIQ Retail Scanner Data provides weekly sales and quantities for products at the level of a UPC for individual retail stores. A UPC (universal product code) corresponds to a barcode and distinct items in a retailer's point-of-sale system. We define products at the brand level, which abstracts away from features like package size and color variants. Thus, identical products may correspond to several UPCs. We follow the dataset documentation to construct revenues and total volume in units, where units are provided in the product details file that links UPCs to the units contained in each package and the unit of measure. We translate the provided units to standard measures, which are milliliter (for liquid volume), ounce (for weight), and count (no standard unit). Typically, all UPCs within a category are reported with the same unit of measure; when multiple unit types are reported, we retain UPCs that are reported with the type of units with the largest share of revenue in that category.

We then aggregate across UPCs, store, and weeks to get total revenues and units at the brand, chain, DMA, and quarter within a category. Chains are defined by NielsenIQ data indicating the store's parent company, and the DMA the store is located in is also provided by NielsenIQ. We construct prices as average unit prices by dividing revenue by total units. To reduce measurement error, we drop products that are extreme outliers in terms of their price—which we implement by dropping observations with a price below the 0.5 percentile or above the 99.5 percentile. We apply this screen before we restrict the data to the 22 DMAs in our baseline sample.

Brands are defined by NielsenIQ and are fairly narrow. In ready-to-eat cereals, "Cheerios," "Honey Nut Cheerios," and "Multigrain Cheerios" are three distinct brands. In cookies, brands include "Oreo," "Oreo Double Stuf", and "Mini Oreo." In yogurt, brands include "Yoplait," "Yoplait Go-gurt," "Yoplait Whips!," "Yoplait Thick & Creamy," and "Yoplait Light Thick & Creamy." As described in the main text, we retain the 20 brands with the largest revenues in each category as separate brands, and we aggregate the remaining UPCs into a single "fringe" brand in the category. Our baseline sample consists of 2,792 distinct products: 21 in each of 133 categories, minus one brand that does not appear in our focal DMAs.[33] Sales are

[33] The 20th-ranked brand in the paper towel category.

highly skewed toward larger brands; the top 20 brands represent 84 percent of revenues. The market share of private labels across categories and time is documented in Figure A.1. The 133 fringe brands are composed of 110,390 distinct NielsenIQ brands, or an average of 830 per category.

NielsenIQ lists products (UPCs) as belonging to one of 1128 distinct product categories ("modules"). The top 133 and 200 modules account for 55 percent and 74 percent of revenues, respectively. Our rankings exclude from consideration four categories that, for some years, exist in the Retail Scanner Data but not the Consumer Panel Data: prerecorded videos, magazines, cookware, and sunscreens.

We treat NielsenIQ product categories as well-defined product markets. Thus, for example, we do not combine the "Light Beer" and "Beer" categories. In principle, these categories could be combined, possibly with richer demand specification that allows for weaker substitution between light beer and beer. Our estimates do not account for cross-category substitution by consumers. Product categories belong to the following high-level departments according to NielsenIQ: "Dry Grocery," "Frozen Foods," "Dairy," "Deli," "Packaged Meat," "Fresh Produce," and "Alcoholic Beverages," "Health and Beauty Care," "Non-food Grocery," and "General Merchandise."

The composition of retailers included in the Retail Scanner Data is fairly stable over our sample period. There are small changes from 2006–2017 and a modest change starting in 2018. We do not find that compositional changes have much impact on our main results. See, for example, Appendix E.2. Due to the terms of the data agreement, we cannot provide details on the composition of retailers or identify individual retail chains.

We make several adjustments to the Consumer Panel Data when we construct consumer demographics. First, we impute household income using the midpoint of the bins provided in the Consumer Panel Data data. It is possible to obtain a comparable income measure for the highest-income bin because additional high-income bins are provided from 2006 to 2009; for this bin, we estimate a midpoint of $137,500. Second, we observe that many fewer consumers are in the top income bin in 2006 than in 2007 and subsequent years. To produce a more consistent demographic representation of consumers, we rescale the NielsenIQ projection weights in 2006 so that the top bin occurs with the same frequency as it does in 2007. We scale down the projection weights for the other bins in 2006 proportionately. The projection weights indicate the representatives of the consumer in the sample, and we use these weights for our demographic draws.

We construct micro-moments from the Consumer Panel Data by calculating the mean value of the demographic variables for consumers that purchase each product in a DMA-year-quarter-category. We use all consumers in the Consumer Panel Data and weight by the projection factors. For the construction of these values, we include only shopping trips to mass merchandisers ("Discount Stores"), grocery stores, and drug stores. We also drop trips to chains that account

Figure A.1: Private Label Shares Over Time

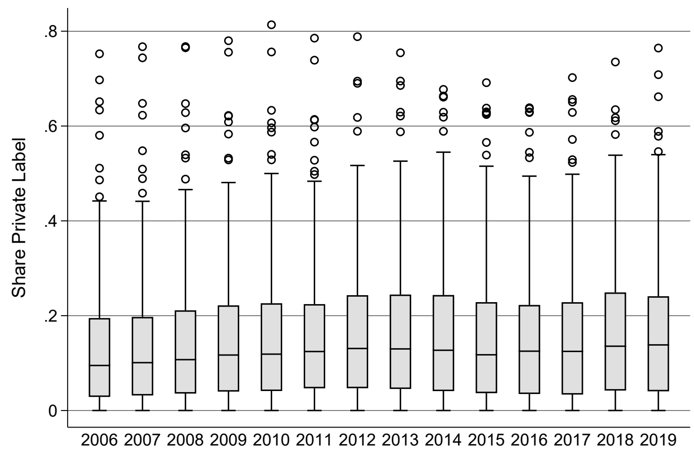

Notes: Figure provides the distribution of private label shares across our 133 baseline categories in each year.

for more than 5 percent of revenues in the Consumer Panel Data but are not in the sample of retailers in the scanner data. Two chains meet this latter screen, one of which is in the three channels we include. We make these adjustments so that the micro-moments we construct use a more similar set of retailers to those in sales data.

A.2 Market Size Calculations

As is standard in applications involving random coefficients logit demand, an assumption on market size is needed in order to convert observed quantities into market shares and then estimate the model. Our approach is to use market sizes that scale with the population of the region and the number of stores operated by the retail chain within the region.

Recall from Section 2.2 that the quantity demanded in our model is given by $q_{jcrt}(\boldsymbol{p}_{crt};\theta) = s_{jcrt}(\boldsymbol{p}_{crt};\theta)M_{crt}$, where $s(\cdot)$ is the market share, $\boldsymbol{p}_{crt}$ is a vector of prices, and M_{crt} is the market size, a measure of potential demand. We apply the following steps separately within each product category:

1. Obtain a time-varying "base" value by multiplying the population (at the region-year level) with the number of stores (at the chain-region-quarter-year level). This obtains $BASE_{crt} \equiv POP_{ry(t)} \times NS_{crt}$ where $POP_{ry(t)}$ is the population in region r and year $y(t)$ and NS_{crt} is the number of stores operated by retail chain c in region r and period t, where a period is a year-quarter.

2. Obtain the total quantity of the inside products across brands: $Q_{crt} = \sum_j q_{jcrt}$.

3. Calculate $\gamma_{cr} = E_t\left[\frac{Q_{crt}}{BASE_{crt}}\right]$ as the average quantity-to-base ratio among the periods

Figure A.2: Distribution of Market Shares of Inside Goods

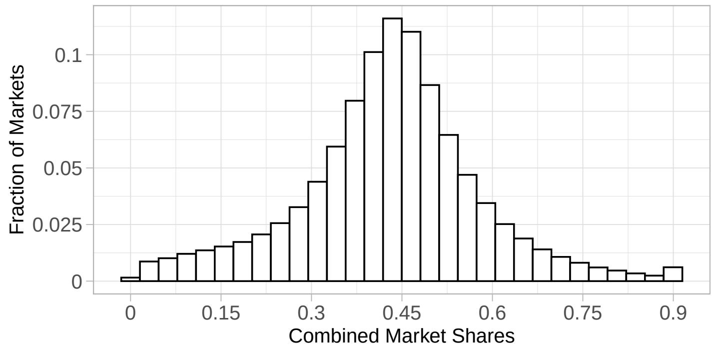

Notes: This figure shows the distribution of market shares of inside goods. Observations are at the chain-region-year-quarter level and reflect the sum of the market shares of all inside goods in a market at a given point in time.

observed for each retail chain and region. This can be used to convert the base value into units that are meaningful in terms of total quantity-sold. In the calculation of γ_{cr}, we exclude a handful of observations for which the base-adjusted quantity is less than 5 percent of the mean, which helps avoid extraordinary small inside good market shares.

4. We set the market size such that the combined share of the inside goods is around 0.45, on average, and we allow the market size to scale with population and number of stores, as captured by the base value. Specifically, we calculate the market size according to

$$M_{crt} = \frac{1}{0.45}\gamma_{cr} BASE_{crt}$$

which generates markets sizes for each retail chain, region, quarter, and year. This yields combined inside shares $\frac{Q_{crt}}{M_{crt}} = 0.45 \frac{Q_{crt}}{BASE_{crt}} \frac{1}{\gamma_{cr}}$.

5. For a small minority of cases (<5 percent of markets), this procedure generates a combined share of the inside goods that exceeds 0.90 in some periods, which is high enough that we encounter numerical problems in estimation. For any category×chain×region combination in which this occurs, we repeat the steps above using the alternative conversion factor $\tilde{\gamma}_{cr} = 0.5 \times \max_t \left(\frac{Q_{crt}}{BASE_{crt}}\right)$, which sets the maximum of the combined shares equal to 0.90.

Figure A.2 shows the distribution of combined market shares of inside goods. By construction the market shares are centered around 0.45 (step 4), and the small peak around 0.9 indicates the imposed maximum that is described in step 5.

We consider alternative definitions for market size as robustness checks in Appendix E.9.

B Estimation Details

This appendix provides details on the estimation procedure. We estimate the parameters in two steps, which is possible because the mean price parameter and the other ("nonlinear") structural parameters are identified by two independent sets of moments. The parameters for estimation are $\theta = (\alpha, \Pi_1, \Pi_2, \sigma)$. We first estimate $\theta_2 = (\Pi_1, \Pi_2, \sigma)$ and then estimate α, the mean price parameter, in the second step. Our micro-moments identify θ_2 but not α (Berry et al., 2004; Berry and Haile, 2022), and the covariance restriction exactly identifies α given θ_2 (MacKay and Miller, 2023). In principle, a single search could be used to estimate the parameters jointly, as is standard practice for applications that rely on instruments for identification. However, our approach has computational benefits, as we explain below.

B.1 First Step

In the first estimation step, we use the micro-moments to pin down the "nonlinear" parameters, i.e., $\theta_2 = (\Pi_1, \Pi_2, \sigma)$. To implement this, we estimate GMM while holding fixed the price parameter at a given value. Because the parameters are identified separately, the specific value chosen for the price parameter has no impact on the micro-moment contributions to the objective function.[34]

For any candidate θ_2, there is a unique vector of the mean product valuations that align the predicted and observed shares (δ). For example, in the special case of $\theta_2 = \overrightarrow{0}$ the mean valuations have a closed-form solution:

$$\delta_{jcrt}\left(\theta_2^{(0)}\right) \equiv \log(s_{jcrt}) - \log(s_{0crt}) \tag{B.1}$$

We proceed to estimate θ_2 based on equation (7) while holding fixed the price parameter. For each candidate θ_2, we recover the mean valuations $\{\delta_{jcrt}(\theta_2)\}$ using the contraction mapping of Berry et al. (1995) with a numerical tolerance of 1e-9. We then calculate the micro-moments with $\{\delta_{jcrt}(\theta_2)\}$ and $\bar{\alpha}$. We choose the parameters $\{\delta_{jcrt}(\theta_2)\}$ that minimize the micro-moment contributions to the objective function. We apply equal weights to each micro-moment in estimation.

B.2 Second Step

In the second step, we hold fixed the estimated nonlinear parameters and choose the price parameter that minimizes the objective based on the covariance restriction moment. In other words, we estimate α taking as given the estimates of θ_2 obtained in the first step. This is possible because micro-moments do not identify the mean price parameter (Berry and Haile, 2022).

[34] We initialize this step with a price parameter $\bar{\alpha}$ such that the average elasticity when $\theta_2 = \overrightarrow{0}$ is equal to -7, which corresponds to the average starting value that we use in the second step (see below).

