

Rising Markups and the Role of Consumer Preferences*

Hendrik Döpper

DICE[†]

Alexander MacKay

Harvard University[‡]

Nathan H. Miller

Georgetown University[§]

Joel Stiebale

DICE[¶]

October 5, 2022

Abstract

We characterize the evolution of markups for consumer products in the United States from 2006 to 2019. We use detailed data on prices and quantities for products in more than 100 distinct product categories to estimate demand systems with flexible consumer preferences. We recover markups under an assumption that firms set prices to maximize profit. Within each product category, we recover separate yearly estimates for consumer preferences and marginal costs. We find that markups increase by about 30 percent on average over the sample period. The change is attributable to decreases in marginal costs that are not passed through to consumers in the form of lower prices. 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, and Ariel Pakes for helpful comments. We thank seminar and conference participants at CRESSE, the Federal Trade Commission, Harvard University, Heinrich Heine University Düsseldorf, Indiana University, the NBER Conference on Megafirms, New York University, Northwestern University, RGS Doctoral Conference, University of California-Berkeley, University of Freiburg, University of Giessen, University of Maryland, University of Pennsylvania, University of Virginia, and Washington University in St. Louis. 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 Nielsen. Nielsen 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.

[†]Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf. Email: doeper@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—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, the growing empirical evidence that markups are rising in the United States and abroad (e.g., De Loecker et al., 2020; Ganapati, 2021a; 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 for a vast number of firms and products. Our objective is to understand the supply and demand conditions that influence firms' pricing decisions. Through an analysis of economic mechanisms, we are able to connect markups to other economic outcomes, such as consumer surplus and deadweight loss, and provide context for various policy considerations. For example, with no changes in demand, rising markups may arise from reduced competition (e.g., due to anticompetitive mergers) or from cost-reducing technological progress.¹ Alternatively, rising markups could reflect shifts in consumer preferences, rather than such supply-side changes.

Although measures of prices are often available, marginal costs are typically unobserved to the researcher. Hence, one must interpret the available data through the lens of economic theory to recover markups. Our approach is to estimate differentiated-products demand systems for more than 100 consumer product categories—such as cereals, shampoo, and over-the-counter cold medications—using prices, quantities, and consumer demographics. With demand estimates in hand, we impute the marginal costs and markups that rationalize prices under the assumption of profit maximization. We repeat this procedure separately for each year over 2006–2019. Our approach is standard in industrial organization (e.g., Berry et al., 1995), although most previous applications 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, 2020). We implement the methodology at scale to obtain markups for thousands of products, across categories, geographic regions, and over time.

We estimate that average markups increase by about 30 percent between 2006 and 2019, with the average Lerner index increasing from approximately 0.45 to 0.60.² 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

¹In environments with incomplete pass-through, cost reductions do not yield corresponding declines in price.

²The Lerner index is calculated as $\frac{p-c}{p}$, where p and c are price and marginal cost, respectively (Lerner, 1934). As long as marginal cost does not exceed price, it can take values from zero to one.

initial markups; however, in percentage terms, the changes that we estimate are similar for high- and low-markup products. Thus, we interpret our results as indicating that the full distribution of product-level markups may be shifting upward over time. Our findings of increasing average markups is consistent with the findings of De Loecker et al. (2020), despite using a different methodology (supply and demand) and data (prices and quantities).

Our paper makes at least three distinct contributions. First, we use models of supply and demand to evaluate changes to markups over time and potential causes, including changes in costs, concentration, and consumer preferences. Second, we identify a secular decline in price sensitivity for consumer products, which is a key driver of the increasing markups we observe. Using auxiliary data, we document that this trend corresponds to a decline in coupon use and time spent shopping. Third, our flexible demand modeling approach allows us to evaluate the implications for consumer welfare across the income distribution.

Rising markups must be due to either price increases or marginal cost reductions. We observe that real prices increase during the early years of the sample period and then fall during the later years. Specifically, from 2006 to 2012, average real prices increase by seven percent. After 2012, average real prices decline and, by 2019, are only two percent higher than in 2006. Although price increases partially account for rising markups initially, by the latter years of the sample, cost reductions account for most of the aggregate markup trend.

In many models with imperfect competition, including the one that we estimate, cost changes are not completely passed through to prices. In such settings, falling marginal costs would typically lead lower prices but higher markups. However, incomplete pass-through cannot, on its own, explain the combination of lower marginal costs and slightly higher prices that emerges from the data and our estimates. Our estimates indicate that demand-side changes help to account for these trends. We find that demand for consumer products has become less elastic over time. In particular, consumer price sensitivity declines by about 30 percent from 2006 to 2019. 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.

We exploit the unique panel structure of our data to explore factors that predict markup trends. In regressions with product and time fixed effects, we find that products with larger increases in markups tend to have greater reductions in both marginal cost and price sensitivity. Indeed, these two factors explain a substantial majority of the differential trends in product-level markups. Changes in consumer demographics and market concentration also are correlated with markups but have much less explanatory power. We then use counterfactual simulations to examine how equilibrium markups would have evolved in response to our estimated changes in price sensitivity and marginal costs if demographics, product assortments, product ownership, and other demand parameters were constant over time. The results confirm that these two factors can account for almost all of the time-series variation in markups.

In many markets, including the consumer products markets we examine, one might expect costs to decline over time as firms improve their production and distribution technologies. Thus, perhaps more surprising is the decline in consumers' price sensitivity. To explore potential mechanisms, we analyze whether changes in price sensitivity are associated with changing retail patterns, such as the growth of online retail and warehouse clubs, or firm-level investments in R&D or marketing. However, we find that these factors account for only a small fraction of the differential category-level trends in price sensitivity. This suggests that lower price sensitivity might instead arise from exogenous shifts in consumer behavior, such as increase in opportunity cost of time. Consistent with this hypothesis, we find that the use of coupons, which involve some small efforts by consumers, has been falling in the U.S. in aggregate since the early 1990s. Over our sample period, total coupons redeemed and coupon redemption rates have fallen by 50 percent and 30 percent, respectively. In addition, according to time use data, time spent shopping on consumer products fell by approximately 20 percent during our sample period.

In our final analyses, we explore consumer surplus and welfare. Our findings indicate that consumer surplus per capita has increased during our sample period despite rising markups. We attribute this to changing preferences, particularly lower price sensitivity. The changes in consumer surplus vary across the income distribution. While consumers with incomes above the median had substantial gains in surplus during the second half of our sample period, the lowest income quartile experienced substantial losses in some time periods and had approximately the same level of consumer surplus at the end of our sample period as they had in 2006.

Changes in markups have been costly for consumers despite the increase in consumer surplus. In a counterfactual simulation, we find that consumer surplus would have been 14 percent higher in 2019 if markups were scaled down to 2006 levels. Furthermore, under the counterfactual of marginal cost pricing, consumer surplus in 2019 increases by 50 percent and total welfare increases by 9 percent. Taken together, these analyses suggest an important impact market of power on resource allocation, aggregate welfare, and the distribution of income—subjects of longstanding interest (e.g., Harberger, 1954).

Our analysis uses detailed product-level sales from the Kilts Nielsen Retail Scanner Data, which consists of a large sample of retail stores. The sales data primarily come from mass merchandisers, grocery stores, and drug stores. Out of a wider set of broad-basket retailers (i.e., also including warehouse clubs and dollar stores), consumer spending on these three retail channels comprised 83 percent of revenues in 2007 and 82 percent in 2019. Thus, our focal channels represent a substantial share of spending on consumer products throughout our sample period. We complement the sales data with the Kilts Nielsen Consumer Panel Data, which contain household-level purchases and demographic information. These data allow us to control for potential selection across retail channels by consumers with different demographics, as well as allowing for differences in product preferences across households.

A significant contribution of this paper is the application of flexible demand models across

categories and over time. We employ the random coefficients logit demand model of Berry et al. (1995) and allow consumer preferences to vary with observable and unobservable demographic characteristics. Typical empirical applications of this model return one set of preference parameters. By contrast, we apply the model across 133 categories, and, critically for our analysis of changing preferences, separately in each of year of our sample. Thus, our approach yields a panel of preference parameters from 1,862 estimated models.

Though we primarily focus on aggregate trends across a broad set of product categories, our empirical approach yields estimates that are consistent with more narrow studies that focus on individual product markets. These comparisons prove useful for assessing the potential simplifications of our model and our identification strategy for the price parameter, which makes use of covariance restrictions (MacKay and Miller, 2022). For 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 estimate the commodity price is roughly half of total marginal costs. For ready-to-eat cereals, we estimate costs and margins in line with those of Backus et al. (2021), who employ additional product characteristics and use an instrumental variables strategy. More broadly, for categories that we can find random coefficients logit estimates, we find that our model yields similar elasticities/markups.

Our research contributes to a growing empirical literature on the evolution of markups. Our finding of increasing markups across a number of categories is broadly consistent with De Loecker et al. (2020); given our distinct modeling approach, we are able to provide insights into specific supply and demand mechanisms. A number of studies recover markups from estimates of demand elasticities, as we do, focusing on specific industries over time. Ganapati (2021b) finds that the markups of wholesalers increased over 1992-2012 due to greater scale economies and the expansion of distribution networks, and with consumers benefiting from lower prices and access to higher quality goods. Grieco et al. (2022) 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. (2022) show that technology adoption in the cement industry over 1974-2016 reduced marginal costs, but that markups increased only modestly because cost reductions were mostly passed through to consumers. Consistent with our results, these studies highlight the role of technological change as a determinant of long run economic outcomes.³

Two other articles explore the relationship between changing consumer preferences and markups. Berry and Jia (2010) find that an increase in consumer price sensitivity helps explain a modest decline in the markups of airline carriers 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 reflects the strength of brand preferences, it may

³Also related is Peltzman (2020), which analyzes accounting data on manufacturing firms over 1982-2012 and finds support for rising markups and increasing total factor productivity.

increase in some sectors even as it decreases in others. Finally, Brand (2021) considers the hypothesis that increases in product variety lead to lower price sensitivity. He estimates demand in nine of the consumer product categories that we consider, both in 2006 and 2017, and finds less elastic demand and higher markups in the later year. Key distinguishing factors in our analysis include both the scope of our analysis—we consider a much broader set of product categories in every year—and our use of individual consumer data to link substitution patterns to variation in demographics in the cross section and over time. In addition, we deal with the issue of price endogeneity.

The paper proceeds as follows: In Section 2, we discuss our approach for recovering markups and specify the model of demand and supply. We discuss the data in Section 3. In Section 4, we describe the estimator and our identification strategy, and we validate the results of our empirical approach for selected industries. Section 5 describes the evolution of markups over time and discusses possible determinants of market power. In Section 6, we investigate the role of changes in price sensitivity and its determinants. In Section 7, we calculate consumer surplus and welfare over time for different scenarios. Section 8 concludes.

2 Methods

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:

$$\frac{P - C'}{P} = -\frac{1}{\varepsilon} \quad (1)$$

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 widely-used measure of markups (Lerner, 1934; Elzinga and Mills, 2011). Knowledge of the demand elasticity identifies the Lerner index. With data on price, one also can recover marginal cost, the additive markup (i.e., $P - C'$), and the price-over-cost markup (i.e., P/C').

The demand approach gained prominence in industrial organization after various methodological advances made it possible to estimate demand systems for markets that contain many differentiated products (e.g., Berry, 1994; Berry et al., 1995). 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 mecha-

nisms. However, in part due to the computation burden of demand estimation, most applications focus on a single industry or consumer product category. An advance of our paper is that it employs a flexible demand model across many product categories simultaneously.

The main alternative is the so-called *production approach* that was pioneered in Hall (1988) 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 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 with accounting data on revenues and expenditures. As with many research designs, challenges arise in implementation. For example, Raval (2020) finds that using different variable inputs can yield different markups, and Bond et al. (2021) demonstrates that markups may not be identified if revenue is used as a proxy for output.⁴ Due to these and other concerns, some scholars have argued that the existing evidence of rising markups is rather suggestive than definitive (e.g., Basu, 2019; Berry et al., 2019; Syverson, 2019).

Importantly, the demand approach we pursue is distinguished from the production approach 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. Thus, we view large-scale evidence on the evolution of markups obtained with the demand approach as a useful complement to the evidence that has been obtained with the production approach (e.g., De Loecker et al., 2020; De Loecker and Eeckhout, 2021).⁵ Implementation of the demand approach comes with its own challenges. As suggested by equation (1), inferences about markups are inextricably linked to the demand elasticities, so an identification strategy is needed to obtain consistent estimates of the demand-side parameters in the presence of price endogeneity. Perhaps more fundamentally, the demand-side approach requires the researcher to specify the structure of the demand system and the nature of competition between firms.

We maintain the assumptions of differentiated-products Bertrand competition and random coefficients logit demand, which have been widely used in the literature to study consumer products. There may be some product categories for which our assumptions may be inappropriate. Our strategy to mitigate any such misspecification bias is to aggregate results across product categories. Implemented at scale, this allows us to explore how product-level markups have evolved, the reasons for any such changes, and the consequences for consumers and firms.

⁴See also Doraszelski and Jaumandreu (2019) and De Ridder et al. (2022).

⁵One 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, 2017).

2.2 Demand Model

For each product category and each year, we apply the random coefficients logit model of Berry et al. (1995). We work with scanner 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 $j = 0, \dots, J_{crt}$ products available for purchase in chain c , region r , and quarter t , including an outside good ($j = 0$). Each affiliated consumer chooses among these products. The indirect utility that consumer i receives from a purchase of product $j > 0$ is

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

where p_{jcrt} is the retail price, the terms $(\xi_{jr}, \xi_{cr}, \xi_t)$ are product \times region, chain \times 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

$$\alpha_i^* = \alpha + \Pi_1 D_i \quad (3)$$

$$\beta_i^* = \beta + \Pi_2 D_i + \sigma v_i \quad (4)$$

where D_i contains the observed demographics and $v_i \sim N(0, 1)$ contains an unobserved consumer demographic. 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 $\theta = (\alpha, \Pi_1, \Pi_2, \sigma)$.

Note that we have omitted subscripts for year and product category. However, as we estimate demand separately for each category-year, all structural parameters and fixed effects are allowed to vary freely by product category and year.

Quantity demanded is given by $q_{jcrt}(p_{crt}; \theta) = s_{jcrt}(p_{crt}; \theta) M_{crt}$, where $s(\cdot)$ is the market share, p_{crt} is a vector of prices, and M_{crt} is the “market size” of the chain-region-period, a measure of potential demand. We refer readers to Nevo (2000b) for equations that characterize market shares and the demand elasticities. We use a market size that is proportional to the population and the number of retail stores operated by the chain within each region. We provide details on the calculation in Appendix B, and we show that our main trends are robust to alternative measures in Appendix E.4.