To do so, we recover $\Delta\xi_{jcrt}(\theta_2)$ as the residual from the OLS regression of $(\delta_{jcrt}(\theta_2) - \alpha p_{jcrt})$ on the fixed effects for each candidate α. We also obtain marginal costs from equation (5), looping over the chain-region-quarter combinations, and then recover $\Delta\eta_{jcrt}(\theta_2)$ as the residual from the OLS regression of marginal costs on the fixed effects. We are then able to calculate the loss function, update the candidate α, and repeat to convergence. We constrain the search to negative values of α. The constraint imposes downward-sloping demand for a consumer with the mean income level.

A complication is that there may be two values for α that satisfy the covariance restriction, with the smaller (more negative) value being the true price parameter under sensible conditions (MacKay and Miller, 2023). Care must then be taken to ensure that the estimator converges to the smaller value. Figure G.7 illustrates this in the context of ready-to-eat cereals. Each panel traces out the contribution of the covariance restriction to the objective function for different values of α. In 2006, a unique negative α satisfies the covariance restriction, and the constraint we place on the parameter space ($\alpha < 0$) is sufficient to recover the correct estimate. In other years, both possible solutions are negative, and thus could be obtained from estimation, even though the larger (less negative) value is implausibly close to zero.[35]

We proceed by selecting starting values of $\alpha^{(0)} = \phi\tilde{\alpha}$ where $\tilde{\alpha}$ is such that the average elasticity is -1 when $\theta_2 = \overrightarrow{0}$, and $\phi = (2, 4, 6, 8, 10, 12)$. Thus, for each year-category, we estimate with six different starting values. As these starting values are quite negative, the estimator tends to converge on the more negative value of the price parameter that satisfies the covariance restrictions. In the category-years for which the estimator finds both solutions, we select the more negative solution as our estimate of α. This appears to be a robust solution given the θ_2 we estimate.

The two-step approach allows us to more readily evaluate the possibility of multiple solutions for the covariance restriction. In addition, the objective function contribution of the covariance restriction moment can be poorly behaved for unreasonable candidate θ_2 parameters that would be considered if estimation of both θ_2 and α were performed simultaneously. Thus, our two-step approach to estimation yields both speed and numerical stability, both of which are important given the scale of the empirical exercise.

B.3 Computation Notes

Our code builds on the BLPestimatoR package for R (Brunner et al., 2020).[36] The package has a slim R skeleton and fast C++ routines for computationally intensive tasks. As micro-moments and covariance restrictions are missing from the package, we added code to cover

[35]The larger values imply that firms are pricing in the inelastic portion of their residual demand curves. A related complication is that the numerical stability of the moment tends to deteriorate as the candidate α approaches the higher solution, which can lead to convergence issues if the estimator considers parameters near the higher solution.

[36]`https://github.com/cran/BLPestimatoR`, last accessed March 26, 2021

that part of estimation. All time-critical parts are in C++. In early experiments, we replicated our results for some categories using the PyBLP package for Python (Conlon and Gortmaker, 2020).[37] We ultimately selected the augmented R package because it allowed us to calculate the micro-moments more quickly; our understanding is that the speed of PyBLP has improved substantially during the course of our research.

In estimation, we use BFGS with a numerical gradient. When searching for θ_2 in the first step of estimation, there are a handful of categories for which BFGS fails to converge, and for those categories we use Nelder-Mead instead. We estimate each category-year combination in parallel using the HILBERT computational cluster at the University of Düsseldorf. There are 2800 estimation routines (200 categories and 14 years). Each routine requires one CPU core and up to 9GB of memory. The longest runs take slightly more than 72 hours and most finish in less than 24 hours. The entire estimation procedure takes around one week.

[37] `https://github.com/jeffgortmaker/pyblp`, last accessed March 26, 2021.

C Validation for Selection of Product Categories

We conduct two validation checks to assess the reasonableness of our approach. First, we examine one product category—ground/whole bean coffee—to assess the ability of our method to capture marginal costs. Coffee is somewhat unique among our product categories in that a single ingredient (coffee beans) accounts for a substantial portion of marginal costs and commodity prices for this ingredient are well-established. Second, we compare the own-price elasticities of demand that we obtain to those obtained in the literature.

C.1 Marginal Cost Estimates

Figure C.1: Prices and Marginal Costs of Coffee Over Time

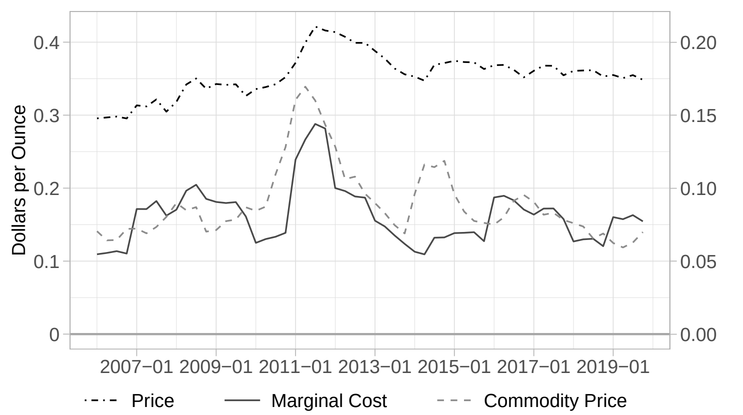

Notes: This figure plots the time series of quantity-weighted prices and marginal costs (solid line) for ground/whole bean coffee. Prices are observed and marginal costs are recovered from the profit-maximization conditions. Also shown is the commodity price index for coffee (dashed gray line), which is scaled following the right axis.

Figure C.1 plots the time series of quantity-average weighted prices (dot-dash line) and marginal costs (solid line) for coffee. Prices are observed, and marginal cost are recovered according to equation (5). The gray dashed line plots the commodity price index for coffee, which is scaled separately on the right axis.[38] Overall, our recovered estimates of marginal costs are strongly correlated with the commodity price index. Our method is able to capture the large spike in commodity prices in 2011, which is reflected in the spike in marginal costs. A regression of average marginal costs on the commodity price yields a coefficient of 0.950 (s.e. $= 0.17$), and the correlation between the two time series is 0.60.[39] This result implies that our marginal cost measure captures the fluctuations in this component of marginal costs. Holding all else equal, a 1 unit increase in a cost component should result in a 1 unit increase

[38] Data on coffee commodity prices were obtained from Macrotrends.net. Available here: `https://www.macrotrends.net/charts/commodities`, last accessed March 1, 2022

[39] Regressing average marginal costs on the one-period lagged commodity price yields a coefficient of 1.020 and a correlation of 0.65. This slightly stronger relationship may reflect the use of contracts. The relationship is weaker with longer lags.

Table C.1: Average Product-Level Own-Price Elasticities of Demand

Category	Our Estimate	Literature Estimate	Citation
Beer	-4.22	-4.74	Miller and Weinberg (2017)
Ready-to-Eat Cereal	-2.29	-2.42	Backus et al. (2021)
Yogurt	-3.06	-4.05	Hristakeva (2022)

Notes: The Miller and Weinberg (2017) estimate is the median product-level elasticity obtained with the RCNL-1 specification. Our corresponding estimate is the median own-price elasticity across all years, combining "Beer" and "Light Beer," which are not distinguished in Miller and Weinberg (2017). The Backus et al. (2021) estimate is the median product-level elasticity obtained with the "prices only" specification; our corresponding estimate is the median own-price elasticity across all years. Hristakeva (2022) reports a mean product-level elasticity from 2001–2010; to make things more comparable, we report our estimated mean own-price elasticity from 2006–2010.

in total marginal cost. We find that, on average, the commodity price is equal to 55 percent of estimated marginal costs. This is consistent with the literature, as Nakamura and Zerom (2010) find that coffee beans account for 45 percent of marginal costs based on data spanning 2001-2004. These results indicate the potential of our empirical approach to recover reasonable marginal cost estimates.

C.2 Elasticity Estimates in the Literature

Next, we compare our product-level own-price elasticities of demand to those obtained in the literature using similar data and models. In Table C.1, we report estimates for beer, ready-to-eat cereal, and yogurt, for which comparisons are possible. As shown, we obtain elasticities for beer, ready-to-eat cereal, and yogurt of -4.22, -2.29, and -3.06, respectively. To provide more comparable estimates, we report the median product-level own price elasticities for beer and ready-to-eat cereal, and the mean own-price elasticity from 2006–2010 for yogurt.[40] For beer, we combine beer and light beer categories to match Miller and Weinberg (2017), who do not distinguish between these categories. Miller and Weinberg (2017) report a median elasticity for beer of -4.74, Backus et al. (2021) reports a median elasticity for ready-to-eat cereal of -2.42, and Hristakeva (2022) reports a mean elasticity for yogurt of -4.05. Thus, we conclude that our methodology can obtain reasonable results that are consistent with analyses that make use of specific institutional details to a greater degree.

To provide a more detailed comparison, consider the empirical approach of Backus et al. (2021), which was developed concurrently. In their analysis of ready-to-eat cereals, Backus et al. (2021) use the Kilts NielsenIQ data over a similar time period (2007-2016) with a smaller sample of DMAs, retailers, and weeks. The supply model is quite similar, and the random coefficients logit demand model includes the same consumer demographics that we include in

[40]Every paper differs in the exact data sample used. For example, Hristakeva (2022) uses data from 2001–2010. Because we find rising markups over time for yogurt, restricting it to the earlier years of our sample provides a closer comparison. None of these papers allow preference parameters to vary over time.

our analysis. One key distinction is that Backus et al. (2021) also collect product characteristics that are included in the demand model. A second key distinction is that, instead of covariance restrictions, Backus et al. (2021) employ two sets of instruments that are constructed from input costs and the characteristics of other products (Berry et al., 1995; Gandhi and Houde, 2020). Despite these differences, we obtain similar elasticities and margins.[41] Furthermore, we run an additional specification for ready-to-eat cereals using product characteristics, and show that this does not materially affect our estimates (Appendix D.2).

[41]For cereals, our average unit price is 0.20 and our average estimated marginal cost is 0.10. We find that average markups for this category are relatively stable over time, which is consistent with the De Loecker et al. (2020) estimates for cereals over our sample period.

D Alternative Model Specifications

D.1 Random Coefficients Logit versus Logit Demand

Figure D.1: Implied Elasticities for Baseline and Logit Estimates

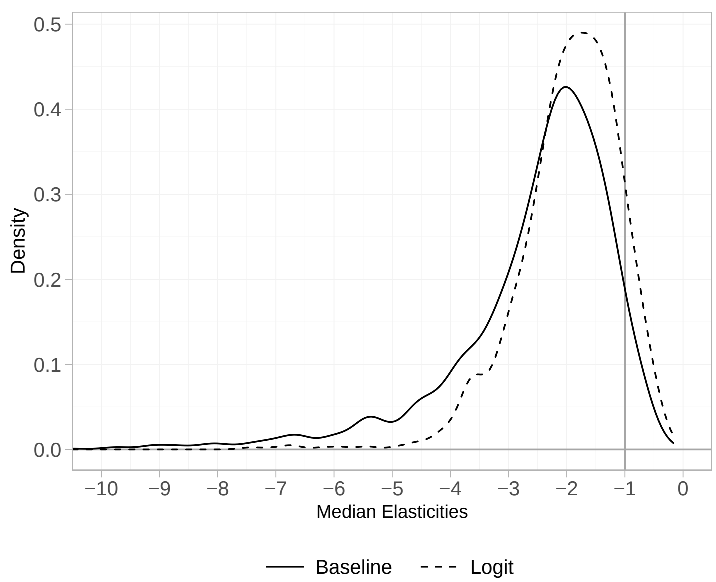

Notes: This figure plots the density of the median own-price elasticity by category and year. The solid black line shows the density of median elasticities using our baseline specification. The dashed line shows the density of median elasticities from a logit specification without random coefficients. Random coefficients allow for richer consumer heterogeneity.

We examine whether the consumer heterogeneity parameters we include in our baseline specification materially change the estimated elasticities and implied markups. For a comparison, we estimate a standard logit demand model ($\Pi_1 = 0$, $\Pi_2 = 0$, $\sigma = 0$) for all categories and years. Figure D.1 plots the density of median elasticities in our baseline model (black line) against those in the logit specification (dashed line).

Relative to the logit specification, our baseline estimates obtain more elastic demand estimates and smaller markups. The mean across the category-year median elasticity estimates is -2.60 in our baseline specification and -1.96 in the logit specification. More than twice as many estimates have a median elasticity > -1 (inelastic demand) with the logit specification. Median category-year markups are 0.120 higher in the logit specification (0.686 versus 0.566). These differences are all statistically significant (p-value < 0.001). We obtain an increasing trend in markups with the logit specification, but the trend is steeper, rising from 0.54 to 0.78.

D.2 Incorporating Additional Product Characteristics

The product fixed effects in our baseline model capture the utility contribution of an arbitrary set of non-price characteristics for the mean consumer. However, the baseline specification does not allow for heterogeneity in consumer tastes across non-price characteristics. This would require additional random coefficients that load on the interaction of consumer demographics with observed characteristics. A specification that incorporated these features would allow for a more flexible treatment of horizontal differentiation in the model.