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. Our specification also incorporates heterogeneity in the utility that consumers receive from the inside goods, which allows the data to determine the extent of substitution between the inside and outside goods.⁶ In principle, product characteristics other than price could be incorporated into the demand model. We do not pursue this because it would require matching to auxiliary datasets on characteristics, which would be difficult to implement at scale.⁷

Data on non-price characteristics would allow for a more flexible treatment of horizontal differentiation in the model. It is generally recognized in industrial organization that this can have benefits for counterfactuals involving specific cross-product substitution patterns, such as merger simulation (e.g., Nevo, 2000a) or studies of entry and exit (e.g., Ciliberto et al., 2021). Whether it has first-order implications for markup trends depends on the prevalence of changes in product ownership, such as those that would be introduced by mergers, entry, or exit. Averaging over the product categories, we do not observe meaningful changes in concentration over the sample period (Figure F.3 in the Appendix).⁸ Furthermore, our analysis includes a screen for within-category product differentiation to account for potentially substantive unobserved product characteristics. We obtain similar results with and without this screen (Figure E.1 in the Appendix). Therefore, we conclude that our treatment of non-price characteristics is unlikely to drive our results.

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).⁹ The first order conditions for profit maximization can be expressed in terms of

⁶An alternative approach that allows data to influence substitution between the inside and outside goods involves specifying a random coefficients nested logit (RCNL) model with the outside good in its own nest (e.g., Grigolon and Verboven, 2014). With the RCNL model, the speed of estimation slows dramatically for higher values of the nesting parameter, making the model inappropriate for our application.

⁷Consider the approach that Backus et al. (2021) take to estimate demand for ready-to-eat cereals. They obtain auxiliary data from Nutritionix about the nutritional content of the products, such as the grain (e.g., wheat or corn) and the sugar content. These data then are consolidated into a handful of principal components that serve as product characteristics in the demand model. For many of the product categories we consider, and all of the non-food categories, nutritional content is unavailable or unlikely to drive consumer substitution.

⁸Bhattacharya et al. (2022) provide a detailed examination of the mergers in the same retail scanner data.

⁹This assumption is maintained elsewhere in the literature (e.g., Miller and Weinberg, 2017; Backus et al., 2021), and is consistent with empirical findings that indicate that retail prices respond to cost shocks (Butters et al., 2022) but not shocks to retailer demand (Arcidiacono et al., 2020). Further, evidence also suggests that retail markups have been relatively stable over the period 1980-2014, despite large changes in demand (Anderson et al., 2018). Finally, we note that the model accommodates nonlinear contracts that specify slotting fees or other fixed transfers.

the additive markup:

$$p_{crt} - c_{crt} = - \left(\Omega_{crt} \circ \left[\frac{\partial s_{crt}(p_{crt})}{\partial p_{crt}} \right]' \right)^{-1} s_{crt}(p_{crt}) \quad (5)$$

where the vectors p_{crt} , s_{crt} , and c_{crt} collect the prices, market shares, and marginal costs of products $j = 1, \dots, J_{crt}$, 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.

An implication of optimal price-setting behavior is that firms find it profitable to adjust their markups with demand conditions, which enter equation (5) through market shares and demand derivatives. Therefore, our model explicitly allows for price endogeneity, which we address in estimation. We decompose marginal cost according to:

$$c_{jcrt} = \eta_{jr} + \eta_{cr} + \eta_t + \Delta\eta_{jcrt} \quad (6)$$

where $(\eta_{jr}, \eta_{cr}, \eta_t)$ are product \times region, chain \times 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 incorporates “cost shifters” that have been used in the literature to estimate demand, including changes in materials costs and distribution costs that affect products and chains differentially.

3 Data

3.1 Data Sources and Estimation Samples

Our primary sources of data are the Retail Scanner Data and Consumer Panel Data of Kilts Nielsen, 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. 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 employ aggregation and a number of screens to construct samples that are suitable for the model laid out in the previous section.

We take as given the consumer product categories (“modules”) that are specified by Nielsen. Within each category are UPCs that consumers are likely to view as 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, and then applying a screen based on observed price dispersion to avoid categories with highly dissimilar products. We discuss our category selection procedure in more detail in Section 3.2.

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
<hr/>							
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
<hr/>							
	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 nominal US \$ 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).

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 Nielsen and are fairly narrow. For example, in ready-to-eat cereals, “Cheerios,” “Honey Nut Cheerios,” and “Multi-grain Cheerios” are three distinct brands.¹⁰ 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.¹¹ 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 85 percent of category revenues and typically include a private label product.¹²

We focus our analysis on the stores that Nielsen classifies as mass merchandisers, grocery

¹⁰Other examples include “Oreo,” “Oreo Double Stuf”, and “Mini Oreo” (cookies) and “Yoplait,” “Yoplait Go-gurt,” “Yoplait Whips!,” “Yoplait Thick & Creamy,” and “Yoplait Light Thick & Creamy” (yogurt).

¹¹In a handful of categories, UPC codes differ in terms of whether units are reported in terms of volume, mass, or count. For those categories, we use only those UPC codes associated with the highest-revenue metric.

¹²To explore the sensitivity of the analysis to the cap of 20 branded products per category, we perform robustness checks with a sample that includes only 15 branded products per category. We obtain very similar results. More brands could be added to the model with additional effort to connect brands to their owner, following the same process that we use for the brands currently in the sample (as discussed later in this section).

stores, or drug stores. These retail channels comprise a substantial part of overall consumer spending on consumer products. Based on auxiliary data on the revenues of large U.S. retailers, we estimate that, in 2019, they accounted for 82 percent of revenues among broad-basket retailers (i.e., mass merchandisers, grocery, drug stores, dollar stores, and warehouse clubs). This share of revenue appears to be stable in our sample period, as the estimated share in 2007 is 83 percent. Among all channels, we estimate that mass merchandisers, grocery stores, and drug stores account for over 50 percent of consumer product spending, where the broader sample includes specialty retailers (e.g., electronics, beauty, apparel). Appendix B.3 provides these summary statistics and describes the auxiliary data.¹³

We use the designated market areas (DMAs) in the Nielsen 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. Table 1 provides summary statistics for a selection of product categories in the estimation sample sorted by number of observations.

We employ household demographic data to account for differences in the composition of consumers across markets and changes within markets over time. Specifically, we generate consumer-specific demographic draws by sampling 2,000 consumers from the Consumer Panel Data for each region and year.¹⁴ We sample with replacement and using the projection weights provided by Nielsen. Among the available demographics, we select two that we expect should 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, and 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.

In estimation, we match the empirical purchasing patterns of households across different de-

¹³The largest broad-basket channel that we omit is warehouse club, which accounts for 9.0 percent of consumer product spending in 2007 and 9.4 percent in 2019. We observe that the revenue share of dollar stores nearly doubles between 2007 and 2019, consistent with the trend documented in Caoui et al. (2022). Nonetheless, dollar stores account for only 1.5 percent of consumer product spending in 2007 and 2.6 percent in 2019. The share of revenues accounted for by retailers that we do not identify as broad-basket declines slightly over time. This reflects a growth of online retailers that is offset by relative declines in other store formats (e.g., department stores, apparel).

¹⁴By sampling at the region-year level, we implicitly assume that the consumers of retail chains within the same region have the same demographics. We take this approach to because we view the consumer panel data as too sparse to reliably sample at the level of a retail chain, region, and year. For a study of consumer demographics and prices as they vary spatially across a city, see Eizenberg et al. (2021).

mographic types, which allows us to control for heterogeneous preferences and for selection by households into different retailers. Specifically, we use the data to construct “micro-moments” that are the average values of observed demographics for consumers that purchase each product in a given region and year, again using the projection weights. Our model 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 these values, we use purchasing data only at a subset of retailers to match the distribution of retailers that appear in the scanner data (e.g., mass merchandisers, grocery stores, and drug stores). Since our sample of households is not restricted in this way, the model provides some adjustment for selection of consumers into the retailers we observe.

We account for multi-product ownership using auxiliary data, as ownership information is not provided in the Nielsen 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. 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. This information is available for about half of the observations in our sample because Compustat covers publicly traded firms.

We deflate prices and incomes using the Consumer Price Index such that they are in real dollars as of the first quarter of 2010.¹⁵

3.2 Selection of Product Categories

Some challenges arise in recovering markups over time using the estimation samples described above. In treating the Nielsen categories as well-defined product markets, we create the potential for model misspecification, due to at least two (related) reasons. The first is that products in different categories might be substitutes. For instance, one might suspect some amount of consumer substitution between products in 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. However, looking holistically across the Nielsen categories, we are skeptical that cross-category substitution is meaningful for most

¹⁵We deflate using the Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average. The inflation data are monthly and seasonally adjusted.

products. Thus, for our research question, it seems more appropriate to use the Nielsen categories rather than making ad hoc adjustments, and that is the approach we take.

The second reason for concern about Nielsen product categories—which we view as more important for our application—is that some categories include products that might be very weak substitutes (or possibly not substitutes at all). The “Batteries” category, for example, has some products that are probably close substitutes, such as various brands of AAA batteries, along with other products that are functionally quite different, such as D batteries. We use a relatively tractable specification of the random coefficients logit model in order to scale estimation across categories, and do not consider the model to be sufficiently flexible to handle such rich patterns of product differentiation. This can be problematic if the same demand parameters—and especially the price parameter—are inappropriate for different classes of products within the same category.

To address this potential concern, we use the within-category distribution of prices as a proxy for within-category product heterogeneity, and remove categories in which the 99th percentile of prices is greater than five times the median price. This screen leaves 133 of the top 200 product categories (by revenue) in the baseline sample. The top 200 categories account for 74 percent of revenues in the Retail Scanner Data; the 133 categories in the baseline sample account for 55 percent. Although our screen for within-category heterogeneity focuses attention on categories for which the model is a likely to be a better fit, it does not drive results; we obtain similar markup trends with screens that are more or less strict. In Appendix E.1, we report the product-level markup trends using all 200 of the product categories.

Also worthy of discussion are the compositional changes that occur in the Nielsen data as retail stores enter and exit the sample. Such churn appears to be inconsequential over 2006–2017, but significant changes do occur over 2018–2019. Because we estimate independent models separately in each year, compositional changes do not affect the trends we observe from 2006–2017. We control for some aspects of compositional changes in 2018–2019 by including (yearly) chain \times region fixed effects in the demand and marginal cost equations and allowing market sizes to scale separately for each retail chain. Moreover, we show in Appendix E.3 that we find nearly identical trends with a balanced panel that includes only brand \times chain \times region combinations that occur in every year of our sample.

4 Empirical Strategy

4.1 Estimation and Identification

We estimate the equilibrium model developed in Section 2 using the generalized method of moments (GMM). We estimate separate models for each category and year, allowing the parameters for estimation, $\theta = (\alpha, \Pi_1, \Pi_2, \sigma)$, to vary arbitrarily across models. The GMM estimator

for θ is:

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

where W is a weighting matrix, $g^{MM}(\theta)$ collects a set of micro-moments that summarizes how well the model matches the correlations between demographics and product purchases that we observe in the Nielsen Panelist dataset, and $g^{CR}(\theta)$ implements 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 . This reflects that micro-moments identify θ_2 but not α (Berry and Haile, 2020), and that the covariance restriction exactly identifies α conditional on θ_2 (MacKay and Miller, 2022). In Appendix A, we explain why this segmentation has computation advantages in our setting and provide additional details on the estimation procedure.

For 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

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

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).¹⁶ In our baseline specification, we use two observed demographic variables and at most 21 products, so there can be up to 42 micro-moments.

In the second step, we identify the price parameter under the assumption that the demand-side and supply-side structural error terms are uncorrelated in expectation: $\mathbb{E}[\Delta \xi_{jcrt} \Delta \eta_{jcrt}] = 0$. We construct the empirical analog of the moment condition:

$$g^{CR}(\theta) = \frac{1}{T} \sum_{c,r,t} \Delta \xi_{crt}(\theta)' \Delta \eta_{crt}(\theta) \quad (9)$$

where the $\Delta \xi_{crt}(\theta)$ and $\Delta \eta_{crt}(\theta)$ terms are recovered for each candidate θ using standard techniques, and T is the number of chain-region-quarter-product combinations for a given year.

¹⁶We allow the average observed demographics to vary by year and category. An alternative approach to the micro-moments would match the implied chain-region demographics to chain-region data, rather than to region-level data. The tradeoff is between the measurement error in the observed component versus the specificity of the moments. However, parameters that fit one set of moments well should also fit the other well.

Estimation of θ_2 is standard and identification strategies for these parameters are reasonably well understood (Berry and Haile, 2020). However, micro-moments that can be used to pin down heterogeneity in preferences cannot recover the mean price parameter and resolve price endogeneity. We address this with covariance restrictions, which are appealing in our setting because they can be implemented at scale for different product categories and years. Alternative approaches typically rely on auxiliary data on cost-shifters or product characteristics, which can be difficult and costly to obtain. An additional benefit of the covariance restrictions approach is that—in contrast to an instruments-based approach—there is no “first-stage” relevance condition that must be satisfied (MacKay and Miller, 2022). Even if product characteristics were available for every category and year, there is no guarantee that, for example, markup-shifter instruments that rely on such characteristics (e.g., Berry et al., 1995; Gandhi and Houde, 2020) would meet the relevance condition for every category of interest.

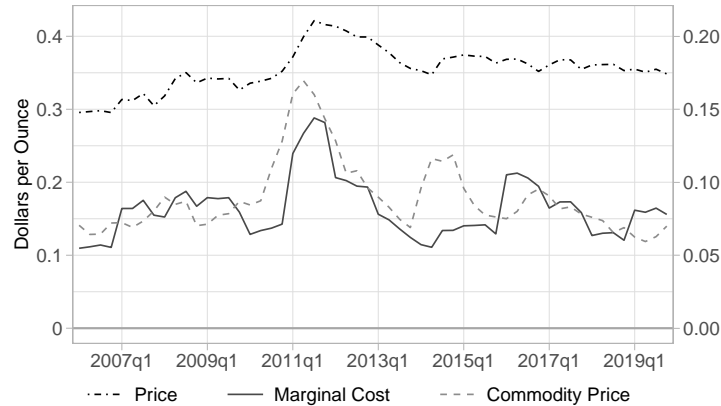
As we have specified the model, the supply-side structural error term ($\Delta\eta_{jert}$) incorporates the variation of some of the cost-shifter instruments that have been used to estimate demand in the recent literature, including product-specific shipping costs (Miller and Weinberg, 2017) and the prices of product-specific ingredients (Backus et al., 2021). These and other plausibly exogenous cost-shifters may be highly correlated with the variation that we exploit in estimation. Because the marginal cost function includes fixed effects at the product \times region, chain \times region, and quarter levels, some potentially confounding sources of variation are absorbed, including the possibility that higher quality products are more expensive to produce.¹⁷ Further, these fixed effects shift arbitrary year-over-year, allowing for longer-run changes in production technology that are correlated with demand.