As reported in Appendix C.2, our baseline estimates obtain nearly identical markups as Backus et al. (2021), who include product characteristics. This suggests that the inclusion of additional non-price characteristics may not be of first-order importance for pinning down markup levels. Here, we document the point estimates for the ready-to-eat cereals category for our baseline estimates, and we provide an additional specification where we follow Backus et al. (2021) and include additional product characteristics when estimating demand.

Panel A of Table D.1 reports the point estimates and standard errors for the mean price parameter and the demographic interactions, including the observed demographics (income and children) and the unobserved $N(0,1)$ draws. Fixed effects are included in estimation but not reported. Panel B of Table D.1 reports the number of observations, the median own-price elasticity, and the median Lerner index. Each column of the table corresponds to a different year, and each year is estimated independently. We use the standard GMM formula to calculate standard errors while clustering at the DMA level, and we apply a small-sample adjustment that scales up the standard errors to account for the fact that we have a small number of clusters.[42]

Our estimated parameters can change somewhat from year to year. For example, from 2006 to 2007, the price parameter changes from -18.19 to -10.75. This change is not due to convergence properties.[43] Changes in parameters reflect differences in the data from year to year. We see more modest changes in the price parameter over the remaining years for this category.

Meaningful year-to-year changes in parameter estimates can occur in other categories, but they appear to be idiosyncratic and are not frequent. Because we pool our results across more than 100 product categories, the presence of such idiosyncratic changes is not, in our view, a critical issue. We do not see any systematic changes in parameters for specific years of our sample.

We also test for the robustness of our estimates to the inclusion of product characteristics.

[42]An earlier version of this paper did not incorporate the additional small-sample adjustment. The adjustment delivers standard errors of the same order of magnitude as a jackknife estimate of standard errors for the price coefficients. MacKay and Miller (2023) demonstrate how the standard errors from the covariance restriction approach can be substantially smaller than IV standard errors because the estimator exploits observed variation in prices and quantities. We view the reported standard errors as indicating that we have a large number of observations and a good deal of variation in the data; inference for coefficients from specific categories is not central to our project.

[43]Figure G.7 shows the objective function remains smooth with a single minimum. Standard errors are small, which suggests that the price coefficient estimate is fairly precise conditional on the nonlinear parameters.

For this purpose, we follow a similar procedure to Backus et al. (2021). We collect data on characteristics at the UPC level, and we merge these characteristics to the UPCs that are associated with each product (brand) in our sample.[44] The characteristics include ingredients, nutritional information, and how the product was marketed. Specifically, we include dummy variables for whether the first ingredient is rice, oat, wheat, corn, protein, almond, or sugar; we include the amount per serving of sugar, fiber, sodium, saturated fat, calories, protein, iron, calcium, and cholesterol; and we include dummy variables for whether the product is marketed as for children, functional/healthy (e.g., heart healthy, antioxidants, etc.), natural, or with low value of "unhealthy" ingredients (e.g., low cholesterol, low fat, etc.). To reduce the dimension of product characteristics, we follow Backus et al. (2021) and project these 20 variables onto the first three principal components ($PC1, PC2, PC3$), which we use in estimation.[45] We interact these variables with our demographics (income and the presence of children) to allow for a product-consumer-specific constant in equation (2). For instance, this can in principle capture that households with children receive higher utility from cereals marketed for children compared to households without children. We do not include the principal components as separate variables without interactions since these are collinear with product fixed effects.

Table D.2 reports the resulting estimates. Many of the product characteristic interactions are statistically significant, but they do not substantially change our conclusions about markups in the ready-to-eat cereal industry. The price coefficients, elasticities, and implied markups are quite similar to those in our baseline estimates in most years.

[44] Our data on characteristics was obtained from Mintel. On average, we merge characteristics from 53 UPCs to each brand, excluding private label (1,039 merged UPCs) and fringe brands (2,559 merged UPCs). The characteristics are fairly stable within these brands.

[45] The first component is correlated with wheat, protein, fiber, and functional/healthy, the second component is correlated with oats, iron, and calcium, and the third is correlated with rice and low values of unhealthy ingredients.

Table D.1: Estimation Results for RTE Cereals

Panel A: Point Estimates and Standard Errors

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Price	-18.193	-10.746	-13.023	-10.369	-10.610	-9.139	-10.399	-10.959	-11.988	-12.003	-13.290	-15.483	-13.419	-16.920
	(0.452)	(0.376)	(0.371)	(0.376)	(0.243)	(0.169)	(0.199)	(0.268)	(0.167)	(0.339)	(0.281)	(0.479)	(0.281)	(0.548)
Demographic Interactions														
Income×Price	0.692	1.377	1.187	0.600	0.331	0.731	0.793	1.249	0.845	0.643	0.687	0.890	0.495	0.300
	(0.023)	(0.022)	(0.020)	(0.018)	(0.017)	(0.018)	(0.017)	(0.019)	(0.021)	(0.022)	(0.022)	(0.023)	(0.022)	(0.024)
Income×Constant	0.161	0.279	0.471	0.260	0.277	0.012	-0.065	-0.104	-0.056	-0.037	0.078	-0.004	0.230	0.305
	(0.025)	(0.073)	(0.070)	(0.059)	(0.041)	(0.014)	(0.007)	(0.015)	(0.005)	(0.017)	(0.013)	(0.022)	(0.016)	(0.042)
Children×Price	-0.455	-1.434	-0.718	1.139	1.634	2.832	3.326	2.373	2.325	2.407	2.908	2.372	2.402	2.174
	(0.059)	(0.046)	(0.045)	(0.040)	(0.038)	(0.042)	(0.043)	(0.044)	(0.045)	(0.050)	(0.053)	(0.057)	(0.059)	(0.055)
Children×Constant	7.087	5.113	5.674	2.615	3.563	0.856	0.656	0.712	0.465	1.114	2.999	1.480	4.320	5.261
	(0.574)	(0.654)	(0.545)	(0.441)	(0.466)	(0.125)	(0.093)	(0.130)	(0.022)	(0.302)	(0.194)	(0.265)	(0.233)	(0.655)
N(0,1)×Constant	5.768	4.209	5.310	2.737	4.453	0.708	0.480	0.525	0.078	2.106	5.979	3.101	8.743	10.405
	(0.540)	(0.670)	(0.588)	(0.590)	(0.619)	(0.351)	(0.355)	(0.524)	(0.519)	(0.651)	(0.357)	(0.559)	(0.476)	(1.256)

Panel B: Other Statistics

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Observations	15,441	16,336	16,604	16,791	17,241	17,329	16,444	16,213	16,443	15,829	15,487	14,365	18,850	17,805
Median Own Elasticity	3.366	2.027	2.573	2.073	2.030	1.743	2.083	2.167	2.344	2.257	2.439	2.767	2.324	2.967
Median Lerner	0.344	0.571	0.454	0.550	0.578	0.628	0.520	0.496	0.455	0.495	0.481	0.411	0.501	0.396

Notes: This table summarizes the results of estimation for the ready-to-eat cereals category for each year in the sample. Panel A provides the parameters and the standard errors, which are clustered at the region level and include a small-sample correction for the number of clusters. Panel B provides the number of product-chain-region-quarter observations, the median own price elasticity of demand, and the median Lerner index.

Table D.2: Alternative Estimation for RTE Cereals Including Product Characteristics

Panel A: Point Estimates and Standard Errors

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Price	-19.323	-11.621	-13.585	-9.272	-9.807	-9.448	-10.316	-10.973	-11.900	-11.740	-12.927	-16.856	-13.269	-17.356
	(0.550)	(0.376)	(0.409)	(0.139)	(0.243)	(0.139)	(0.169)	(0.218)	(0.167)	(0.259)	(0.490)	(0.537)	(0.251)	(0.566)
Demographic Interactions														
Income×Price	1.817	2.573	2.228	1.986	1.224	2.041	1.881	2.371	1.659	1.709	1.478	1.815	0.844	1.041
	(0.027)	(0.026)	(0.025)	(0.025)	(0.023)	(0.025)	(0.025)	(0.028)	(0.029)	(0.029)	(0.031)	(0.031)	(0.027)	(0.029)
Income×Constant	0.082	0.268	0.366	-0.202	-0.046	-0.252	-0.293	-0.329	-0.221	-0.231	-0.056	-0.077	0.265	0.361
	(0.044)	(0.097)	(0.089)	(0.013)	(0.031)	(0.009)	(0.006)	(0.014)	(0.007)	(0.016)	(0.057)	(0.037)	(0.019)	(0.060)
Children×Price	1.288	0.801	-0.290	-0.054	0.863	3.132	1.989	0.846	0.748	0.121	1.135	0.778	1.310	-0.724
	(0.079)	(0.058)	(0.060)	(0.050)	(0.050)	(0.058)	(0.056)	(0.059)	(0.063)	(0.066)	(0.072)	(0.084)	(0.068)	(0.070)
Children×Constant	5.174	3.983	4.311	1.214	1.528	0.593	0.756	0.950	0.693	1.285	2.327	2.248	3.724	5.667
	(0.502)	(0.495)	(0.449)	(0.058)	(0.227)	(0.050)	(0.036)	(0.090)	(0.023)	(0.161)	(0.577)	(0.313)	(0.159)	(0.557)
N(0,1)×Constant	6.953	5.539	6.067	0.314	1.747	0.501	0.141	0.594	0.078	1.864	5.008	4.722	8.284	12.090
	(0.737)	(0.758)	(0.718)	(0.380)	(0.458)	(0.205)	(0.451)	(0.365)	(0.578)	(0.427)	(1.318)	(0.717)	(0.370)	(1.291)
Product Characteristics														
Income×PC1	0.016	0.022	0.015	0.029	0.024	0.027	0.023	0.023	0.014	0.020	0.017	0.019	0.014	0.016
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Children×PC1	-0.122	-0.112	-0.119	-0.110	-0.083	-0.056	-0.079	-0.069	-0.087	-0.103	-0.114	-0.092	-0.077	-0.107
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Income×PC2	-0.018	-0.026	-0.025	-0.019	-0.012	-0.016	-0.012	-0.016	-0.013	-0.007	-0.002	-0.008	0.011	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Children×PC2	-0.011	-0.025	-0.025	-0.033	-0.027	-0.039	-0.031	-0.001	-0.011	0.020	-0.009	0.010	0.003	0.019
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Income×PC3	-0.027	-0.019	-0.007	-0.006	0.008	-0.005	0.003	0.015	-0.002	-0.017	-0.010	-0.005	-0.011	-0.011
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Children×PC3	-0.217	-0.226	-0.255	-0.223	-0.205	-0.154	-0.190	-0.166	-0.173	-0.179	-0.179	-0.195	-0.170	-0.211
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)

Panel B: Other Statistics

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Observations	15,441	16,336	16,604	16,791	17,241	17,329	16,444	16,213	16,443	15,829	15,487	14,365	18,850	17,805
Median Own Elasticity	3.488	2.068	2.633	1.844	1.871	1.749	2.083	2.196	2.350	2.252	2.396	3.021	2.308	3.085
Median Lerner	0.333	0.560	0.446	0.593	0.604	0.621	0.517	0.491	0.454	0.494	0.486	0.384	0.504	0.382

Notes: This table summarizes the results of estimation for the ready-to-eat cereals category for each year in the sample. Panel A provides the parameters and the standard errors, which are clustered at the region level and include a small-sample correction for the number of clusters. Panel B provides the number of product-chain-region-quarter observations, the median own price elasticity of demand, and the median Lerner index.

Figure D.2: Implied Elasticities Under Alternative Identification Restrictions

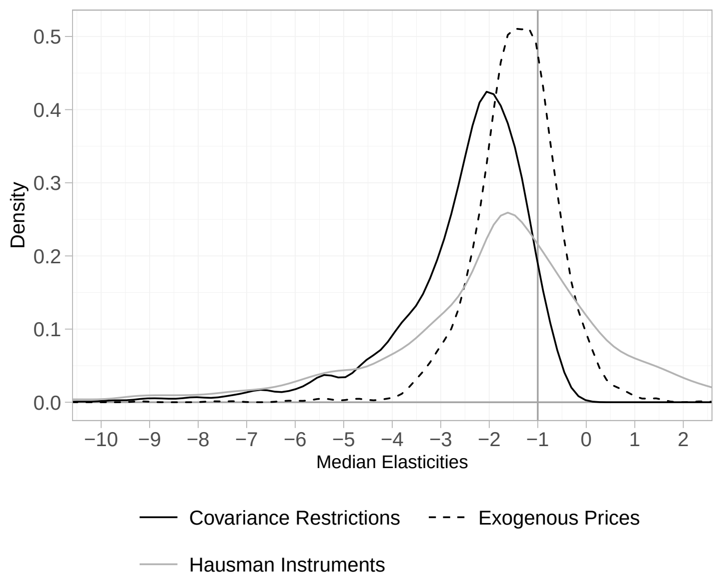

Notes: This figure plots the density of the median own-price elasticity by category and year under different identification assumptions. The solid black line shows the density of implied elasticities using covariance restrictions. The dashed line shows the density of implied elasticities assuming exogenous prices. The solid gray line shows the density of implied elasticities using Hausman instruments. The vertical line indicates an elasticity of -1.

D.3 Alternative Identification Strategies

For the third validation check, we examine the distribution of median own-price elasticities across all of the 1,862 category-year combinations in our baseline sample. We compare the results to those obtained under two alternative assumptions that can identify the price parameter and be applied at scale. The first alternative assumption is that prices are exogenous.

The second alternative approach to estimation uses instruments based on the average price of the same product in other regions (Hausman, 1996). This approach is valid if cost shocks are correlated across regions due to shared manufacturing or distribution facilities, for example, but demand shocks are uncorrelated across regions. These conditions may not be satisfied in many empirical settings. For example, validity can be threatened if firms employ region-wide or national advertising campaigns. Thus, Hausman instruments are at best subject to scrutiny when employed (Berry and Haile, 2021; Gandhi and Nevo, 2021).

Figure D.2 plots the densities of median own-price elasticities. The solid black line summarizes the results that we obtain with covariance restrictions (our baseline assumption). As shown, the peak of the distribution with covariance restrictions occurs at an elasticity slightly more negative than -2. Relative to our estimates, the distributions of elasticities with exogenous prices (the dashed line) and Hausman instruments (the solid gray line) are shifted to the right, yielding more inelastic demand overall.

Using covariance restrictions, demand is never upward-sloping, and only 5.5 percent of the category-year combinations have inelastic demand (i.e., a median elasticity greater than -1). By contrast, 28.8 percent of the category-year estimates exhibit inelastic demand with exogenous prices; with Hausman instruments, it is 34.1 percent. Furthermore, both of those approaches yield several estimates with upward-sloping demand. These results suggest the covariance restrictions approach generates reasonable demand elasticities, and that it is a distinctly good way to approach estimation in our context.

The differences in the distributions are consistent with price endogeneity arising from firms adjusting prices in response to demand shocks. Typically, firms will charge higher prices for larger demand shocks. This will show up as a bias term and lead to less elastic or even upward-sloping demand under the (misspecified) assumption of exogenous prices. Covariance restrictions systematically correct for this form of price endogeneity, yielding more elastic demand relative to those obtained under the assumption of exogenous prices. By contrast, Hausman instruments yield more elastic demand than exogenous prices in some cases and more inelastic demand in others.

D.4 Supply Model

Figure D.3: Changes in Demand Over Time

(a) Elasticity Trend (b) Price Sensitivity Trend

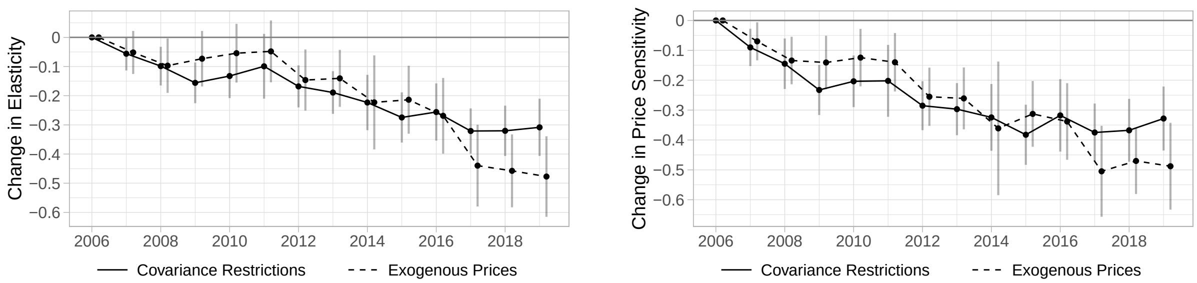

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of the log absolute value of the own-price elasticity (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a black line and employ covariance restrictions to estimate mean price parameters. The dashed line corresponds to estimates that instead employ an assumption that prices are exogenous.

We examine whether the estimated trends in demand, in terms of more inelastic demand and reduced price sensitivity, are robust to the supply model and the covariance restrictions that we invoke to identify the mean price parameter. As described in the text, the other demand-side parameters are identified by micro-moments. Thus, here we focus on the mean price parameter, which also has implications for the implied elasticities.

We show that a similar trend is obtained when we estimate demand using the assumption that prices are exogenous, which does not invoke the supply model to pin down the demand parameters. Though elasticity estimates under this approach are often unreasonable in terms of levels (see Section C), a change in the estimated parameters would be consistent with a rotation of the demand curve.

Figure D.3 shows that we find similar trends in elasticities (panel (a)) and the mean price parameter (panel (b)) under the assumption that prices are exogenous. This finding indicates that the reduced-form relationship between prices and quantities is becoming more "vertical" (on a price-quantity graph) over time, consistent with a rotation in the demand curve. The covariance restriction approach finds a similar trend while correcting for price endogeneity. The fact that the trends are similar suggests that our finding of reduced price sensitivity is not sensitive to the particular supply-side assumptions we invoke in estimation.[46]

[46] Of course, as indicated in the main text, a model of firm behavior is required to calculate markups and evaluate whether they are increasing. Regardless of whether firms actually exert market power, a finding of less elastic demand points to a increase in market power *potential*. We thank Chad Syverson for offering this interpretation.

E Robustness Checks

In this section, we present a series of alternative specifications and robustness checks to evaluate the sensitivity of our main findings to particular assumptions. First, we show how the main trend in markups is not sensitive to particular choices of measurement, in terms of which categories are included in our baseline sample and our choice of the Lerner index as our markup measure. We then show that the product-level trend in markups looks nearly identical with a balanced panel, confirming that the trend is not due to compositional shifts in products over time.

We then explore robustness to our sample of categories, retailers, geographic markets, and brands. We estimate our models on alternative samples where we consider all of the top 200 categories by revenues, additional large retail chains that are present only in the consumer panel data, more or fewer DMAs, and a different definition of fringe brands.

Further, we provide robustness checks that help address whether our findings are determined by the level at which our model is specified. We consider an alternative time aggregation and an alternative market definition. For the first alternative, we aggregate time periods to semiannual instead of quarterly observations. For the second, we specify markets at the DMA level across retailers instead of assuming that each retailer is active in a separate market. We also consider different approaches for determining market size.

In each of the above cases, we find similar trends in markups and price sensitivity to our baseline specification.

E.1 Markup Measure

Figure E.1: Markups Over Time: Price-Over-Cost Markups

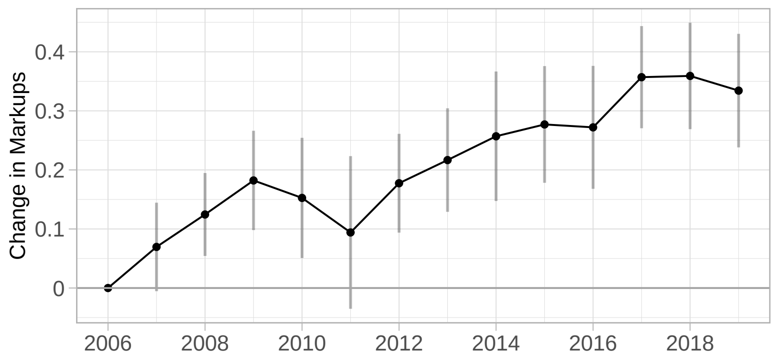

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. Markups are defined as price over marginal cost (p/c) as in De Loecker et al. (2020).

Throughout the paper, we use the Lerner index, $(p - c)/p$, as our measure of markups, which is a typical measure used in the industrial organization literature and in antitrust analysis (Elzinga and Mills, 2011). Other papers studying markups, particularly those in the macroeconomic literature, have used p/c, or price-over-cost markups (e.g., De Loecker et al., 2020). Both measures reflect the same fundamental relationship, but they are measured on different scales. The Lerner index is typically on $[0, 1]$, while price-over-cost markups are typically on $[1, \infty)$.

This distinction between the two does not matter for the trends we find in our analysis, which are typically reported in log changes. Figure E.1 replicates our product-level markup trends, corresponding to panel (a) of Figure 2 in the main text, using the price-over-cost markup measure. The trends are nearly identical.

E.2 Product and Retailer Composition

Figure E.2: Balanced Panel

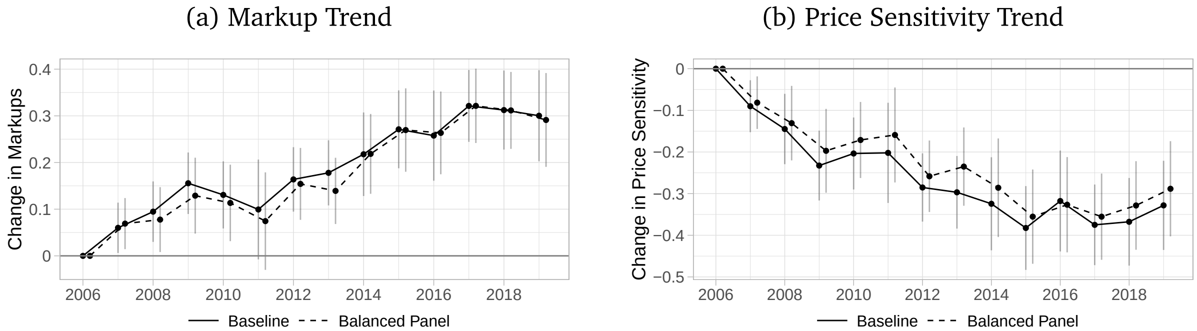

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a solid line. The dashed line corresponds to an alternative set of estimates from a panel that is balanced by brand×chain×region.

In our main specification, we use an unbalanced panel to maximize sample size and capture changes in aggregate markups due to entry and exit of products. As we discuss in section 3.1, some compositional changes in the NielsenIQ data occur during our sample period due to coverage of certain retail chains. Although our demand estimation controls for chain×region fixed effects, and these fixed effects can change with each year, a possible concern is that retail chains entering the sample may have different growth rates of markups.

In Figure E.2, we therefore replicate trends of markups and price sensitivity using a balanced panel of brand×chain×region combinations. The trends are similar to those reported in panel (a) of Figure 2 and panel (b) of Figure 3. The baseline trends are reproduced in the figure for comparison.

E.3 Sample of Categories

Figure E.3: Markups Over Time: Alternative Samples

(a) Markup Trend (b) Price Sensitivity Trend

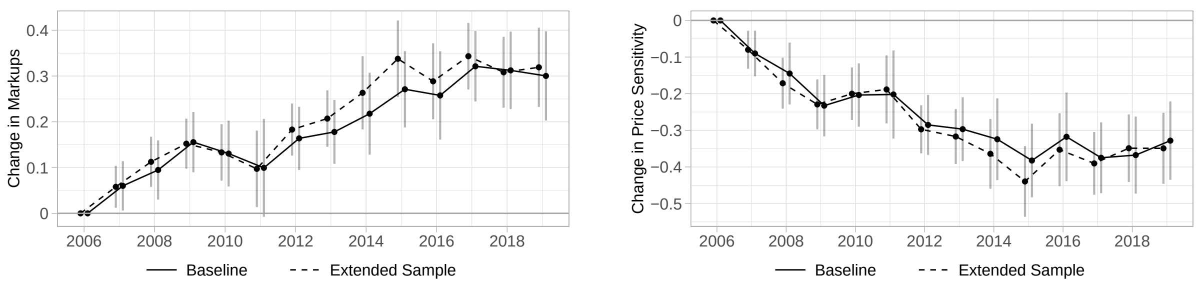

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a solid line. The dashed line corresponds to an alternative set of estimates from an extended sample (200 product categories).

In Section 3.1, we describe a category selection procedure in which we first choose the top 200 product categories by revenue, and then screen out categories with large values of within-category price dispersion. All of our baseline results are obtained with the 133 product categories that reflect that screen.

In Figure E.3, we replicate our product-level markup trends plot using an extended sample of all top 200 categories by revenue. The baseline trend is plotted for comparison. We find similar trends in markups with either selection procedure, with a change of approximately 30 log points from 2006 to 2019.