The covariance restrictions approach to estimation differs in some ways from an instrument-based approach (MacKay and Miller, 2022). In particular, the covariance restrictions approach uses all of the endogenous price and quantity data in estimation, rather than only the portions that are attributable to excluded instruments. Although this eliminates the first-stage relevance requirement, it does require the joint estimation of parametric models of supply and demand. Thus, a misspecification of the marginal cost function could affect demand estimates. However, because a fully specified supply-side model is required to recover markups, we view it as sensible to also employ the supply model to estimate structural parameters.

As shown by MacKay and Miller (2022), reduced price sensitivity would suggest that the ratio of the variation in quantities to the variation in prices is falling over time. Intuitively, this reflects demand that is less sensitive to price variation. A change in the price coefficient corresponds to a rotation of the inverse demand curve; more inelastic demand will result in

¹⁷Our supply-side assumption imposes that marginal costs are constant in output. For consumer products, we view this as a reasonable approximation and the assumption is often maintained in the literature (Villas-Boas, 2007; Chevalier et al., 2003; Hendel and Nevo, 2013; Miller and Weinberg, 2017; Backus et al., 2021). The literature suggests a positive correlation between demand-side and supply-side structural error terms may exist in other settings, including automobile manufacturing (Berry et al., 1995) and airlines (Ciliberto et al., 2021).

Figure 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.

a more “vertical” inverse demand and inverse supply curves on a price-quantity plot and a lower relative variance. Indeed, in our data, within-market price dispersion is increasing while within-market share dispersion is falling, and the changes to the relative variance are highly correlated with the changes in the price sensitivity we estimate.¹⁸

4.2 Assessment

We conduct three 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. Third, we plot the distribution of elasticities that we obtain with our baseline estimates, and also compare this distribution to two alternative approaches that have been used in the literature.

Marginal Cost Estimates Figure 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

¹⁸Simulations in MacKay and Miller (2022) suggest that the covariance restriction approach can be robust to modest supply-side misspecification. As an empirical robustness check, we explore an alternative approach that does not require our supply-side model in estimation. In Appendix E.5, we calculate trends in demand (elasticity and price sensitivity) under the assumption that prices are exogenous. Though this often provides biased elasticities (see Section 4.2), we interpret the robustness check as being consistent with rotations in the demand curve, i.e., with demand becoming less elastic.

Table 2: Average Product-Level Own-Price Elasticities of Demand

Category	Our Estimate	Literature Estimate	Citation
Beer	-4.06	-4.74	Miller and Weinberg (2017)
Ready-to-Eat Cereal	-2.29	-2.42	Backus et al. (2021)
Yogurt	-3.12	-4.05	Hristakeva (2020)

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 (2020) 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.

index for coffee, which is scaled separately on the right axis.¹⁹ Overall, our recovered estimates of marginal costs are strongly correlated with the commodity price index. A regression of average marginal costs on the commodity price yields a coefficient of 0.990 ($p < 0.001$), and the correlation between the two time series is 0.62. Our method is able to capture the large spike in commodity prices in 2011, which is reflected in the spike in marginal costs. We find that, on average, the commodity price is equal to 56 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.

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 2, 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 -4.06, -2.29, and -3.12, 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.²⁰ 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 (2020) 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.

¹⁹Data on coffee commodity prices were obtained from Macrotrends.net. Available here: <https://www.macrotrends.net/charts/commodities>, last accessed March 1, 2022

²⁰Every paper differs in the exact data sample used. For example, Hristakeva (2020) 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.

(2021), which was developed concurrently. In their analysis of ready-to-eat cereals, Backus et al. (2021) use the Kilts Nielsen 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 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 very similar elasticities and margins.²¹

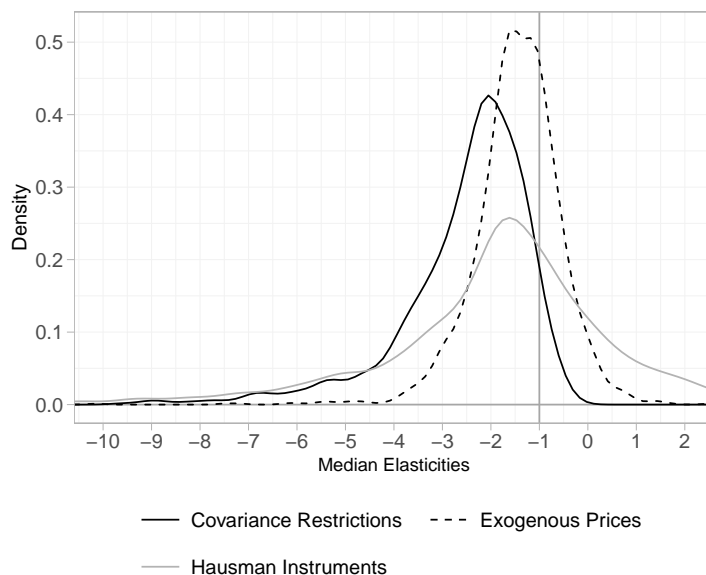
Alternative Identification Strategies For the third validation check, we examine the distribution of median own-price elasticities across all of the 1,862 category-year combinations in our baseline sample. We compare the results to those obtained under two alternative assumptions that can identify the price parameter and be applied at scale. The first alternative assumption is that prices are exogenous. For a given model of supply and demand, price exogeneity holds if both (a) firms do not adjust markups in response to demand shocks and (b) demand shocks are uncorrelated with marginal cost shocks. If the latter condition fails, then prices are endogenous (i.e., correlated with demand shocks) even if firms do not respond directly to demand. Thus, a covariance restriction is necessary for consistent estimation under an assumption of exogenous prices. However, profit maximization generally implies that prices are endogenous, and our covariance restrictions approach to estimation corrects for price endogeneity.

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 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, consistent with price endogeneity arising from firms adjusting prices in response to demand shocks. Though covariance restrictions systematically correct for price endogeneity, Hausman instruments do not, and instead yield more elastic demand than exogenous prices in some cases and more inelastic demand in others.

²¹For cereals, our average unit price is 0.20 and our average estimated marginal cost is 0.10.

Figure 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 .

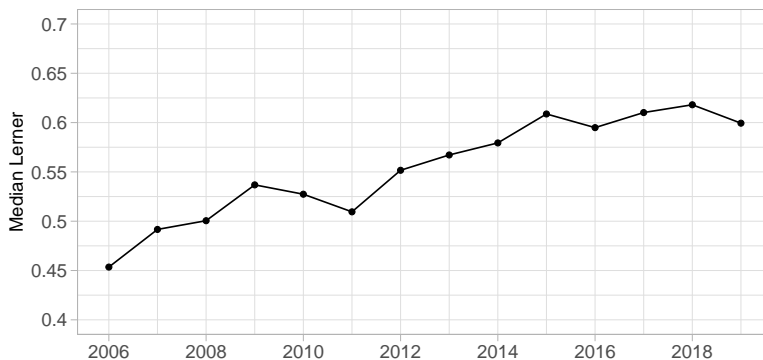
Using covariance restrictions, demand is never upward-sloping, and only 5 percent of the category-year combinations have inelastic demand (i.e., a median elasticity greater than -1). By contrast, 29 percent of the category-year estimates exhibit inelastic demand with exogenous prices; with Hausman instruments, it is 34 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.

Of course, our ultimate interest is in the evolution of markups across the many different categories in our estimation sample, and we turn to that exercise next.

5 The Evolution of Markups in Consumer Products

In this section, we document the evolution of markups across consumer products over time. We start by reporting median markups at the product category level before we discuss how the distribution of markups has shifted. We then move the analysis to the product level which allows us to distinguish within-product variation from variation across products and to decompose the evolution of markups into changes in prices and marginal costs.

Figure 3: 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.

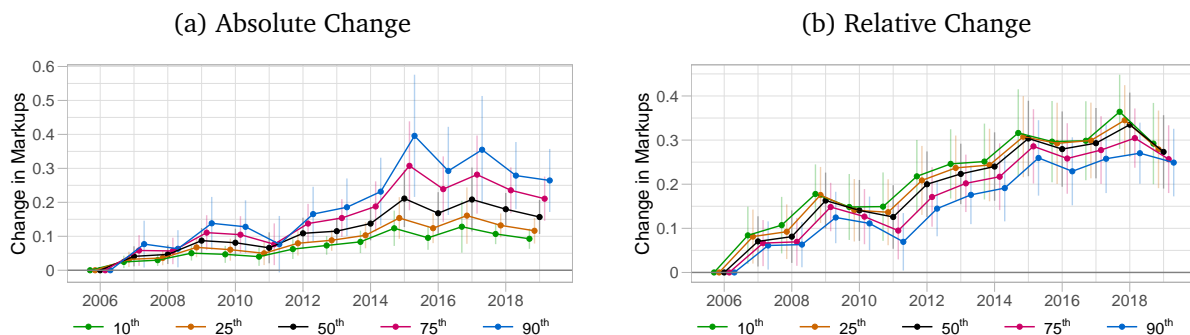
5.1 Aggregate Markup Trends

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 3 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.

Next, we analyze how the distribution of markups within product categories has shifted over time. For this purpose, we regress different percentiles of the markup distribution on year dummies and document the coefficients and confidence intervals in panel (a) of Figure 4. We use the year 2006, the first year of our estimation sample, as the base category. Hence, the estimated coefficients can be interpreted as the change in markups in each year relative to 2006. The results indicate that, while all quartiles of the distribution have increased over time, the upper part of the markup distribution has changed by a higher amount, especially during the second half of our sample period. In panel (b), we repeat the exercise by using the log of the Lerner index, $\ln(\frac{p-c}{p})$. The results show that the *relative* increase in markups is in fact quite similar across the distribution and even slightly more pronounced for lower quartiles. Overall, our estimates indicate that the full distribution of markups is shifting upward over time.²²

²²We 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 4: Changes in the Distribution of Markups



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.

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

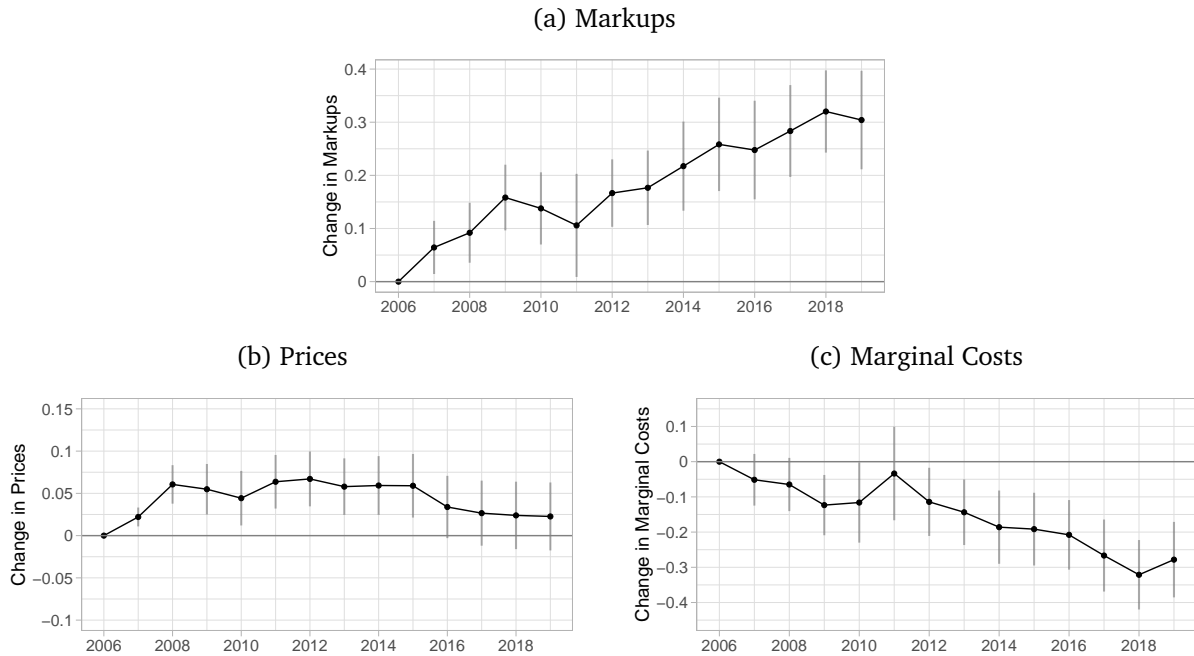
The aggregate trends in markups that we document could be explained by firms charging higher markups on existing products or by market entry and exit of brands with different levels of markup. Further, to the extent that within-product changes explain rising markups, this could be due to higher prices, reductions in marginal costs, or both. To evaluate these mechanisms, we analyze the change in markups, prices, and marginal costs at the product level, where our unit of observation is a unique product-chain-region-quarter-year combination.

For markups, we regress the log of the Lerner index on quarter, year, and product-chain-region fixed effects, using revenues as weights.²³ The results of this regression are documented in panel (a) of Figure 5. The figure displays point estimates and 95 percent confidence intervals for the year fixed effects. The estimates indicate an increase in product-level markups of about 30 percent between 2006 and 2019. The estimated annual growth rate in product-level markups is 2.2 percent per year. With the inclusion of product-chain-DMA fixed effects, the nonparametric time trend only captures variation within products. Thus, the estimated change over time is not affected by entry and exit or a reallocation of market shares across products. This indicates that the aggregate markup trends are mainly driven by changes within products over time. We find similar results if we instead use price-over-cost (p/c) markups, as studied by De Loecker et al. (2020). See Appendix E.2.

Table F.1 in the Appendix provides the full regression results that corresponds to panel (a) of Figure 5, alongside 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 in-

²³We 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.