E.4 Sample of Retailers

Figure E.4: Including Additional Retail Chains

(a) Markup Trend (b) Price Sensitivity Trend

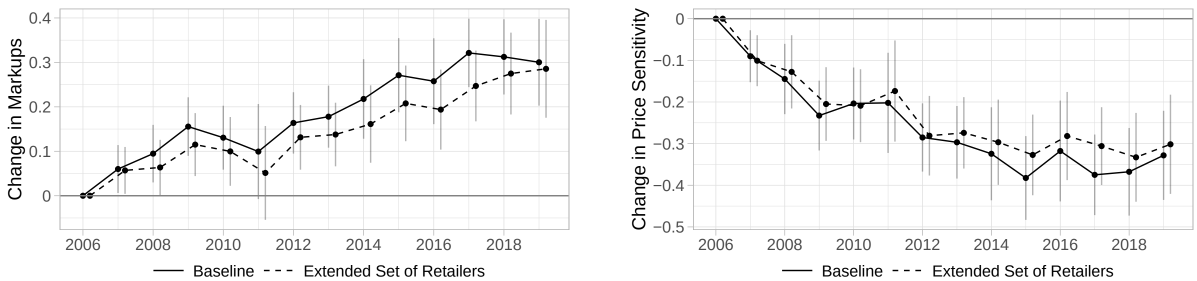

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-time level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The dashed line corresponds to estimates that additionally include data from two large retailers that are available in the consumer home scan panel data but not in the retailer scanner data.

Our baseline specification uses data from all retailers that are available in the retail scanner data. The advantage of the retail scanner data set is that all purchases in included stores are recorded which enables us to measure prices and quantities very precisely. A disadvantage of the retail scanner data set is that not all retail chains provide data. In contrast, the Consumer Panel Data includes information from all retailers. However, from this data, prices and quantities can be less precisely measured, as purchases are only available for a sample of consumers. For some products, prices have to measured from a few observations only, and some brand-DMA-retailer-time combinations are not observed. Therefore, our preferred specification uses data from the retail scanner data set. To check the sensitivity of our results towards the inclusion of retailers, we re-estimated our demand model on a sample that we complement with information on purchases from large retail chains that are available in the consumer level data but are not observed in the retail scanner data set. Specifically, we construct product-level price and quantity data for retailers with greater than a 5 percent revenue share in the consumer panel across all of our 133 product categories.[47] Figure E.4 shows that trends in markups and price sensitivity vary little with the inclusion of the additional retailers.

The added retailers account for a small share of observations but a disproportionately larger share of revenues. We have also re-run our estimation routine weighting each observation in the GMM objective function by log market size, which varies by retailer and DMA (see section A.2 for the calculation of market size). We obtain similar estimates to our baseline results.

[47]The added retailers have lower product-level prices on average, but there is no differential trend in prices relative to our baseline sample.

E.5 Sample of Geographic Markets

Figure E.5: Smaller and Larger Selection of DMAs

(a) Markup Trend (b) Price Sensitivity Trend

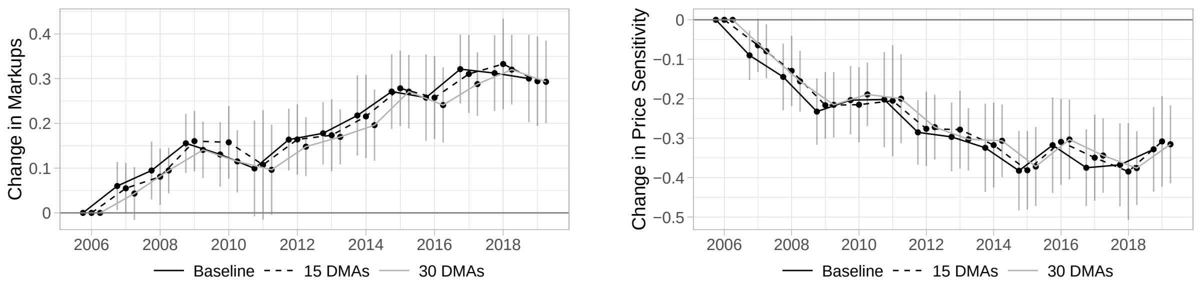

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) on year dummies controlling for year and unit-of-observation fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a solid black line and are derived from a sample of 22 DMAs. The dashed line corresponds to estimates based on 15 DMAs. The grey line corresponds to estimates based on an extended sample of 30 DMAs.

Our baseline specification uses data from 22 DMAs for which at least 500 panelists are available in the consumer level data in every year. To check the robustness of our results towards the selection of DMAs, we reran our demand model on a restricted sample of the 15 largest DMAs and an expanded sample of 30 DMAs. In the expanded sample, at least 500 panelist are available in each year except 2006. Relative to the baseline sample, the number of observations increase by 30 percent and total revenues increase by 17 percent with the expanded sample. Figure E.5 shows that trends in markups and price sensitivity are similar across estimation samples.

E.6 Sample of Brands

Figure E.6: Smaller number of top brands

(a) Markup Trend (b) Price Sensitivity Trend

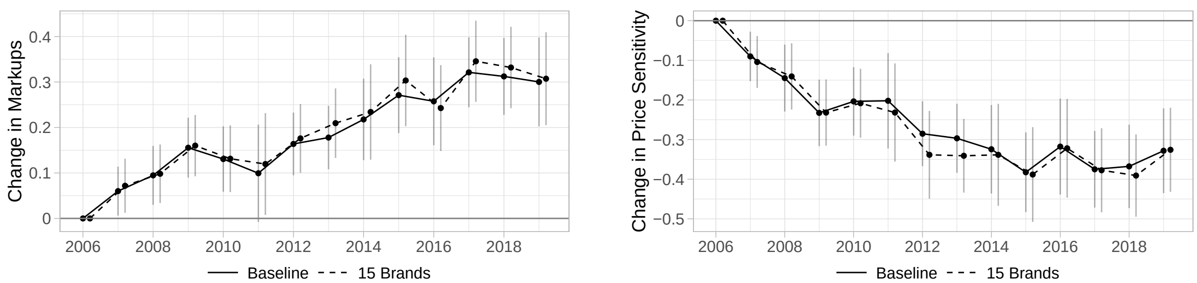

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a solid black line. The dashed line corresponds to estimates where we treat the top 15 instead of the top 20 brands as distinct product categories.

For our baseline specification, we treat the top 20 brands within each product category as distinct products and account for multi-product firm pricing among those brands. We aggregate the remaining brands into a fringe product that is assumed to be priced by an independent firm. To check whether our results are robust towards the definition of the fringe product, we reran our demand model treating only the top 15 brands as distinct products. Figure E.6 shows that trends in markups and price sensitivity are very similar to the baseline specification with the top 20 brands.

E.7 Time Aggregation

Figure E.7: Alternative Aggregation of Time Periods

(a) Markup Trend (b) Price Sensitivity Trend

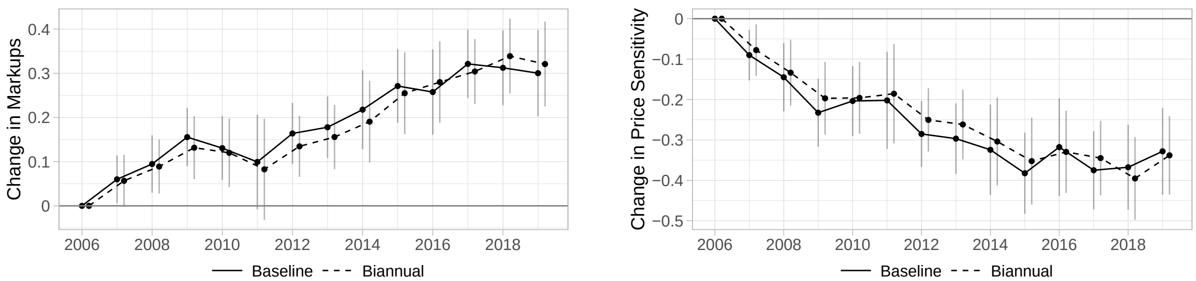

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-time level on year dummies controlling for product-chain-DMA and time period fixed effects. The year 2006 is the base category. The baseline estimates are derived from quarterly data and are plotted with a solid black line. The dashed line corresponds to estimates that use data aggregated to semiannual observations. Time period dummies correspond to quarterly and six-month periods, respectively.

Our baseline specification uses data aggregated to the quarterly data. Time aggregation involves a trade-off between the number of observations that are used to identify demand parameters and the sensitivity to short-run fluctuations induced, for instance, by temporary sales. To check the robustness of our results towards aggregation of time periods, we reran our demand model with semiannual data. Figure E.7 shows that this alternative aggregation leads to very similar trends in estimated markups and price sensitivity.

E.8 Market Definition

Figure E.8: Alternative Market Definition

(a) Markup Trend

(b) Price Sensitivity Trend

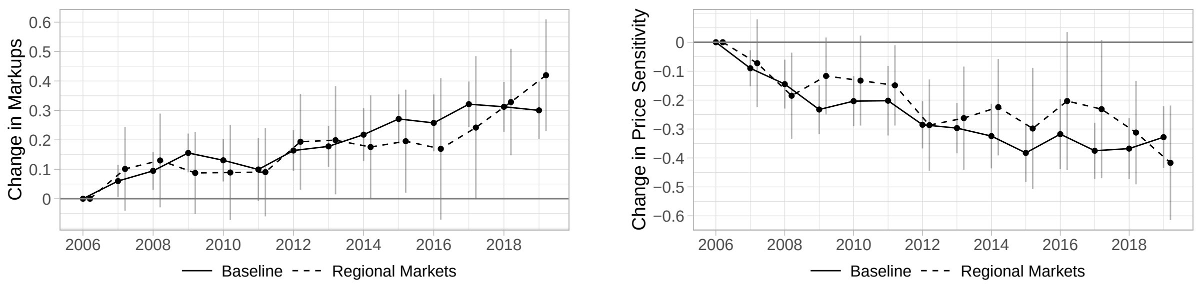

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) on year dummies controlling for year and unit-of-observation fixed effects. The year 2006 is the base category. In the baseline specification, the unit of observation is a product-chain-DMA combination. In the regional markets specification, a unit of observation is a product-DMA observation. The baseline estimates are plotted with a solid black line. The gray line corresponds to estimates using the alternative market definition which aggregates across retailers.

In our baseline specification, we assume that consumers are affiliated with a single retail chain and define markets at the product category-DMA-retailer-time level. Although this choice is consistent with the previous literature (e.g., Backus et al., 2021), a potential concern is that changes in consumers' search effort across retail chains affects our estimates. To check the robustness of our results towards this market definition, we reran our demand model based on a regional market definition, i.e., a product category-DMA-time combination. Figure E.8 shows that this alternative market definition leads to similar trends in markups and price sensitivity. The coefficients obtained from the alternative market definition are somewhat lower in most years, although the differences are not statistically significant, and less precisely estimated.

E.9 Market Size

Figure E.9: Smaller and Larger Market Size

(a) Markup Trend

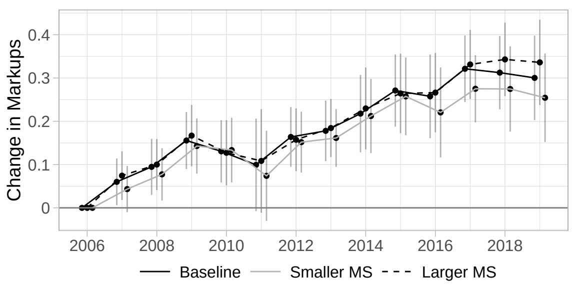

(b) Price Sensitivity Trend

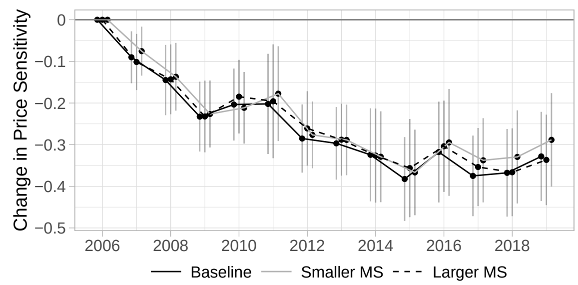

(c) Markup Trend

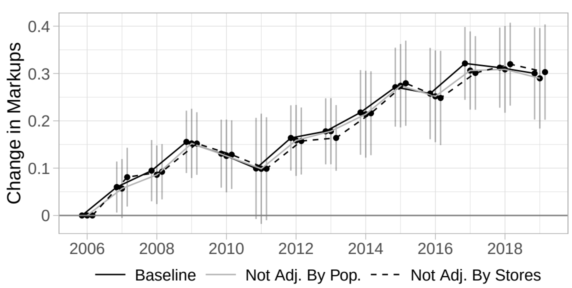

(d) Price Sensitivity Trend

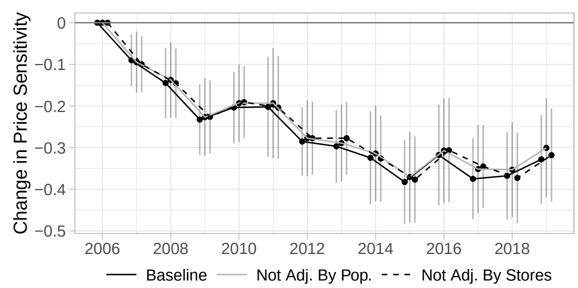

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panels (a) and (c)) and price sensitivity (panels (b) and (d)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a solid black line. In panels (a) and (b), the grey and dashed lines correspond to estimates that use alternative values for the average market size. Smaller (larger) market size refers to a specification where we rescale market size such that the average combined market share of inside goods equals 0.6 (0.3). In panels (c) and (d), the gray line corresponds to estimates using an alternative market size calculation that does not vary with population over time, and the dashed line corresponds to estimates using an alternative market size calculation that does not vary with the number of stores.

As discussed in Section 2.2, we need an assumption about market size to measure market shares of products. In Appendix A.2, we describe how we scale market size to obtain an average market share of inside goods of 0.45 and market growth that varies with the growth of population at the regional level and the number of stores at the region-chain level.

To check the robustness of our results towards assumptions about the relevant market, we reran our demand estimation using several alternative definitions of market size. First, we rescale market size to obtain an average combined market share of inside goods of either 0.3 or 0.6, which are smaller and larger than our baseline target value of 0.45. Second, we assume that market size does not vary with population growth; and third, we assume that market size does not change with the number of stores each retail chain has in a DMA. Figure E.9 shows that these alternative assumptions lead to similar trends in markups and price sensitivity. Thus, the trends we estimate do not hinge on the precise definition of market size.

F Exploring Alternative Mechanisms

Given the important role of price sensitivity in markups, we next examine potential factors that could explain the change over time. In the main text, we provide evidence that consumers are becoming less price sensitive over time due to exogenous factors (Section 5). In this appendix, we consider whether this change could reflect shifts in purchases across and within retail channels or whether this change may be due to firm-level investments that affect consumer behavior, such as increased marketing or product variety.

F.1 Consumer Spending Across Retail Channels

Purchases of the consumer products that are in our data primarily come from five retail channels, which NielsenIQ classifies as mass merchandisers, grocery stores, drug stores, warehouse clubs, and dollar stores. We refer to retailers in these channels as *broad-basket* retailers, indicating the broad assortment of product categories they sell. To provide context about aggregate spending on consumer products and the relative size of these channels, we use auxiliary data on retailer revenues for large U.S. retailers.

Specifically, we obtain retailer-level revenue data for the largest 100 U.S. retailers. The data are compiled annually by the National Retail Federation, which is the largest retail trade association. The earliest estimates we can find are from 2007, one year after the start of our sample. For 2007 and 2019, we categorize each retailer into one of the following types: mass merchandisers, grocery stores, drug stores, warehouse clubs, dollar stores, and other consumer product stores. Other consumer product stores include convenience stores, department stores, online retailers, and retailers that specialize in a more narrow set of categories (e.g., electronics, beauty, or apparel).[48] We also identify retailers that are restaurants, home improvement stores, and auto parts stores, and we drop these from the analysis because they do not primarily sell consumer products. Because the included retailers also sell products outside of the scope of our analysis (e.g., prescription drugs), the aggregate data may not provide an exact picture of how the retail shares of consumer products evolve over time. Nonetheless, we think the auxiliary data provide useful information. The included retailers represent \$1.4 trillion in revenues in 2007 and \$2.0 trillion in 2019.

Table F.1 reports the share of consumer product spending by broad-basket retail channels in 2007 and 2019. These shares in each of these channels have been fairly stable over our sample period. By these calculations, broad-basket retailers account for 63 percent of consumer product spending in 2007 and 67 percent in 2019. Within other consumer product channels,

[48]For Walmart, we adjust the provided estimates to separate Walmart U.S. (mass merchandiser) and Sam's Club (warehouse club) into distinct channels. For Amazon, we adjust the provided estimates in 2019 to include revenues from online sales and third-party seller services in the United States (other), and we separate out Whole Foods (grocery). We use data from Statista for Walmart (https://www.statista.com/statistics/269403/net-sales-of-walmart-worldwide-by-division/), and we obtain 2019 Amazon estimates from Amazon's 2021 10-K filing.

Table F.1: Share of Revenue by Retail Channel

	2007	2019
Broad-Basket Retail Channels		
Mass Merchandisers	0.214	0.218
Grocery Stores	0.219	0.217
Drug Stores	0.088	0.117
Warehouse Club	0.090	0.094
Dollar Stores	0.015	0.026
Other Consumer Product Retail Channels		
Convenience Stores, Department Stores, Apparel, etc.	0.374	0.328

Notes: This table displays the share of revenues of broad-basket retailers out of all consumer product spending. We compare broad-basket retailers to "specialized" retailers such as convenience stores, department stores, apparel stores, beauty stores, electronic stores, and online retailers. To construct these estimates, we take the revenues of the largest 100 U.S. retailers. We exclude from this list retailers that do not have consumer products as their primary source of revenue: restaurants, home improvement stores, and auto parts stores. The included retailers represent $1.4 trillion in revenues in 2007 and $2.0 trillion in 2019.

online retailers grew substantially, reaching roughly 6 percent of revenues in 2019. However, this increase was offset by relative declines in other store formats, such as department stores and apparel.

These data indicate that there are no broad shifts in consumer spending across the channels in our data during our sample period.[49] As noted in Section 3.1, the retailers in the Retail Scanner Data are disproportionately sampled from the first three broad-basket channels: mass merchandisers, grocery stores, and drug stores. Combined, these channels represent 52 percent of consumer product spending in 2007 and 55 percent in 2019, and the channel shares are fairly stable over time. Thus, the revenue growth in these channels has paralleled the average revenue growth among other large U.S. retailers.

F.2 Analysis of Potential Shifts within Data Sample

Building on the previous analysis, we assess changes the share of revenues across retail channels using the Kilts NielsenIQ Consumer Panel Data for the 133 categories in our baseline sample. For each product category and year, we include revenues for the five broad basket retailers from the previous section (mass merchandisers, grocery, drug stores, warehouse clubs, and dollar stores), as well as online retail. Using these data, we obtain qualitatively patterns in channel shares to the auxiliary data presented in Appendix F.1. Warehouse clubs, dollar stores, and online retailers are undersampled in the Retail Scanner Data. These channels realize relatively

[49] The revenue share of dollar stores roughly doubles between 2007 and 2019, consistent with the trend documented in Caoui et al. (2023). Nonetheless, dollars stores account for only 1.5 percent of consumer product spending in 2007 and 2.6 percent in 2019.

small growth in shares over this period. The average cross-category share in 2019 was 12.0 percent for warehouse clubs, 2.2 percent for dollar stores, and 1.9 percent for online retailers. In 2006, these values were 12.2 percent, 1.3 percent, and 0.5 percent, respectively. Consistent with the findings in Appendix F.1, among the six channels, mass merchandisers, grocery stores, and drug stores capture 86.0 percent share on average in 2006 and 83.9 percent in 2019. Thus, the aggregate compositional shifts in these channels are fairly small for the product categories we study.

Further, we do not find evidence that shifts in consumer spending to retailers outside of our price/quantity data is driving our results. The portion of expenditures in the Consumer Panel Data that are captured by retailers in the Retail Scanner Data is flat from 2006 to 2013, when price sensitivity is falling. In part due to changes in the composition of participating retailers, this portion is lower from 2014 to 2017 and higher in 2018 and 2019. To address the potential for the sample composition to impact our findings, we perform a robustness check with a balanced panel of retailers in Appendix E.2. We perform another robustness check in which we supplement our baseline sample with large retailers that are in the Consumer Panel Data but not in the Retail Scanner Data, which we discuss in Appendix E.4. In both cases, we find very similar trends in markups and price sensitivity.

Finally, our estimated demand parameters provide some evidence that selection over time into different types of retailers may not be driving the trend in price sensitivity we observe. Specifically, we find no trend over time in the coefficients that load onto the interaction of price and household income (Figure G.3). This indicates that, based on income, there is no disproportionate selection of greater price sensitive consumers to retailers outside of our sample.[50]

Taken together, we think it is unlikely that compositional shifts would account for the 30 percent decline in price sensitivity we estimate over this period. Nonetheless, we explore this further with a regression analysis that exploits panel variation in subsection F.3. Some categories are disproportionately affected by the growth of alternative retail channels. For example, less than one percent of beer was sold online in each year of the sample, whereas the share of online revenues for dry dog food increased from less than 2 percent to over 15 percent during the sample period. If we see a greater decrease in price sensitivity for categories disproportionately affected by the shift to online, that might suggest that consumer selection may be playing some role.

F.3 Category-Level Variation

We also investigate whether firm-level investments may yield consumers that are less price sensitive, either through perceived or realized changes to their products. To explore this, we merge our estimates with financial data on marketing and R&D expenses obtained from Compustat.

[50]The random coefficients model endogenizes the consumer's decision to buy from the retailers in our sample, so we are also able to control for some types of selection directly with the model.

Table F.2: Potential Mechanisms

	(1) Price Sensitivity	(2) Log Abs. Elasticity	(3) Marginal Cost	(4) Perceived Quality
Log Share Warehouse Clubs	-0.038	-0.013	0.150	-0.053
	(0.066)	(0.061)	(0.188)	(0.150)
Log Share Dollar Stores	0.059^{**}	0.060^{**}	0.062	0.079
	(0.027)	(0.026)	(0.075)	(0.088)
Log Share Online	-0.078	-0.056	-0.133	-0.412^{***}
	(0.047)	(0.044)	(0.136)	(0.144)
Log Marketing Spend	0.007	0.015	0.118^{**}	0.041
	(0.021)	(0.020)	(0.056)	(0.058)
Log R&D	-0.006	-0.006	-0.068	0.023
	(0.025)	(0.022)	(0.061)	(0.084)
Log Num. UPCs	0.103^{**}	0.092^{**}	0.426^{***}	0.473^{***}
	(0.051)	(0.046)	(0.124)	(0.152)
Brand-Category FEs	X	X	X	X
Time Period FEs	X	X	X	X
Observations	1,799	1,799	1,799	1,799
R^2	0.942	0.615	0.132	0.173
R^2 (Within)	0.014	0.013	0.017	0.027

Notes: Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These measures are obtained from annual reports of the parent companies. We also consider whether changes in product variety may account for the changes we observe. We measure product variety as the (log) number of UPCs offered by each brand in each market. We aggregate our data to the category-year level, taking a simple average of each measure. Thus, we seek to evaluate whether categories with disproportional increases in marketing, R&D, or variety also realized greater declines in price sensitivity.

To explore these relationships, we regress price sensitivity ($\ln(-\alpha_t)$) on the logged values of the above measures. We include category fixed effects and year dummies, so that the coefficients reflect time-series variation within each category that departs from the aggregate trend. Following the discussion in Appendix F.2, we also include category-level values of (log) expenditure shares at warehouse clubs, dollar stores, and online, allowing us to assess potential relationships with category-specific trends across retail channels.

Column (1) of Table F.2 reports the results. We find no significant relationships between share sold in warehouse clubs, marketing expenditures, or R&D expenditures. We find a negative, statistically insignificant relationship between the share sold online and consumer price sensitivity, and a positive, statistically significant relationship between share sold in dollar stores and price sensitivity. Given the coefficient magnitudes and the absolute size of these channels (shares of less than 2.5 percent in 2019), we think these results most likely reflect other mechanisms, e.g., online retailers entering categories with less price sensitive consumers. In support of other mechanisms, a regression with price elasticity as the dependent variable, reported in column (2), returns a coefficient on online sales that is roughly 25 percent smaller. If online sales were skimming off more price sensitive consumers, we would expect elasticities to have a stronger relationship with online sales than the (mean) price sensitivity parameter, as the elasticity also incorporates self-selection based on demographic characteristics (e.g., lower-income consumers). We do not find evidence for this selection. Likewise, the point estimates for share at dollar stores is similar in columns (1) and (2).

We find a significantly positive relationship between variety and price sensitivity, which indicates that greater variety is weakly correlated with *greater* price sensitivity.[51] Since price sensitivity has decreased over time while variety has increased, we think it is likely that this coefficient reflects other factors. Together, all five measures only explain 1.4 percent of the residual variation in price sensitivity, suggesting that neither retail shopping patterns nor firm-level investments are driving the changes in price sensitivity over time.