Figure 5: Product-Level Changes in Markups, Prices, and 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.

creases in average 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 aggregate (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.²⁴

Using our detailed data on prices and our demand estimates, we can decompose the increase in markups into changes in prices and marginal costs (equation (5)). For this purpose, 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 5. Panel (b) shows that real prices increase at the beginning of our sample period, but decline in later years. The average real price for products in our sample increases by 7 percent over 2006 to 2012, but real prices are only 2 percent higher in 2019 than in 2006. Panel (c) of the figure reports the yearly coefficients for log marginal costs. We estimate that marginal costs decline by 2.1 percent per year on average. In 2017–2019, marginal costs are roughly 25 log points lower than in 2006.²⁵

²⁴Table F.2 in the Appendix shows results using unweighted regressions. The results are similar. As Table F.3 shows, we also obtain similar results if we focus on a balanced panel of products, indicating that the overall trends are not primarily driven by the entry and exit of products with different markup growth rates.

²⁵Figure F.1 in the Appendix uses nominal (i.e., non-deflated) prices and marginal costs, and shows that nominal

Thus, although higher real prices account for part of the increase in markups during the first half of our sample, the higher markups we observe at the end of our sample arise from lower real marginal costs, not higher real prices. Overall, our estimates suggest that declines in real marginal costs have not been passed on to consumers.

5.3 Changes in Demand

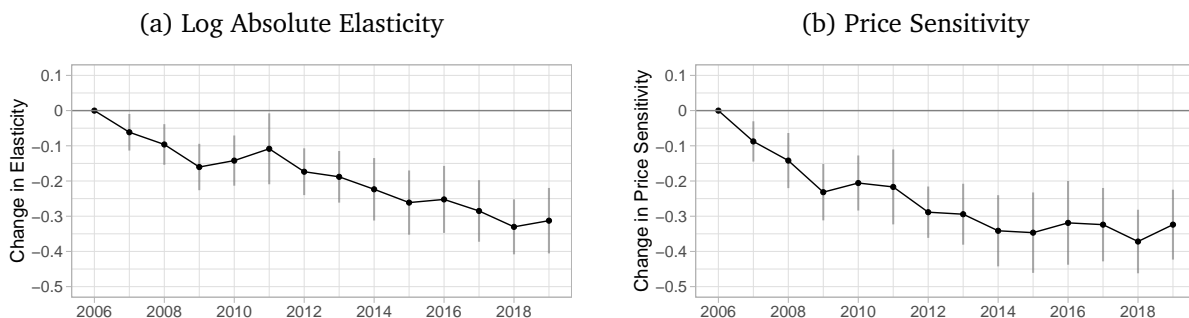
Why might lower marginal costs not lead to lower prices? Incomplete pass-through arises in many models of imperfect competition, including the one that we estimate. However, on its own, incomplete pass-through cannot explain the combination of lower marginal costs with slightly higher prices. Economic theory suggests other possibilities that could contribute to this phenomenon, including changes in demand. More inelastic demand would put upward pressure on markups, as evident in the first-order conditions in equations (5) and (1). Another possible explanation for increasing markups is the consolidation of brand ownership, which might occur due to mergers and acquisitions. We do not observe meaningful increases in concentration in our data (see Figure F.3 in the Appendix).

To investigate the possibility of demand-side changes, we first regress the logarithm of the absolute value of own-price elasticities at the product level on the same set of fixed effects used above. We present the results in panel (a) of Figure 6. The displayed coefficients show that price elasticities have declined in magnitude, indicating that demand indeed becomes less responsive to prices over time. Price elasticities capture several underlying aspects of consumer preferences and may also reflect supply-side factors such as quality and competition. However, in our sample the main driver appears to be changes in the mean price coefficients that we estimate for each category and year in the data. These parameters implicitly adjust for changing consumer demographics and selection by consumers into retailers and products. 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 that the declines in price sensitivity were large through 2012, corresponding with the increase in real prices we observe over the same period. In econometric and simulation-based exercises that we present shortly, this reduced price sensitivity emerges as an important determinant of rising markups.

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 F.2 in the Appendix shows that perceived product quality declines over the sample period. Improvements in the outside good—which includes shopping through online retailers

marginal costs are relatively constant over time.

Figure 6: Changes in Demand



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.

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). As we discuss below, these changes have relatively little impact on markup trends.

To summarize, our decomposition of effects indicates that the increase in markups was driven by lower real marginal costs, without commensurate reductions in real prices. Firms were able to charge higher markups because consumers became less price sensitive over time.

5.4 Panel Data Analysis

To evaluate the relative importance of demand and supply channels in driving changes in product-specific markups, we use a regression analysis that exploits the unique panel structure of our estimates across products and over time. Specifically, we regress product-level log markups on consumer preference parameters, marginal costs, consumer demographics, and market concentration. We use category and year fixed effects, such that the regression coefficients capture deviations from aggregate trends. We focus on the ability of the regressors to explain changes in product-level 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 by for each category, so that they have a variance of one.²⁶ 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

²⁶Standardization improves comparability across categories and also eases interpretation of the coefficients. We choose this approach to standardization, rather than logs, so as to include observations with negative values.

Table 3: Factors Predicting Cross-Category Variation in Markup Trends

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Marginal Cost (Standardized)	-0.564*** (0.024)					-0.450*** (0.023)	-0.449*** (0.023)
Price Sensitivity		-0.721*** (0.030)				-0.392*** (0.022)	-0.393*** (0.022)
Quality (Standardized)			-0.142*** (0.022)			0.006 (0.006)	0.007 (0.006)
Income (Log)				0.052** (0.025)		0.059*** (0.013)	0.058*** (0.013)
Children at Home				-0.175*** (0.064)		-0.076*** (0.026)	-0.083*** (0.027)
Parent HHI					0.236 (0.186)		0.236*** (0.046)
Brand HHI					0.091 (0.178)		-0.097** (0.048)
Retailer HHI					0.203*** (0.077)		0.074*** (0.025)
Brand-Category-DMA-Retailer FEs	X	X	X	X	X	X	X
Time Period FEs	X	X	X	X	X	X	X
Observations	14,407,410	14,407,410	14,407,410	14,407,410	14,407,353	14,407,410	14,407,353
R^2 (Within)	0.719	0.468	0.047	0.000	0.003	0.826	0.827

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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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.²⁷

Table 3 summarizes the results. Each regression includes fixed effects for each product-market (i.e., brand \times category \times retailer \times region) and time period (year \times quarter). Thus, the coefficients reflect the correlations of within-product changes over time. Standard errors are clustered at the product category level. The R^2 (within) statistic shows how much of the residual variation in markups—i.e., the portion not absorbed by fixed effects—is accounted for by the explanatory variables.

The results indicate that changes in marginal costs and price sensitivity are highly correlated with rising markups, and can explain the bulk of the variation in within-product markup changes. Column (1) indicates that marginal cost reductions alone can explain 72 percent of the within-product variation in markups (within $R^2 = 0.719$). The coefficient implies that a one standard deviation reduction in marginal costs is associated with a 56 percent increase in markups. Similarly, column (2) indicates that declines in price sensitivity alone can explain 47 percent of the within-product variation in markups; the coefficient indicates that a 10 percent decrease in price sensitivity is associated with a 7.2 percent increase in markups.

²⁷We use the consumer panel data to construct HHI measures. Our results are qualitatively similar if we instead use the retail scanner data.

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 F.4 in the Appendix.

Columns (3), (4), and (5) examine perceived quality, consumer demographics, and concentration. Although some of the coefficients are statistically significant, each of these measures explains little of the variation in log markups, with within R^2 values less than 0.05.

In column (6), we combine price sensitivity, quality, and marginal costs with demographic characteristics. 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. Increases in income remain positively associated with greater markups. Changes in price sensitivity, marginal costs, and demographics explain most of the variation in markups over time. The within R^2 is 0.83.

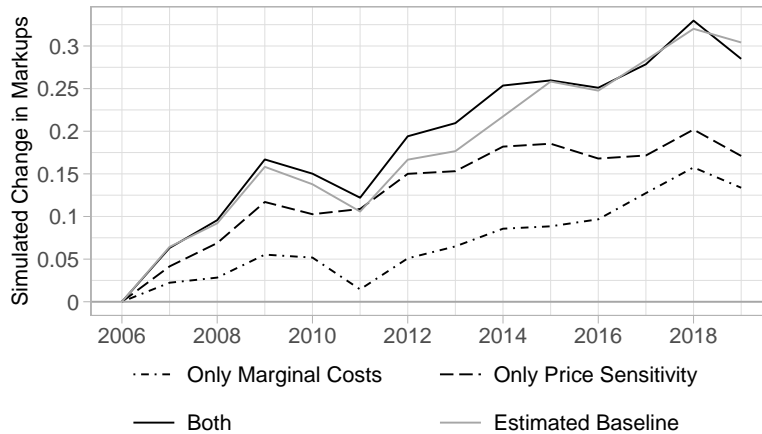
In column (7), we add our measures of concentration to the specification. We find that changes in retailer concentration remain positively correlated with changes in markups, and the coefficient for parent-manufacturer concentration increases and becomes statistically significant. Yet, these coefficients remain modest. The parent-retailer coefficient of 0.236 in column (7) indicates that a 0.02 change in parent company HHI—i.e., a 200-point change on a 0 to 10,000 scale—is associated with a 0.5 percent increase in markups. The relationship between markups and changes in concentration at the retailer level and brand level (which ignores multi-product ownership) is weaker. Overall, the inclusion of concentration measures does little to change the explanatory power of the regression, as the R^2 barely changes.

5.5 Impacts of Marginal Costs and Price Sensitivity on Markups

The previous subsection shows that reductions in marginal costs and price sensitivity are highly correlated with the variation in markup growth across products. Here, we use counterfactual simulations to show the hypothetical causal impact of these two factors on markup trends, holding everything else fixed. Specifically, taking 2006 data as a starting point, we change price sensitivity and/or marginal costs, holding fixed product assortments, consumer demographics, and demand parameters. Given these hypothetical changes, we use equation (5) to solve for equilibrium prices and compute markups. In each simulated year, we apply a uniform relative change to scale product-specific values by the estimated aggregate changes documented in Figures 5 and 6. Thus, we ask to what extent aggregate changes in marginal costs and price sensitivity can explain aggregate trends in markups.

The results of the counterfactual simulations are depicted in Figure 7. The dash-dotted line shows that, relative to 2006, estimated changes in marginal costs would have increased

Figure 7: 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 5 and 6. 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.

markups by about 13 percent in 2019 if preferences, demographics, product assortments, and ownership had not changed. Changes in price sensitivity (holding marginal costs and other factors fixed) would have increased markups by more than 15 percent towards the end of the sample period, as indicated by the dashed line. The solid black line shows that simulated markups increase by about 28 percent from 2006 to 2019 if we adjust both price sensitivity and marginal costs at the same time. The trajectory of simulated markup changes tracks overall markup trends, depicted by the gray line, very closely. Hence, changes in price sensitivity and marginal costs account for nearly all of the time-series variation in markups. Consistent with trends documented in Figures 5 and 6, changes in markups can be mainly attributed to changes in price sensitivity in the first half of our sample period, while marginal costs are the main driver of rising markups in the second half of our sample.

Economic theory provides a tight theoretical connection between changes in marginal costs and markups. In typical models of imperfect competition, a decline in marginal costs will not be fully passed on to consumers (i.e., cost pass-through is less than one). If costs fall faster than prices, then markups increase. Thus, the relationship that we find between markups and marginal costs is partly a result of imperfectly competitive product markets and declining costs. This logic applies to more general settings: in otherwise stable economic environments, declining costs will yield higher markups due to imperfect competition.

In many markets, we expect costs to decline over time due to innovations in production/distribution technology and operational efficiencies. Our empirical setting is no exception, as many manufacturers sought ways to reduce costs over this time period. For example, 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 2019.²⁸ Overall, our finding of modest declines in marginal costs is consistent with secular increases in productivity across the economy.

There is also a tight theoretical connection between price sensitivity and markups. All else equal, firms will charge higher prices to less price sensitive consumers. However, in contrast to our finding of declining marginal costs, it is perhaps more surprising that we find that consumer price sensitivity has fallen over time. In the following section, we examine the role of price sensitivity in more detail and discuss potential explanatory factors for the time trend.

6 Price Sensitivity and Markups

In this section, we explore the role that price sensitivity plays in explaining changing markups. First, we provide an econometric decomposition that isolates the price sensitivity parameter from observable features of the market that also determine markups in equilibrium. We use this decomposition to provide further evidence for the special role of consumer price sensitivity. We then explore potential mechanisms that could be driving changes in price sensitivity.

We apply an econometric decomposition developed to examine the role that the mean price parameter plays in our analysis. As shown by MacKay and Miller (2022), we can write the product-level additive markups as a function of the mean price parameter (α) and an inverse supply ($\lambda(\cdot)$) for a broad class of oligopoly models. In our model of random coefficients logit demand and Bertrand pricing, this takes the form:

$$p_{jcr,t} - c_{jcr,t} = -\frac{1}{\alpha} \lambda_{jcr,t}(s_{cr,t}, p_{cr,t}, \Gamma_{cr,t}; \Pi_1, \Pi_2, \sigma), \quad (10)$$

where $s_{cr,t}$ and $p_{cr,t}$ vectors of market shares and prices at the chain-region-quarter level, and $\Gamma_{cr,t}$ denotes the matrix of partial demand derivatives (with respect to prices). From an econometric standpoint, $\lambda_{jcr,t}(\cdot)$ is a function of market shares, prices, and consumer-specific choice probabilities; it does not depend on the mean price parameter. In Appendix C, we provide the specific functional form of $\lambda(\cdot)$.

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

$$\bar{L} = \frac{\bar{p} - \bar{c}}{\bar{p}} = -\frac{1}{\alpha} \frac{\bar{\lambda}}{\bar{p}}, \quad (11)$$

²⁸The Procter & Gamble Company 2019 Annual Report. Available here: <https://www.pg.com/annualreport2019/download/PG-2019-Annual-Report.pdf>

In logs, we obtain:

$$\ln \bar{L} = \underbrace{-\ln(-\alpha)}_{-1 \times \text{Price Sensitivity}} + \underbrace{\ln\left(\frac{\bar{\lambda}}{\bar{p}}\right)}_{\text{Structural Factors}}, \quad (12)$$

where we can decompose the (log) category markups into price sensitivity (i.e., $\ln(-\alpha)$) and a term that captures the net effect of other structural factors: the qualities and marginal costs of products, the ownership of products (i.e., market concentration), the parametric assumptions, and the nonlinear preference parameters. This term can be obtained from directly observable data on product ownership, market shares, prices, and consumer purchasing patterns such as the micro-moments that we use in the first stage of estimation.