Though we focus on explaining price sensitivity, we also run regressions with marginal costs and perceived quality as the dependent variables. We report results in columns (3) and (4). We find a positive and significant relationship with marginal costs and marketing, suggesting that cost decreases were also correlated with less spending on marketing. We also find a large and highly significant negative relationship between perceived quality and online sales. As perceived quality captures the value to consumers above and beyond outside options (including online sales), this is consistent with the trends we find in Section 4. Online retail became an increasingly popular option over the time period, lowering the (relative) utility of in-store purchases. Conversely, we find no effect of warehouse clubs on perceived quality, though the point estimate is negative.

We find that product variety is positively correlated with marginal costs and perceived quality. As both marginal costs and quality are falling over time, while variety is rising, this suggests that greater variety may have helped to mitigate the substitution of consumers to other channels (i.e., online), albeit at higher costs.[52]

Overall, this analysis suggests that firm-level investments and changes in the composition of retail shopping across channels cannot account for the change in consumer price sensitivity that we document.

[51] Brand (2021) finds the opposite relationship.

[52] This is related to the explanation offered by Brand (2021), who suggests that increased variety may lead to less price sensitivity. However, we do not find that increases in variety are related to lower price sensitivity, and we do not find that changes in quality, which are correlated with variety, drive changes in markups. In the time series, quality declines over time, and we estimate a net relationship with markups very close to zero when controlling for other factors (Table 2). Thus, product variety does not appear to be driving the trends we observe.

G Additional Figures and Tables

Figure G.1: Changes in the Distribution of Markups

(a) Absolute Change

(b) Relative Change

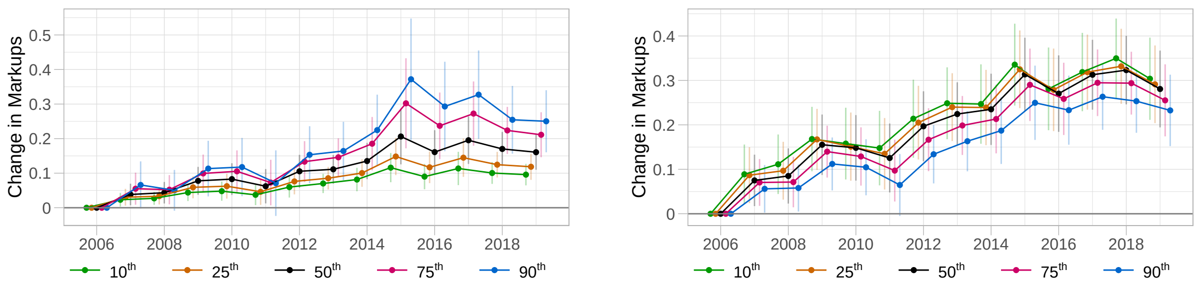

Notes: This figure shows coefficients and 95 percent confidence intervals of regressions of percentiles of the markup distribution at the product category level on year dummies using the year 2006 as the base category. In panel (a), outcomes are percentiles of the level of the Lerner index, $(p - c)/p$, in panel (b), outcomes are measured in logarithms.

Figure G.2: Product-Level Changes in Nominal Prices and Marginal Costs

(a) Prices

(b) Marginal Costs

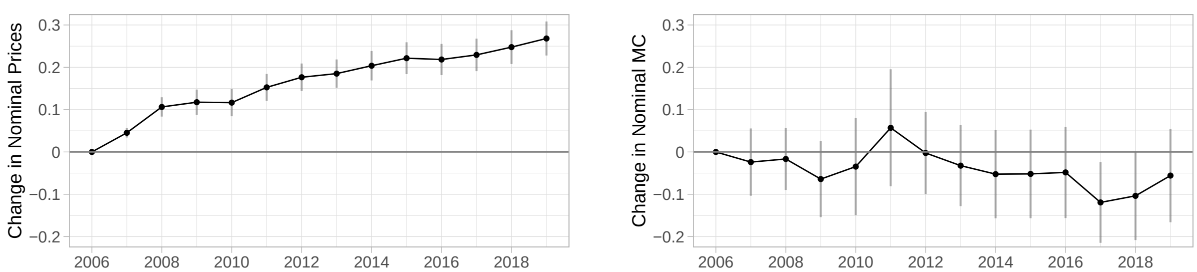

Notes: This figure shows coefficients and 95 percent confidence intervals of regressions of the log of nominal prices and marginal costs at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.

Figure G.3: Changes in Demand Parameters

(a) Relative Quality

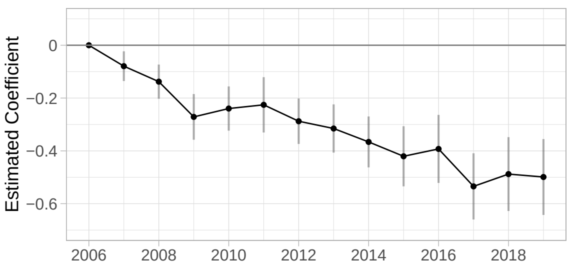

(b) Random Coefficient: $N(0,1)\times$ Intercept

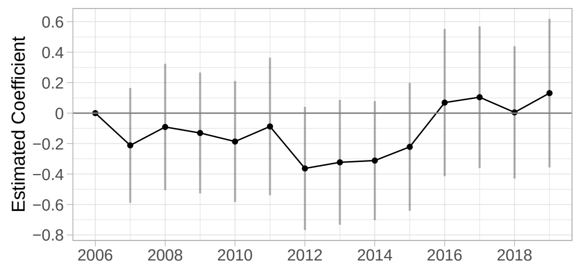

(c) Income $\times$ Price Coefficient

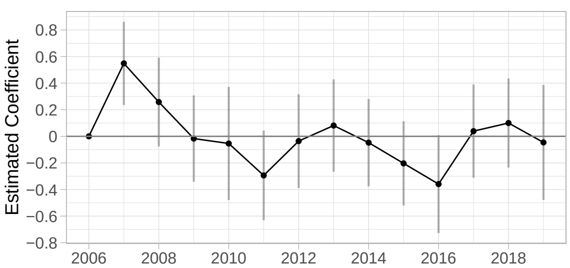

(d) Income $\times$ Intercept Coefficient

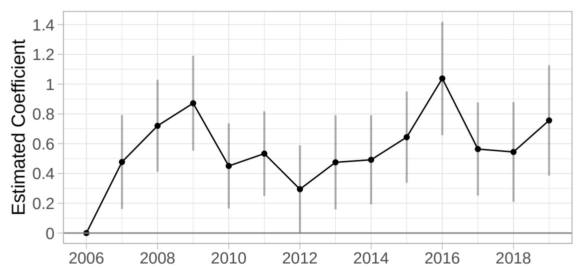

(e) Children $\times$ Price Coefficient

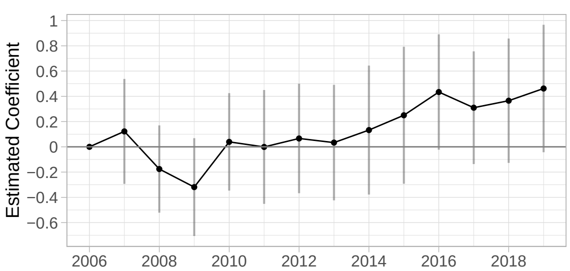

(f) Children $\times$ Intercept Coefficient

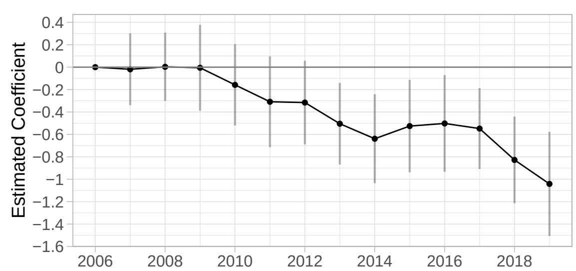

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of standardized demand parameters on year dummies controlling for product-chain-DMA and quarter fixed effects. Observations are at the product-chain-DMA-quarter-year level. The year 2006 is the base category.

Figure G.4: Substitution Among Products and to the Outside Good

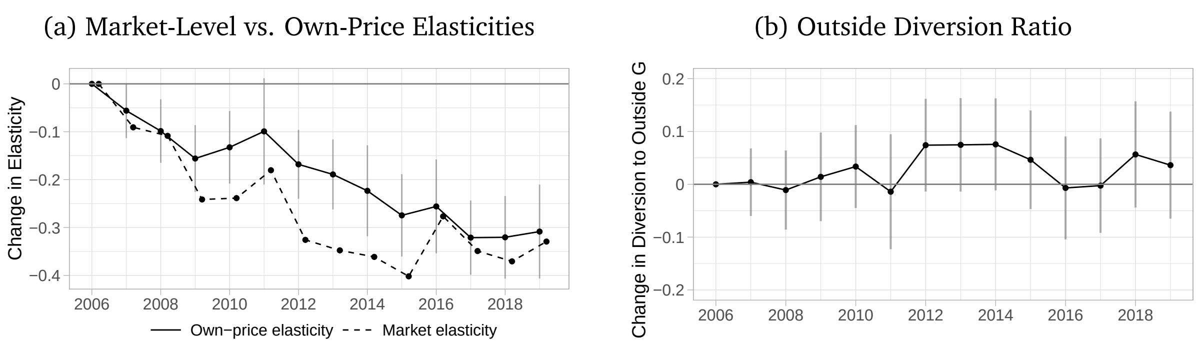

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of own-price and market-level elasticities (panel (a)) and diversion ratios (panel (b)) on year dummies controlling for product-chain-DMA and quarter fixed effects. Market-level elasticities are calculated as the change in log total quantities over the change in log average prices. Observations are at the product-chain-DMA-quarter-year level. The year 2006 is the base category.

Figure G.5: Changes in Market Concentration

(a) Parent HHI

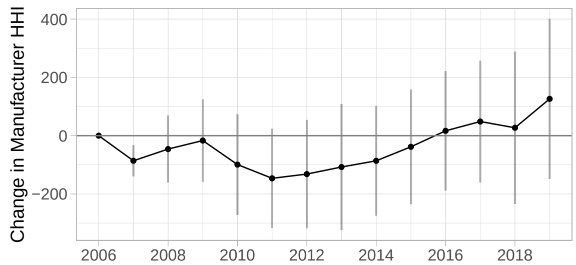

(b) Brand HHI

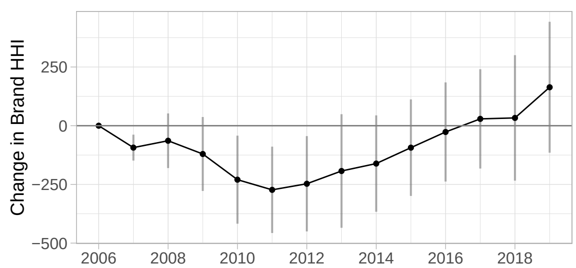

(c) Retailer HHI

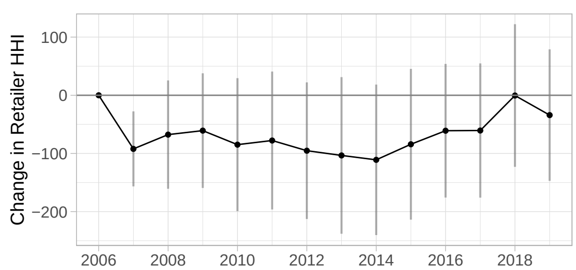

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of HHI measures on year dummies controlling for product-chain-DMA and quarter fixed effects, with 2006 as the base category. We measure HHI as the sum of squared market shares, where we first adjust market shares so that inside shares sum to one. For this figure, HHI is measured on a 0 to 10,000 scale. Observations are at the product-chain-DMA-quarter-year level.

Figure G.6: Consumer Surplus Over Time By Income Group, Deciles

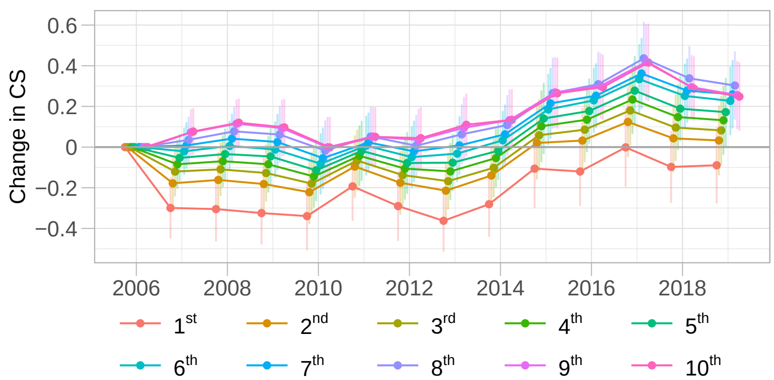

Notes: This figure reports coefficients and 95 percent confidence intervals of a regression of the log of consumer surplus by purchase on year dummies, controlling for category fixed effects, separately for different deciles of the income distribution.