The decomposition suggests a regression-based approach to explore the degree to which price sensitivity explains variation in markups across both product categories and over time. We start with cross-sectional regressions—separately for 2006, 2017, and 2019—in which the dependent variable is the category-level aggregate Lerner Index (in logs) and the independent variable is price sensitivity. 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.

Table 4 summarize the results. The R^2 in columns (1)-(3) indicates that 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. This suggests that other structural factors, such as product qualities and multi-product ownership, are relatively more important in explaining variation in markups across categories. Further, this highlights that our demand specification is sufficiently rich to attribute much of the variation in markups across categories to structural factors that are uncorrelated with consumer price sensitivity.²⁹ As our prior results indicate that decreasing price sensitivity is correlated with higher markups, one might suspect its explanatory power also to increase over time. Consistent with that, the R^2 in 2017 is higher than that of 2006; the lower R^2 in 2019 may be attributable to the compositional shift in the scanner data (which we control for in the analysis of the previous section.)

Column (4) summarizes the results of the panel regression. We find that changes in price sensitivity over time explain 58 percent of the variation in markups over time. Thus, to understand rising markups among the consumer products that we examine, it appears necessary to have an understanding of consumer price sensitivity and how it has changed over time. That is, 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

²⁹This 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 $\ln(\lambda_t/p_t) = 0$).

Table 4: Price Sensitivity and Markups Across Product Categories

	(1) 2006 Log \bar{L}	(2) 2017 Log \bar{L}	(3) 2019 Log \bar{L}	(4) Δ Log \bar{L}
Price Sensitivity	-0.134*** (0.027)	-0.200*** (0.029)	-0.090*** (0.029)	
Δ Price Sensitivity				-0.575*** (0.012)
Observations	133	133	133	1,729
R^2	0.162	0.268	0.070	0.571

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. 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. Column (4) captures the time series variation by estimating the model in first differences from 2007 through 2019. The regressions are motivated by the decomposition in equation (12). Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

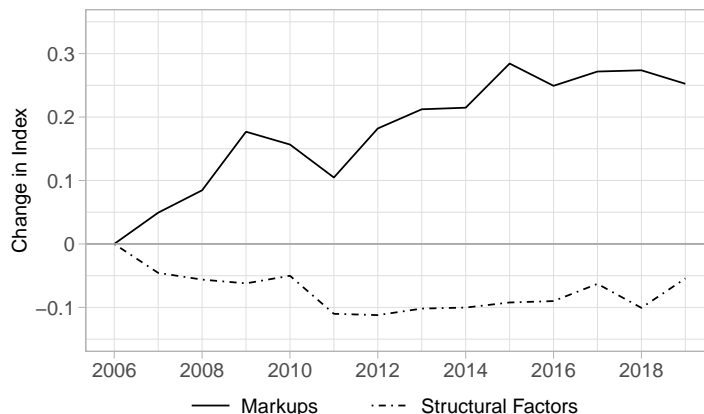
in price sensitivity. This points to a strength of our modeling approach: as we estimate demand separately for each category and each year, our estimates of price sensitivity can adjust flexibility over time with the shifts in the empirical variation in the data.

In fact, our analysis implies that a decline in consumer price sensitivity is necessary to generate higher markups in our sample. From 2006 to 2019, the average structural component of equation (12) decreased by 0.05. In other words, if price sensitivity had not changed over this period, then the observed changes in other structural features of the market would have implied a five percent decrease in the log Lerner index.³⁰ Figure 8 presents the time series of aggregate log Lerner index as well as the structural components, which are reported relative to 2006 values. The average log Lerner index increased by 0.25 from 2006 to 2019, as shown by the solid black line, while the structural factors decreased from 2006 to 2011 and remained below 2006 levels thereafter. Consistent with our earlier results, this decomposition illustrates that the overall change in markups is tied to changes in consumer price sensitivity.

Why does consumer price sensitivity decline? One possibility is that price-sensitive consumers increasingly select out of mass merchandisers, grocery stores, and drug stores and into other channels that offer lower prices, such as warehouse clubs or dollars stores. However, such an explanation seems to be at odds with aggregate consumer spending patterns. As documented in Table B.1 in the Appendix, the focal channels in our data comprise the vast majority of broad-basket retail spending in 2007 (83 percent) and 2019 (82 percent). Furthermore, supplementary regressions show that category-level increases in spending among these other channels does not meaningfully explain category-level changes in price sensitivity (see Ap-

³⁰In our empirical model, the structural component can be obtained from the first step in our estimation routine, where we pin down heterogeneity in demand using micro-moments (Berry and Haile, 2020). Thus, our finding of decrease in the structural component is not sensitive to price endogeneity and does not rely on the moments used to pin down the mean price parameters. See Appendix A for details.

Figure 8: Decomposition of Markup Trends



Notes: This figure shows the changes to the aggregate log Lerner Index (black line) and the structural factors (dash-dotted line) specified by equation (12). The structural factors incorporate observable changes in prices and the distribution of market shares. The difference between the two lines is captured by changes to price sensitivity.

pendix D). An alternative possibility is that firms make investment decisions that serve to 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 D, we 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 consumer selection across retail channels or firm-level investment decisions.

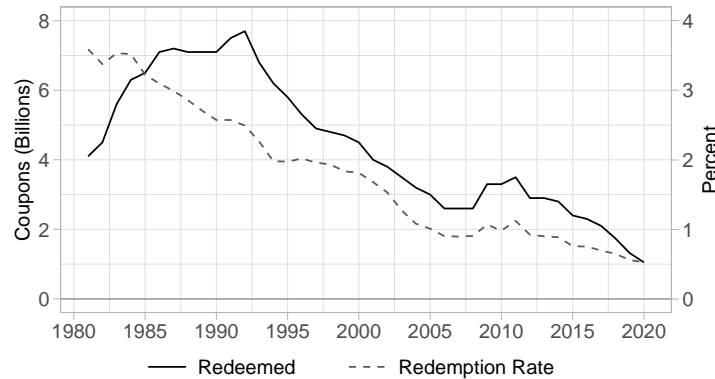
Changes in price sensitivity may reflect exogenous shifts in preferences that are not the result of changes to supply. To explore this possibility, we examine other information about consumer shopping patterns. In particular, we look at the use of coupons and estimates of time spent shopping for consumer 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.³¹

Figure 9 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

³¹Industry estimates were obtained from reports by two companies, NCH Marketing from 1981 through 2002, and Inmar Intelligence from 2003 through 2020.

Figure 9: 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.

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.

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.³²

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. 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.³³

³²The 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.

³³Griffith et al. (forthcoming) 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.

7 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.

Following Small and Rosen (1981), we calculate consumer surplus as the total expected value that consumers receive from a set of products, given the distribution of the consumer-specific logit error terms (but not their realizations). With the observed set of products, consumer surplus is given by:

$$CS = -\frac{1}{N} \sum_i \frac{1}{\alpha_i} \ln \left(\sum_j \exp(w_{ij}) \right) \quad (13)$$

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.³⁴ This represents the additional consumer surplus provided by the inside goods, relative to a counterfactual in which only the outside option is available to consumers (as the outside option alone provides zero consumer surplus by assumption). Thus, it can be interpreted as the added value of the focal products under consideration, or, identically, the equivalent variation that would compensate consumers for the loss of these product-market combinations.³⁵

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 estimation results do not identify fixed costs and, as they are not incorporated into our measure of producer surplus, our results do not inform whether brand manufacturers earn economic profit.³⁶ 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.

Table 5 shows per capita consumer, producer surplus, and welfare for 2006 and 2019 using observed prices (“Baseline”) and prices under different counterfactual scenarios. To compute

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.

³⁴In calculating consumer surplus, we use the average price coefficient within each consumer’s income decile to avoid dividing by numbers very close to zero. In practice, this matters only for a single category, and we obtain nearly identical results if we use the average price coefficient within income quartiles or across all consumers.

³⁵We do not evaluate trends in overall welfare, which would necessitate taking a stance on utility for the outside good. We focus on the relationship between markups and welfare within the products and markets of our sample.

³⁶The findings of De Loecker et al. (2020), which look at firm-level accounting statements, indicate that profits have increased along with markups.

Table 5: Annual Surplus and Welfare Per Capita

(a) 2006 Preferences and Costs

Specification	CS	PS	W	% change CS	% change W
Baseline	628	261	889	0.0	0.0
Prices Scaled to 2019 Levels	603	263	867	-3.8	-2.4
Markups Scaled to 2019 Levels	551	267	818	-12.2	-8.0
Prices Equal to Marginal Costs	956	0	956	52.4	7.6

(b) 2019 Preferences and Costs

Specification	CS	PS	W	% change CS	% change W
Baseline	974	371	1345	0.0	0.0
Prices Scaled to 2006 Levels	1006	350	1356	3.3	0.8
Markups Scaled to 2006 Levels	1106	280	1386	13.5	3.1
Prices Equal to Marginal Costs	1460	0	1460	49.9	8.6

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.

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 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.

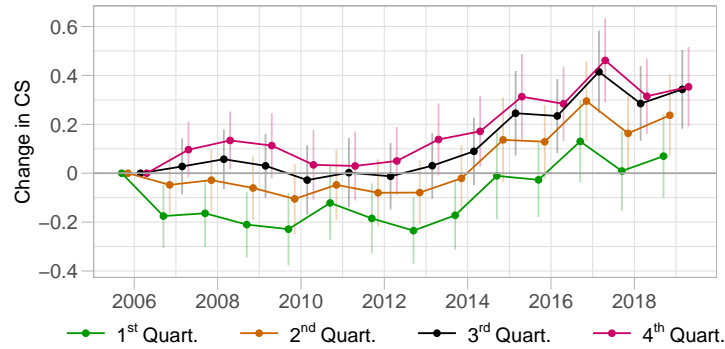
Comparing the baseline scenarios, the results indicate that per capita consumer surplus increased by about 50 percent (i.e., about 3 percent annually) between 2006 and 2019, from \$628 to \$974. As average prices did not decline and perceived quality did not increase, the increase in consumer surplus 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 \$261 to \$371 per capita. Thus, approximately three quarters 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 \$956 to \$889 (about 7 percent). In 2019, markups reduced welfare by about 8 percent.³⁷

The changes in markups over this period are economically meaningful. Holding fixed the 2006 preferences, marginal costs, and product assortments, increasing markups to 2019 lev-

³⁷These estimates of deadweight loss are similar in magnitude to those reported in recent study of publicly-traded firms in the United States (Pellegrino, 2021).

Figure 10: 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.

els would reduce consumer surplus by 12 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 (3.8 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 10. 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 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 F.4 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.

8 Conclusion

This paper analyzes the evolution of market power in consumer products in the United States between 2006 and 2019. For this purpose, we combine retail scanner data on quantities and prices with consumer level data across more than 100 product categories. This approach allows us to 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 increase by about 30 percent during our sample period. In contrast to previous research on the evolution of market power, we estimate similar changes across different quartiles of the markup distribution. In addition, we find similar increases in markups within product categories over time which implies that the results are not driven by a reallocation of market shares towards products with higher markups. We decompose changes in markups into changes in prices and changes in marginal costs. Overall, the nominal prices of products rise at a similar rate as inflation during our sample period. Thus, real prices remain almost constant, and the increase in markups we estimate is primarily due to falling (real) marginal costs. Our results suggest that prices do not decrease along with marginal costs because of changes in consumer preferences. Our estimates suggest that consumers became about 30 percent less price sensitive over the sample period.

The results of a counterfactual simulation exercise indicate that changes in price sensitivity and marginal costs account for nearly all of the time series variation in aggregate markup changes between 2006 and 2019. We also find that these two factors explain most of the cross-category variation in markup trends, while changes in ownership, demographics and perceived quality only play a minor role. Due to decreased price sensitivity, consumer surplus increased during our sample period despite rising markups. The increase in consumer surplus is, however, concentrated among consumers with relatively high income. Nonetheless, changes in markups have been costly for consumers. In a counterfactual simulation, we find that consumer surplus would have been 14 percent higher in 2019 if markups had not changed relative to 2006. If firms would set price equal to marginal costs, consumer surplus in 2019 increases by 50 percent and total welfare increases by 9 percent.

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.
- ANDERSON, E., S. REBELO, AND A. WONG (2018): “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.
- BACKUS, M., C. CONLON, AND M. SINKINSON (2021): “Common Ownership and Competition in the Ready-to-Eat Cereal Industry,” Working paper.
- BASU, S. (2019): “Are Price-Cost Markups Rising in the United States? A Discussion of the Evidence,” *Journal of Economic Perspectives*, 33, 3–22.
- BERRY, S. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25, 242–262.
- BERRY, S., M. GAYNOR, AND F. SCOTT MORTON (2019): “Do Increasing Markups Matter? Lessons from Empirical Industrial Organization,” *Journal of Economic Perspectives*, 33, 44–68.
- 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.
- BERRY, S. T. AND P. A. HAILE (2020): “Nonparametric Identification of Differentiated Products Demand Using Micro Data,” Working paper.
- (2021): “Foundations of Demand Estimation,” *Handbook of Industrial Organization, Volume 4*, edited by Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, Elsevier.
- BHATTACHARYA, V., G. ILLANES, AND D. STILLERMAN (2022): “Have Mergers Raised Prices? Evidence from U.S. Retail,” Working Paper.
- BOND, S., A. HASHEMI, G. KAPLAN, AND P. ZOCH (2021): “Some Unpleasant Markup Arithmetic: Production Function Elasticities and their Estimation from Production Data,” *Journal of Monetary Economics*, 121, 1–14.
- 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 (2022): “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. 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 (2017): “Estimating Market Power: Evidence from 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.
- DE RIDDER, M., B. GRASSI, AND G. MORZENTI (2022): “The Hitchhiker’s Guide to Markup Estimation,” Working Paper.
- DELLAVIGNA, S. AND M. GENTZKOW (2019): “Uniform Pricing in US Retail Chains,” *Quarterly Journal of Economics*, 134, 2011–2084.
- DORASZELSKI, U. AND J. JAUMANDREU (2019): “Using Cost Minimization to Estimate Markups,” Working paper.
- EIZENBERG, A., S. LACH, AND M. OREN-YIFTACH (2021): “Retail Prices in a City,” *American Economic Journal: Economic Policy*, 13, 175–206.
- 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. (2021a): “Growing Oligopolies, Prices, Output, and Productivity,” *American Economic Journal: Microeconomics*, 13, 309–327.
- (2021b): “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 (2022): “The Evolution of Market Power in the US Automobile Industry,” Working paper.
- GRIFFITH, R., W. M. JIN, AND V. LECHENE (forthcoming): “The Decline of Home-cooked Food,” *Fiscal Studies*.
- GRIGOLON, L. AND F. VERBOVEN (2014): “Nested Logit or Random Coefficient Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation,” *Review of Economics and Statistics*, 96, 916–935.
- HALL, R. E. (1988): “The Relation Between Price and Marginal Cost in U.S. Industry,” *Journal of Political Economy*, 96, 921–947.
- 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. (2020): “Vertical Contracts with Endogenous Product Selection: An Empirical Analysis of Vendor-Allowance Contracts,” Working paper.
- 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 (2022): “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., M. OSBORNE, G. SHEU, AND G. SILEO (2022): “The Evolution of Concentration and Markups in the Unites States Cement Industry,” Working paper.
- MILLER, N. H. AND M. C. WEINBERG (2017): “Understanding the Price Effects of the Miller-Coors Joint Venture,” *Econometrica*, 85, 1763–1791.
- 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.
- PELLEGRINO, B. (2021): “Product Differentiation and Oligopoly: A Network Approach,” Working paper.
- PELTZMAN, S. (2020): “Productivity, Prices and Productivity in Manufacturing: a Demsetzian Perspective,” Working paper.
- RAVAL, D. (2020): “Testing the Production Approach to Markup Estimation,” Working paper.
- SMALL, K. A. AND H. S. ROSEN (1981): “Applied Welfare Economics with Discrete Choice Models,” *Econometrica*, 105–130.
- SYVERSON, C. (2019): “Macroeconomics and Market Power: Context, Implications, and Open Questions,” *Journal of Economic Perspectives*, 33, 23–43.
- VILLAS-BOAS, S. B. (2007): “Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data,” *Review of Economic Studies*, 74, 625–652.

Appendix

A 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 and Haile, 2020), and the covariance restriction exactly identifies α given θ_2 (MacKay and Miller, 2022). 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.

A.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.³⁸

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 $\theta_2 = \vec{0}$ the mean valuations have a closed-form solution:

$$\delta_{jcr}t(\theta_2^{(0)}) \equiv \log(s_{jcr}t) - \log(s_{0cr}t) \tag{A.1}$$

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 $\{\delta_{jcr}t(\theta_2)\}$ using the contraction mapping of Berry et al. (1995) with a numerical tolerance of 1e-9. We then calculate the micro-moments with $\{\delta_{jcr}t(\theta_2)\}$ and $\bar{\alpha}$. We choose the parameters $\{\delta_{jcr}t(\theta_2)\}$ that minimize the micro-moment contributions to the objective function. We apply equal weights to each micro-moment in estimation.

³⁸We initialize this step with a price parameter $\bar{\alpha}$ such that the average elasticity when $\theta_2 = \vec{0}$ is equal to -7, which corresponds to the average starting value that we use in the second step (see below).

A.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, 2020). To do so, we recover $\Delta\xi_{jcrct}(\theta_2)$ as the residual from the OLS regression of $(\delta_{jcrct}(\theta_2) - \alpha p_{jcrct})$ 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 $\Delta\eta_{jcrct}(\theta_2)$ 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, 2022). Care must then be taken to ensure that the estimator converges to the smaller value. Figure F.5 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.³⁹

We proceed by selecting starting values of $\alpha^{(0)} = \phi\tilde{\alpha}$ where $\tilde{\alpha}$ is such that the average elasticity is -1 when $\theta_2 = \vec{0}$, 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.

³⁹The 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.

A.3 Computation Notes

Our code builds on the BLPestimatorR package for R (Brunner et al., 2020).⁴⁰ 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 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).⁴¹ 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.

⁴⁰<https://github.com/cran/BLPestimatorR>, last accessed March 26, 2021

⁴¹<https://github.com/jeffgortmaker/pyblp>, last accessed March 26, 2021.

B Data Details

B.1 Market Size Calculations

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

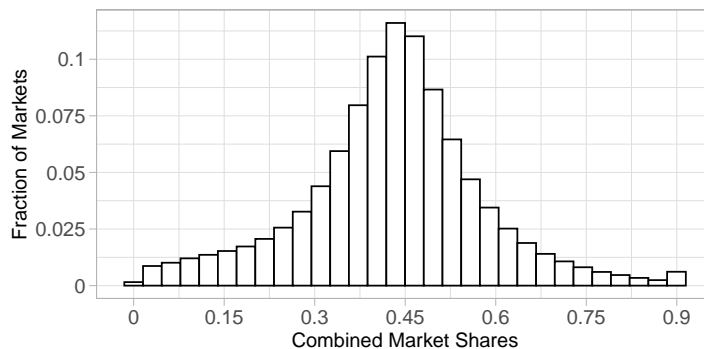
1. Obtain a time-varying “base” value by multiplying the population (at the region-year level) with the number of stores (at the chain-region-quarter-year level). This obtains $BASE_{crqy} \equiv POP_{ry} \times NS_{crqy}$ where POP_{ry} is the population in region r and year y and NS_{crqy} is the number of stores operated by retail chain c in region r , quarter q , and year y .
2. Obtain the total quantity of the inside products across brands: $Q_{crqy} = \sum_j q_{jcrqy}$.
3. Calculate $\gamma_{cr} = E_{q,y} \left[\frac{Q_{crqy}}{BASE_{crqy}} \right]$ as the average quantity-to-base ratio among the periods 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

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

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

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 $\tilde{\gamma}_{cr} = 0.5 \times \max_{q,y} \left(\frac{Q_{crqy}}{BASE_{crqy}} \right)$, which sets the maximum of the combined shares equal to 0.90.

Figure B.1: 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.

Figure B.1 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 provide robustness checks in Appendix E.4.

B.2 Other Notes on Estimation Data

We make a number of adjustments to the Nielsen data as we construct the estimation samples. First, we drop two large chains from the Consumer Panel Data that do not appear in the Retail Scanner Data. Second, 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; we estimate a midpoint of \$137,500. Third, 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 Nielsen 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. Fourth, 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 restricting attention to the 22 DMAs. Fifth, we exclude four categories from the ranking that, for some years, exist in the scanner data but not the consumer panel data: prerecorded videos, magazines, cookware, and sunscreens. Finally, product categories belong to the following high-level departments according to Nielsen: “Dry Grocery,” “Frozen Foods,” “Dairy,” “Deli,” “Packaged Meat,” “Fresh Produce,” and “Alcoholic Beverages,” “Health and Beauty Care,” “Non-food Grocery,” and “General Merchandise.”

B.3 Auxiliary Data on Revenues by Retail Channel

Table B.1: Share of Revenue by Retail Channel

	2007	2019
<i>Focal Channels</i>		
Mass Merchandisers	0.214	0.218
Grocery Stores	0.219	0.217
Drug Stores	0.088	0.117
<i>Other Broad-Basket Retail Channels</i>		
Warehouse Club	0.090	0.094
Dollar Stores	0.015	0.026
<i>Other Consumer Product Retail Channels</i>		
Convenience Stores, Department Stores, Apparel, etc.	0.374	0.328
<i>Combined Share of Focal Channels</i>		
Among All Consumer Products	0.522	0.552
Among Broad-Basket Retailers	0.833	0.822

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.

As described in the main text, we focus our analysis on retailers that Nielsen classifies as mass merchandisers, grocery stores, or drug stores. 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).⁴² 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

⁴²For 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.

consumer products. The included retailers represent \$1.4 trillion in revenues in 2007 and \$2.0 trillion in 2019.

Table B.1 reports the share of consumer product spending in our focal channels (mass merchandisers, grocery stores, and drug stores) and other broad-basket retailers (warehouse clubs and dollar stores) in 2007 and 2019. Our focal channels are the three largest consumer product channels in 2019, and their shares have been fairly stable over our sample period. Combined, the channels represent 83 of spending within broad-basket retailers in 2007 and 82 percent in 2019. Out of all consumer product spending, the focal channels represent 52 percent of spending in 2007 and 55 percent in 2019.

Thus, the focal channels capture the majority of consumer product spending, and their revenue growth has paralleled the average revenue growth among other large U.S. retailers. The largest broad-basket channel that we omit is warehouse club, which accounts for 9.0 percent of revenues in 2007 and 9.4 percent in 2019. The revenue share of dollar stores roughly doubles between 2007 and 2019, consistent with the trend documented in Caoui et al. (2022). Nonetheless, dollars stores account for only 1.5 percent of consumer product spending in 2007 and 2.6 percent in 2019.

The share of revenues allocated to other consumer product channels declined slightly over our sample, from 37 percent in 2007 to 33 percent in 2019. Within this category, 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.

C Derivation of the Econometric Decomposition

In this appendix, we obtain the structural decomposition used in Section 6, following MacKay and Miller (2022). The decomposition is available for a wide class of models, but we focus on random coefficients logit demand with differentiated-products Bertrand competition.

First, it is helpful to re-express the indirect utility that consumer i receives from product $j > 0$ (in chain c , region r , and quarter t) as follows:

$$u_{ijcrt} = \delta_{jcrt}(p_{jcrt}; \beta, \alpha) + \mu_{ijcrt}(p_{jcrt}, D_i, v_i; \Pi_1, \Pi_2, \sigma) + \epsilon_{ijcrt} \quad (\text{C.1})$$

where the mean utility of each product, $\delta_{jcrt}(\cdot)$, and contribution of demographics to consumer-specific deviations, $\mu_{ijcrt}(\cdot)$, respectively are given by

$$\begin{aligned} \delta_{jcrt}(p_{jcrt}; \beta, \alpha) &= \beta + \alpha p_{jcrt} + \xi_{jr} + \xi_{cr} + \xi_t + \Delta \xi_{jcrt} \\ \mu_{ijcrt}(p_{jcrt}, D_i, v_i; \Pi_1, \Pi_2, \sigma) &= p_{jcrt} \Pi_1 D_i + \Pi_2 D_i + \sigma v_i \end{aligned}$$

The indirect utility of the outside good remains $u_{i0crt} = \epsilon_{i0crt}$. The probability with which consumer i selects product j can be expressed

$$s_{ijcrt}(\delta_{crt}, p_{jcrt}, D_i, v_i; \Pi_1, \Pi_2, \sigma) = \frac{\exp(\delta_{jcrt}(p_{jcrt}; \beta, \alpha) + \mu_{ijcrt}(p_{jcrt}, D_i, v_i; \Pi_1, \Pi_2, \sigma))}{1 + \sum_{k=1}^{J_{crt}} \exp(\delta_{kcrt}(p_{kcrt}; \beta, \alpha) + \mu_{ikcrt}(p_{kcrt}, D_i, v_i; \Pi_1, \Pi_2, \sigma))} \quad (\text{C.2})$$

where $\delta_{crt} = (\delta_{1crt}, \delta_{2crt}, \dots)$ is the vector of mean utilities. Finally, the market share of product j is obtained by integrating over the joint distribution of consumer demographics:

$$s_{jcrt}(\delta_{crt}, p_{jcrt}; \Pi_1, \Pi_2, \sigma) = \frac{1}{I} \sum_i s_{ijcrt}(\delta_{crt}, p_{jcrt}, D_i, v_i; \Pi_1, \Pi_2, \sigma)$$

For a broad class of oligopoly models, the first order conditions for profit maximization can be expressed in terms of product-level additive markups as follows:

$$p_{jcrt} - c_{jcrt}(\chi_{crt}; \theta) = -\frac{1}{\alpha} \lambda_{jcrt}(q_{crt}, p_{crt}, \Gamma_{crt}; \theta^*), \quad (\text{C.3})$$

where q_{crt} and p_{crt} are vectors of quantities and prices (typically data), Γ_{crt} denotes the matrix of demand derivatives, and θ^* includes all the demand parameters *except* for the mean price parameter (α). Let the set of products sold by the same firm as product j be given by $\mathcal{J}_{f(j)}$. Then, with random coefficients logit demand and Bertrand competition, we have:

$$\lambda_{jcrt} = \frac{s_{jcrt}}{\frac{1}{I} \sum_i s_{ijcrt} (1 - s_{ijcrt})} - \sum_{k \in \mathcal{J}_{f(j)} \setminus j} \frac{s_{kcrt}}{\frac{1}{I} \sum_i s_{ijcrt} s_{ikcrt}} \quad (\text{C.4})$$

where the denominators integrate over the (product of) consumer-specific choice probabilities. From an econometric standpoint, $\lambda_{j_{crt}}$ is free from the mean price parameter (α) because it depends only on market shares and consumer-specific choice probabilities. The market shares are data. From equation (C.2), the consumer-specific choice probabilities depend on $\mu_{crt}(\cdot)$, which obtains immediately from data and $\theta^* = (\Pi_1, \Pi_2, \sigma)$, and on $\delta_{crt}(\cdot)$, which obtains from the contraction mapping of Berry et al. (1995), again given data and θ^* . Related is the observation of Berry and Haile (2020) that micro-moments summarizing how demographics affects consumer choice patterns cannot identify the mean price parameter.

D Additional Regression Analysis

Given the important role of price sensitivity in markups, we next examine potential factors that could explain the change over time. In this appendix, we consider changes to the composition of retail markets and firm-level decisions that could affect consumer behavior in the studied markets. In the main text, we provide evidence that consumers are becoming less price sensitive over time due to exogenous factors (Section 6).

To assess changes in the composition of retail markets, we construct the share of revenues by retail channel in each product category and each year, including warehouse clubs, dollar stores, and online retail, in addition to mass merchandisers, grocery, and drug stores. We use all available data from the Kilts Nielsen consumer panel dataset to construct these measures. We then test whether changes in price sensitivity are correlated with growth in any of the additional three segments, leveraging cross-category idiosyncratic variation over time. 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.

As documented in Appendix B.3, the aggregate compositional shifts in our focal channels are fairly small. Here, using the consumer panel data, we obtain similar channel shares. Within these six channels, 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. The focal channels capture 83.9 percent share on average in 2019.

The next category of mechanisms we investigate addresses 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. 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.