Figure G.7: Contribution of Covariance Restriction to Objective Function: Ready-to-Eat Cereals

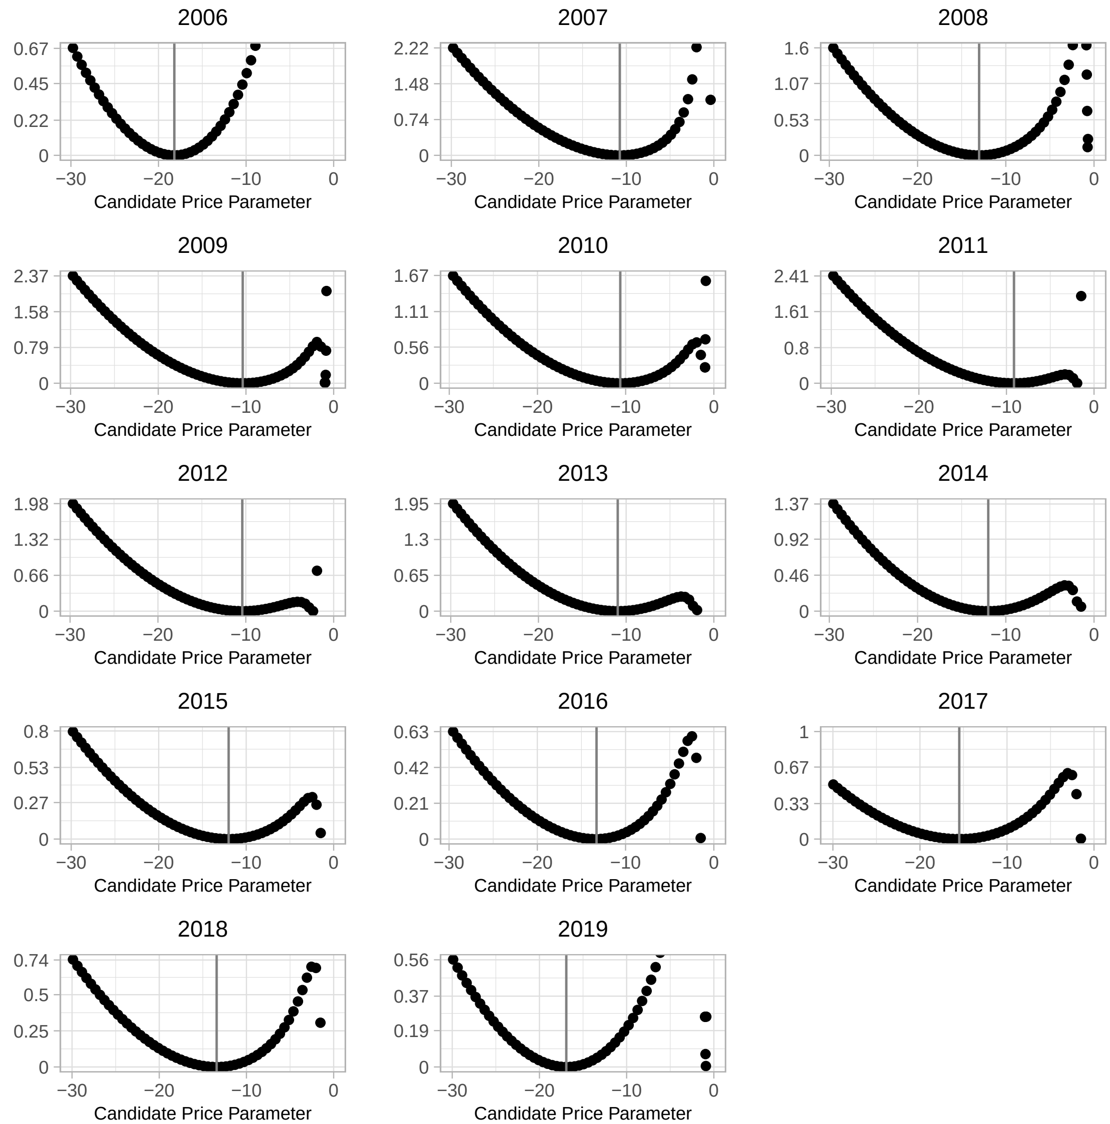

Notes: This figure plots the contribution of the covariance restriction to the objective function, scaled by ten thousand, for different candidate price parameters over the range $[-30, 0]$. Other parameters are held fixed at the levels obtained in the first step of estimation.

Table G.1: Product-Level Markups Over Time, Sales-Weighted Regressions

	(1) Log Markup	(2) Log Markup	(3) Log Markup	(4) Log Markup	(5) Log Markup	(6) Log Markup
Trend	0.019***		0.018***		0.023***	
	(0.003)		(0.003)		(0.003)	
Year=2007		0.057**		0.059**		0.060**
		(0.027)		(0.027)		(0.027)
Year=2008		0.094***		0.095***		0.095***
		(0.032)		(0.032)		(0.033)
Year=2009		0.161***		0.160***		0.156***
		(0.033)		(0.033)		(0.033)
Year=2010		0.136***		0.134***		0.131***
		(0.036)		(0.036)		(0.036)
Year=2011		0.101*		0.098*		0.099*
		(0.053)		(0.053)		(0.054)
Year=2012		0.165***		0.158***		0.164***
		(0.034)		(0.035)		(0.035)
Year=2013		0.181***		0.170***		0.178***
		(0.035)		(0.036)		(0.035)
Year=2014		0.213***		0.203***		0.218***
		(0.045)		(0.046)		(0.045)
Year=2015		0.255***		0.242***		0.271***
		(0.043)		(0.042)		(0.042)
Year=2016		0.231***		0.216***		0.258***
		(0.049)		(0.049)		(0.049)
Year=2017		0.289***		0.275***		0.321***
		(0.039)		(0.038)		(0.039)
Year=2018		0.267***		0.265***		0.312***
		(0.045)		(0.043)		(0.043)
Year=2019		0.246***		0.245***		0.300***
		(0.052)		(0.049)		(0.049)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	14,406,731	14,406,731	14,406,731	14,406,731	14,406,731	14,406,731
R^2	0.013	0.016	0.359	0.361	0.782	0.783

Notes: Dependent variable is the log of the Lerner index. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.2: Product-Level Markups Over Time, Unweighted Regressions

	(1) Log Markup	(2) Log Markup	(3) Log Markup	(4) Log Markup	(5) Log Markup	(6) Log Markup
Trend	0.018***		0.020***		0.023***	
	(0.003)		(0.003)		(0.003)	
Year=2007		0.072**		0.070**		0.078***
		(0.030)		(0.030)		(0.030)
Year=2008		0.092***		0.090***		0.106***
		(0.033)		(0.033)		(0.034)
Year=2009		0.143***		0.140***		0.161***
		(0.038)		(0.037)		(0.038)
Year=2010		0.139***		0.136***		0.158***
		(0.041)		(0.041)		(0.042)
Year=2011		0.094**		0.091**		0.116***
		(0.044)		(0.043)		(0.044)
Year=2012		0.168***		0.165***		0.194***
		(0.039)		(0.039)		(0.039)
Year=2013		0.184***		0.179***		0.207***
		(0.036)		(0.035)		(0.035)
Year=2014		0.217***		0.215***		0.245***
		(0.043)		(0.042)		(0.043)
Year=2015		0.303***		0.307***		0.342***
		(0.051)		(0.049)		(0.051)
Year=2016		0.248***		0.253***		0.288***
		(0.048)		(0.046)		(0.048)
Year=2017		0.293***		0.303***		0.340***
		(0.045)		(0.043)		(0.045)
Year=2018		0.251***		0.267***		0.304***
		(0.043)		(0.042)		(0.043)
Year=2019		0.219***		0.238***		0.276***
		(0.044)		(0.043)		(0.044)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	14,406,731	14,406,731	14,406,731	14,406,731	14,406,731	14,406,731
R^2	0.011	0.014	0.352	0.355	0.758	0.761

Notes: Dependent variable is the log of the Lerner index. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.3: Product-Level Markups Over Time, Balanced Panel, Sales-Weighted Regressions

	(1) Log Markup	(2) Log Markup	(3) Log Markup	(4) Log Markup	(5) Log Markup	(6) Log Markup
Trend	0.022***		0.020***		0.024***	
	(0.003)		(0.003)		(0.003)	
Year=2007		0.060**		0.060**		0.062**
		(0.029)		(0.028)		(0.028)
Year=2008		0.105***		0.104***		0.102***
		(0.035)		(0.035)		(0.035)
Year=2009		0.167***		0.165***		0.156***
		(0.035)		(0.035)		(0.036)
Year=2010		0.147***		0.143***		0.133***
		(0.038)		(0.038)		(0.038)
Year=2011		0.119**		0.114*		0.105*
		(0.058)		(0.058)		(0.059)
Year=2012		0.187***		0.177***		0.170***
		(0.034)		(0.036)		(0.035)
Year=2013		0.195***		0.184***		0.179***
		(0.036)		(0.038)		(0.037)
Year=2014		0.236***		0.223***		0.223***
		(0.047)		(0.049)		(0.048)
Year=2015		0.276***		0.261***		0.270***
		(0.043)		(0.044)		(0.043)
Year=2016		0.260***		0.245***		0.264***
		(0.051)		(0.052)		(0.051)
Year=2017		0.315***		0.298***		0.326***
		(0.039)		(0.040)		(0.040)
Year=2018		0.299***		0.282***		0.312***
		(0.046)		(0.045)		(0.045)
Year=2019		0.302***		0.285***		0.319***
		(0.050)		(0.049)		(0.050)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	4,810,064	4,810,064	4,810,064	4,810,064	4,810,064	4,810,064
R^2	0.020	0.021	0.401	0.402	0.767	0.768

Notes: Dependent variable is the log of the Lerner index. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.4: Factors Predicting Cross-Category Variation in Markup Trends (Category Level)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Marginal Cost (Standardized)	-0.245^{***}					-0.133^{***}	-0.131^{***}
	(0.012)					(0.010)	(0.011)
Price Sensitivity		-0.688^{***}				-0.443^{***}	-0.443^{***}
		(0.032)				(0.049)	(0.048)
Quality (Standardized)			-0.207^{***}			0.002	0.002
			(0.013)			(0.009)	(0.009)
Income (Log)				0.690		-0.019	-0.005
				(2.388)		(0.919)	(0.968)
Children at Home				-8.502		-5.904^{*}	-6.188^{*}
				(6.166)		(3.338)	(3.145)
Parent HHI					1.365^{**}		0.725^{**}
					(0.588)		(0.312)
Brand HHI					-0.246		0.049
					(0.301)		(0.137)
Retailer HHI					1.596^{*}		0.263
					(0.815)		(0.359)
Category FEs	X	X	X	X	X	X	X
Year FEs	X	X	X	X	X	X	X
Observations	1,861	1,861	1,861	1,861	1,861	1,861	1,861
R^2 (Within)	0.644	0.691	0.450	0.002	0.025	0.791	0.800

Notes: Dependent variable is the log of the mean Lerner index within a category-year. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.5: Price Sensitivity and Markups Across Product Categories

	(1) 2006 Log $\overline{L}$	(2) 2017 Log $\overline{L}$	(3) 2019 Log $\overline{L}$	(4) Δ Log $\overline{L}$
Price Sensitivity	-0.138^{***}	-0.180^{***}	-0.122^{***}	
	(0.027)	(0.028)	(0.031)	
Δ Price Sensitivity				-0.590^{***}
				(0.012)
Observations	133	133	133	1,729
R^2	0.164	0.234	0.104	0.593

Notes: This table reports regression results that examine the cross-sectional and time series relationships of price sensitivity and markups, as measured by the log aggregate Lerner index at the category-year level. The regressions are motivated by the decomposition in equation (12). All regressions include a constant. Columns (1), (2), and (3) capture cross-sectional variation using the years 2006, 2017, and 2019 for the 133 product categories in our baseline sample. We include 2006 and 2019 because they are the first and last years of the sample, and we include 2017 due to more substantial compositional changes in the Retail Scanner Data in 2018–2019, as discussed in Appendix A. Variation in price sensitivity explains a modest fraction of the cross-sectional variation in markups: 16 percent in 2006, 27 percent in 2017, and 7 percent in 2019. Consistent with changes in markups due to price sensitivity, the R^2 in 2017 is higher than that of 2006; the lower R^2 is 2019 may be attributable to the compositional shift in the scanner data. Column (4) captures the time series variation by estimating the model in first differences from 2007 through 2019. Changes in price sensitivity over time explain 57 percent of the category-level variation in markups over time. Overall, the results indicate that variation in factors other than price sensitivity can explain a large portion in category-level variation in markups. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.