Column (1) of Table D.1 reports the results. We find no significant relationships between share sold in warehouse clubs, marketing expenditures, or R&D expenditures. We find a nega-

Table D.1: Potential Mechanisms

	(1) Price Sensitivity	(2) Log Abs. Elasticity	(3) Marginal Cost	(4) Perceived Quality
Log Share Warehouse Clubs	-0.014 (0.064)	0.023 (0.063)	0.142 (0.193)	-0.090 (0.158)
Log Share Dollar Stores	0.064** (0.029)	0.058** (0.028)	0.049 (0.079)	0.073 (0.084)
Log Share Online	-0.090* (0.047)	-0.074* (0.041)	-0.121 (0.136)	-0.449*** (0.144)
Log Marketing Spend	0.012 (0.021)	0.017 (0.020)	0.125** (0.054)	0.049 (0.056)
Log R&D	-0.006 (0.023)	-0.006 (0.020)	-0.059 (0.057)	0.016 (0.072)
Log Num. UPCs	0.100* (0.051)	0.089* (0.046)	0.384*** (0.127)	0.455*** (0.155)
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	0.943	0.603	0.122	0.173
R^2 (Within)	0.015	0.013	0.015	0.028

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

tive, marginally significant 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 the shares of these channels (less than 3 percent for the categories in 2019), we think these results most likely reflect other mechanisms, e.g., online retailers entering categories with higher markups and less price sensitive consumers. To that end, a regression with price elasticity as the dependent variable, reported in column (2), returns a coefficient on online sales is that roughly 20 percent smaller. If online sales were skimming off more price sensitive consumers, we would expect consumers to self-select based on demographic characteristics that also affect demand elasticities. In this case, we would expect elasticities to have a stronger relationship with online sales than the (mean) price sensitivity parameter, which we do not find.

We find a marginally significant positive relationship between variety and price sensitivity, which indicates that greater variety is weakly correlated with *greater* price sensitivity.⁴³ 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 1.5 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

⁴³Brand (2021) finds the opposite relationship.

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 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 5. 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.⁴⁴

Overall, the regression results suggest that the composition of retail shopping across new channels and firm-level investments cannot account for the change in consumer price sensitivity that we document.

⁴⁴This 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 3). Thus, product variety does not appear to be driving the trends we observe.

E Robustness Checks

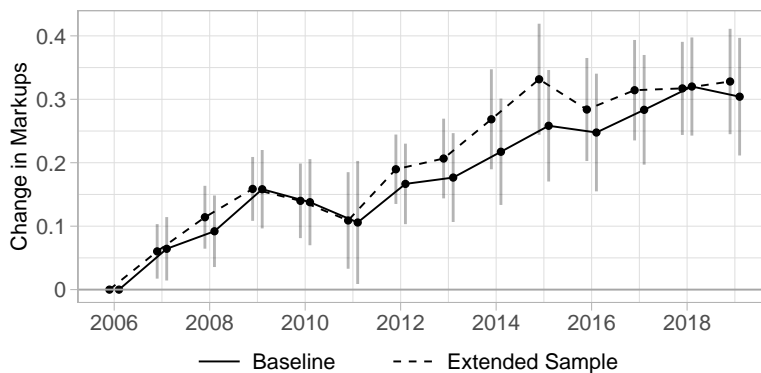
In this section, we present a series of robustness checks to evaluate the sensitivity of our main findings to particular assumptions. First, we show how the main trend in markups is not sensitive to particular choices of measurement, in terms of which categories are included in our baseline sample and our choice of the Lerner index as our markup measure. We then show that 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 also find very similar trends in markups and price sensitivity with different approaches to measuring market size.

Finally, 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.⁴⁵ 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 4.2), a trend in these parameter would be consistent with a rotation of the demand curve. We find a similar decline in the mean price parameter under this alternative assumption, indicating that our findings of falling price sensitivity are robust to the particular supply-side assumptions we invoke in estimation.

⁴⁵As described in the text, the other demand-side parameters are identified by micro-moments.

E.1 Category Selection

Figure E.1: Markups Over Time: Alternative Samples



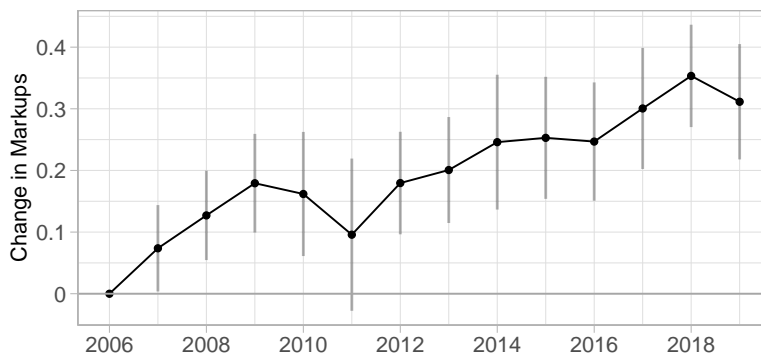
Notes: This figure displays the changes in product-level markups over time for our baseline sample (133 product categories, solid line) and the extended sample (200 product categories, dashed line). The 133 product categories in the baseline sample are selected based on a proxy for within-category product heterogeneity. Point estimates and 95 percent confidence intervals are obtained from regressions of the log of the Lerner index $(p - c)/p$ 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.

In Section 3, 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.1, 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 very similar trends in markups with either selection procedure, with a change of approximately 30 log points from 2006 to 2019.

E.2 Markup Measure

Figure E.2: Markups Over Time: Price-Over-Cost Markups



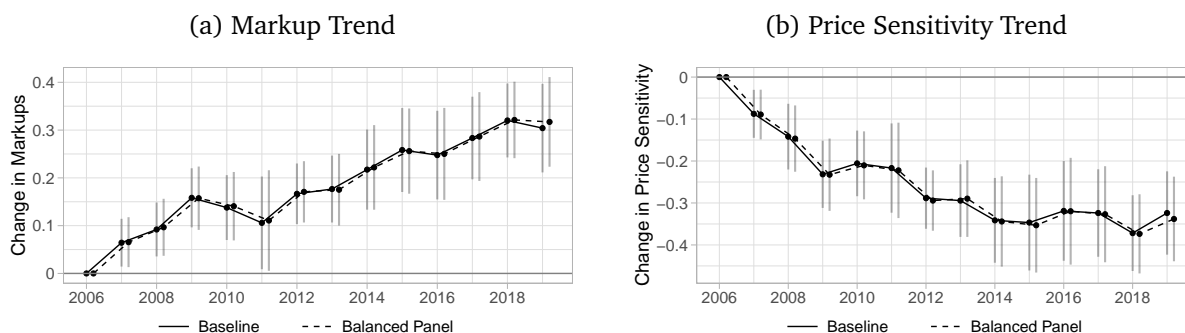
Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. Markups are defined as price over marginal 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.2 replicates our product-level markup trends, corresponding to panel (a) of Figure 5 in the main text, using the price-over-cost markup measure. The trends are nearly identical.

E.3 Balanced Panel

Figure E.3: Balanced Panel



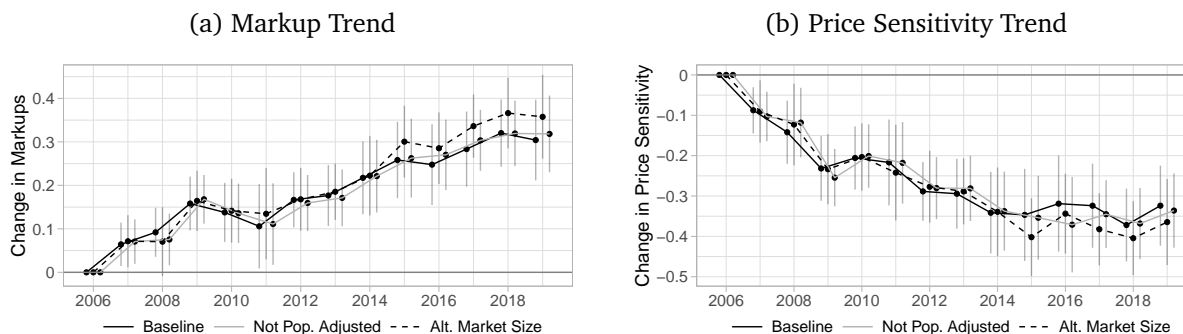
Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. 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 $\text{brand} \times \text{chain} \times \text{region}$.

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, some compositional changes in the Nielsen data occur during our sample period due to coverage of certain retail chains. Although our demand estimation controls for $\text{chain} \times \text{region}$ 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.3, we therefore replicate trends of markups and price sensitivities using a balanced panel of $\text{brand} \times \text{chain} \times \text{region}$ combinations. The trends are very similar to those reported in panel (a) of Figure 5 and panel (b) of Figure 6. The baseline trends are reproduced in the figure for comparison.

E.4 Market Size

Figure E.4: Alternative Market Size Measures



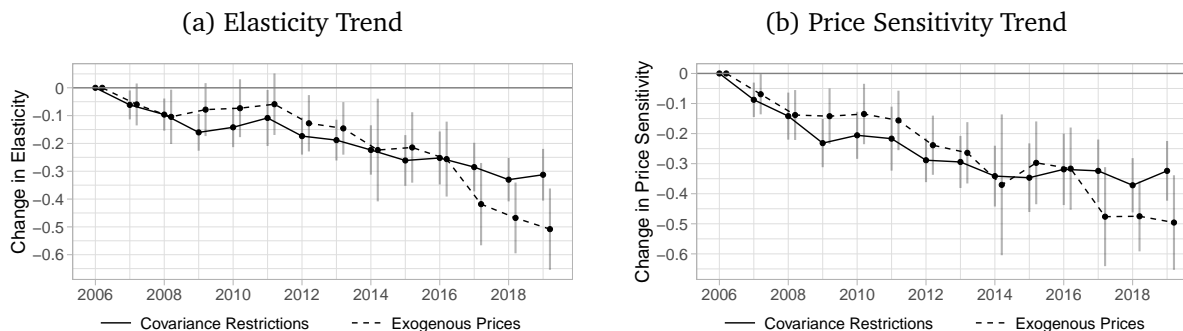
Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. The baseline estimates are plotted with a solid black line. The 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 that use an alternative value for the average market size.

As discussed in Section 2.2, we need an assumption about market size to measure market shares of products. In Appendix B.1, 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.

To check the robustness of our results towards assumptions about the relevant market, we reran our demand estimation using two alternative definitions of market size. First, we rescale market size to obtain an average combined market share of inside goods of 0.3. Second, we assume that market size does not vary with population growth. Figure E.4 shows that these alternative assumptions lead to very similar trends in markups and price sensitivity. Thus, the trends we estimate do not hinge on the precise definition of market size.

E.5 Changes in Demand Over Time

Figure E.5: Changes in Demand Over Time



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.

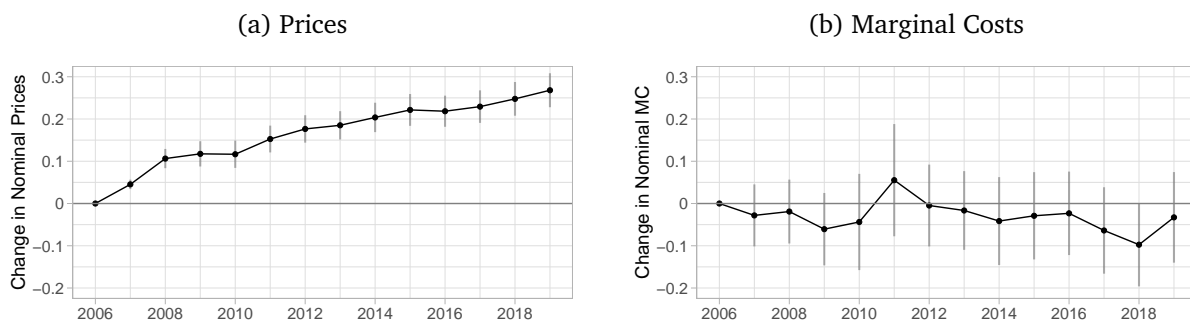
We show that a similar trend is obtained when we estimate demand using the assumption that prices are exogenous, which does not invoke the supply model to pin down the demand parameters. Though elasticity estimates under this approach are often unreasonable in terms of levels (see Section 4.2), a change in the estimated parameters would be consistent with a rotation of the demand curve.

Figure E.5 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

⁴⁶Of 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.

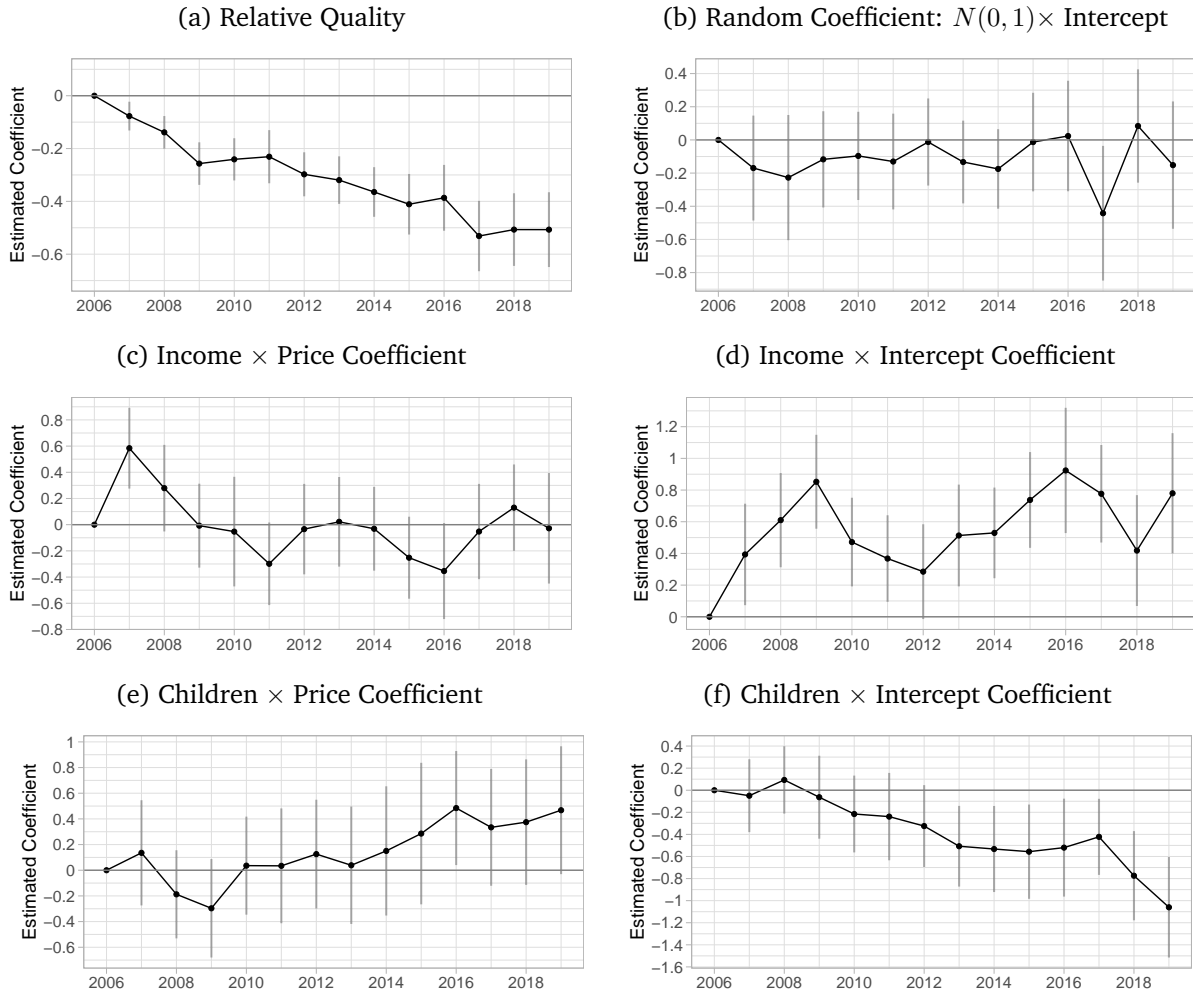
F Additional Figures and Tables

Figure F.1: Product-Level Changes in Nominal Prices and 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.

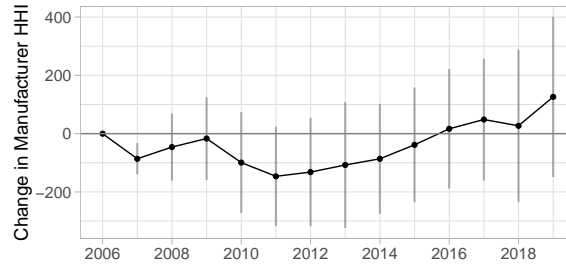
Figure F.2: Changes in Demand Parameters



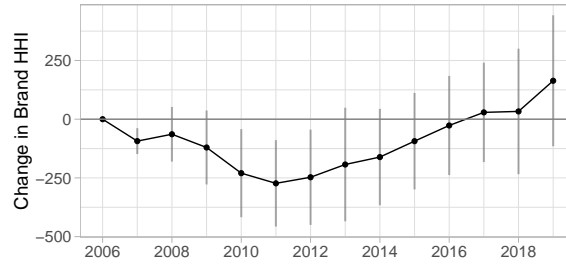
Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of standardized demand parameters on year dummies controlling for product-chain-DMA and quarter fixed effects. Observations are at the product-chain-DMA-quarter-year level. The year 2006 is the base category.

Figure F.3: Changes in Market Concentration

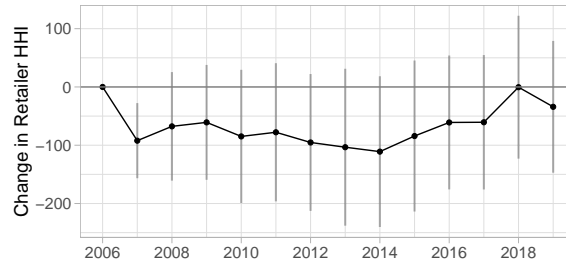
(a) Parent HHI



(b) Brand HHI

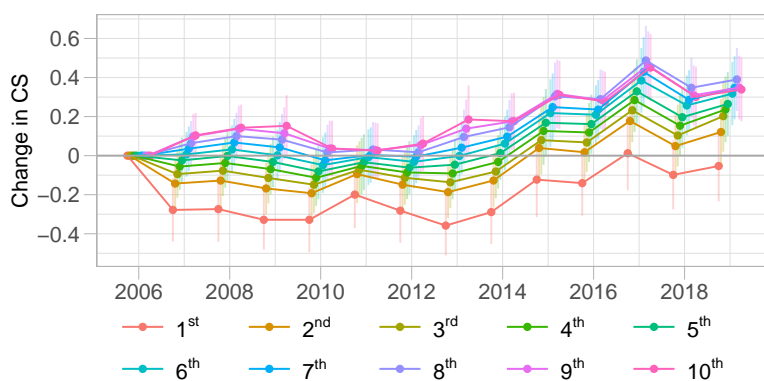


(c) Retailer HHI



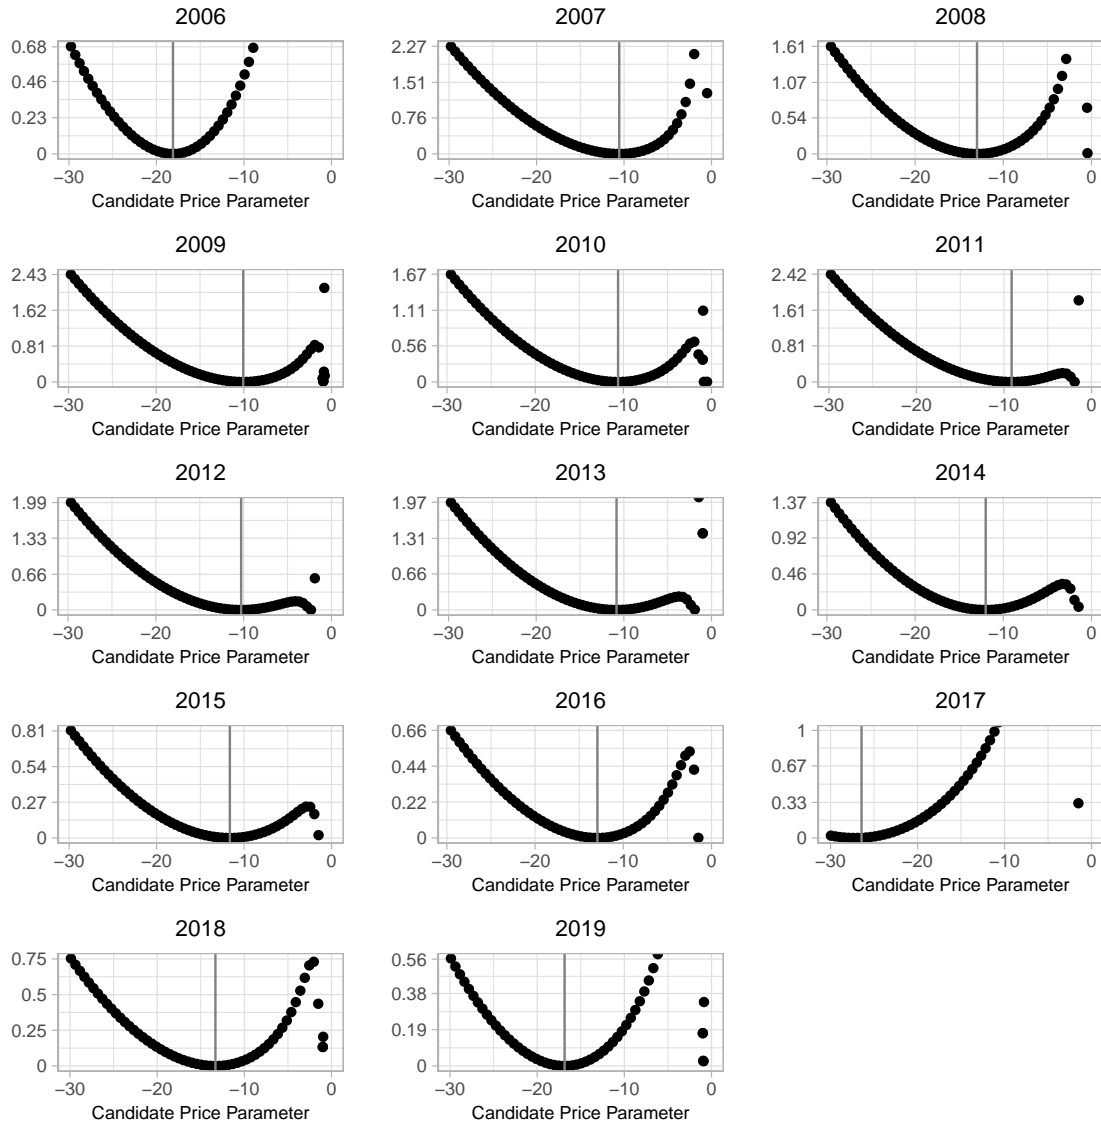
Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of HHI measures on year dummies controlling for product-chain-DMA and quarter fixed effects, with 2006 as the base category. We measure HHI as the sum of squared market shares, where we first adjust market shares 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 F.4: 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 F.5: 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 F.1: Product-Level Markups Over Time, Sales-Weighted Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Markup	Log Markup	Log Markup	Log Markup	Log Markup	Log Markup
Trend	0.018*** (0.003)		0.017*** (0.003)		0.022*** (0.003)	
Year=2007		0.061** (0.025)		0.063** (0.025)		0.064** (0.025)
Year=2008		0.092*** (0.028)		0.093*** (0.028)		0.092*** (0.029)
Year=2009		0.164*** (0.031)		0.163*** (0.031)		0.158*** (0.031)
Year=2010		0.144*** (0.034)		0.142*** (0.034)		0.138*** (0.034)
Year=2011		0.108** (0.048)		0.105** (0.048)		0.106** (0.049)
Year=2012		0.169*** (0.032)		0.161*** (0.033)		0.167*** (0.032)
Year=2013		0.180*** (0.036)		0.169*** (0.036)		0.177*** (0.035)
Year=2014		0.214*** (0.042)		0.203*** (0.043)		0.217*** (0.042)
Year=2015		0.243*** (0.045)		0.230*** (0.045)		0.258*** (0.044)
Year=2016		0.222*** (0.048)		0.208*** (0.047)		0.248*** (0.047)
Year=2017		0.253*** (0.044)		0.238*** (0.043)		0.283*** (0.044)
Year=2018		0.277*** (0.041)		0.274*** (0.040)		0.320*** (0.039)
Year=2019		0.255*** (0.049)		0.253*** (0.047)		0.304*** (0.047)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	14,407,410	14,407,410	14,407,410	14,407,410	14,407,410	14,407,410
R^2	0.013	0.014	0.357	0.359	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 F.2: Product-Level Markups Over Time, Unweighted Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Markup	Log Markup	Log Markup	Log Markup	Log Markup	Log Markup
Trend	0.018*** (0.003)		0.020*** (0.003)		0.023*** (0.003)	
Year=2007		0.079*** (0.028)		0.078*** (0.028)		0.086*** (0.028)
Year=2008		0.098*** (0.029)		0.095*** (0.029)		0.111*** (0.030)
Year=2009		0.157*** (0.034)		0.154*** (0.034)		0.175*** (0.035)
Year=2010		0.145*** (0.038)		0.142*** (0.037)		0.165*** (0.038)
Year=2011		0.103** (0.040)		0.099** (0.039)		0.125*** (0.040)
Year=2012		0.176*** (0.037)		0.173*** (0.037)		0.202*** (0.037)
Year=2013		0.192*** (0.033)		0.187*** (0.032)		0.215*** (0.033)
Year=2014		0.226*** (0.042)		0.223*** (0.040)		0.254*** (0.041)
Year=2015		0.306*** (0.052)		0.310*** (0.050)		0.345*** (0.052)
Year=2016		0.258*** (0.046)		0.262*** (0.045)		0.298*** (0.047)
Year=2017		0.278*** (0.049)		0.288*** (0.046)		0.326*** (0.049)
Year=2018		0.269*** (0.039)		0.286*** (0.038)		0.322*** (0.040)
Year=2019		0.227*** (0.041)		0.247*** (0.041)		0.282*** (0.042)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	14,407,410	14,407,410	14,407,410	14,407,410	14,407,410	14,407,410
R ²	0.011	0.014	0.353	0.356	0.760	0.763

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

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Markup	Log Markup	Log Markup	Log Markup	Log Markup	Log Markup
Trend	0.021*** (0.003)		0.019*** (0.003)		0.022*** (0.003)	
Year=2007		0.064** (0.027)		0.064** (0.026)		0.065** (0.026)
Year=2008		0.100*** (0.030)		0.100*** (0.030)		0.097*** (0.030)
Year=2009		0.168*** (0.033)		0.166*** (0.033)		0.157*** (0.034)
Year=2010		0.155*** (0.036)		0.151*** (0.036)		0.141*** (0.036)
Year=2011		0.126** (0.053)		0.120** (0.053)		0.111** (0.053)
Year=2012		0.188*** (0.032)		0.178*** (0.033)		0.171*** (0.032)
Year=2013		0.192*** (0.038)		0.180*** (0.039)		0.175*** (0.038)
Year=2014		0.235*** (0.043)		0.222*** (0.046)		0.222*** (0.045)
Year=2015		0.263*** (0.045)		0.248*** (0.046)		0.256*** (0.045)
Year=2016		0.246*** (0.050)		0.232*** (0.050)		0.250*** (0.049)
Year=2017		0.276*** (0.047)		0.259*** (0.047)		0.286*** (0.047)
Year=2018		0.309*** (0.041)		0.291*** (0.041)		0.321*** (0.040)
Year=2019		0.301*** (0.047)		0.283*** (0.046)		0.317*** (0.047)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	4,821,264	4,821,264	4,821,264	4,821,264	4,821,264	4,821,264
R ²	0.018	0.019	0.398	0.399	0.764	0.766

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Marginal Cost (Standardized)	-0.238*** (0.010)					-0.137*** (0.010)	-0.135*** (0.010)
Price Sensitivity		-0.667*** (0.027)				-0.406*** (0.037)	-0.408*** (0.037)
Quality (Standardized)			-0.203*** (0.011)			-0.000 (0.008)	0.001 (0.008)
Income (Log)				-2.373 (2.190)		-0.391 (0.773)	-0.467 (0.785)
Children at Home				-5.100 (6.916)		-2.296 (2.838)	-2.707 (2.711)
Parent HHI					1.042*** (0.358)		0.487*** (0.133)
Brand HHI					-0.374 (0.294)		-0.049 (0.116)
Retailer HHI					1.705** (0.850)		0.439 (0.318)
Category FEs	X	X	X	X	X	X	X
Year FEs	X	X	X	X	X	X	X
Observations	1,862	1,862	1,862	1,862	1,862	1,862	1,862
R^2 (Within)	0.707	0.726	0.496	0.002	0.016	0.848	0.852

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