

Rising Markups and the Role of Consumer Preferences*

Hendrik Döpper

DICE[†]

Alexander MacKay

Harvard University[‡]

Nathan H. Miller

Georgetown University[§]

Joel Stiebale

DICE[¶]

October 13, 2021

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 25 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

*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 the Heinrich Heine University of Düsseldorf.

[†]Düsseldorf Institute for Competition Economics (DICE), Heinrich-Heine University Düsseldorf. Email: doep-per@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 associated 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.

Our approach is to estimate differentiated-products demand systems for more than 100 consumer product categories using prices, quantities, and consumer demographics. With demand estimates in hand, we follow Rosse (1970) and impute the marginal costs and markups that rationalize observed prices under the assumption of profit maximization. We repeat this procedure separately for each year over the period 2006–2019, allowing us to examine the evolution of markups in product categories that differ along important observable dimensions, such as changes in consumer price sensitivity, market concentration among manufacturers and retailers, and the income of consumers purchasing the products. 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) or beer (Miller and Weinberg, 2017). We implement the methodology at scale to obtain markups for hundreds of products.

Using the Lerner index as our measure of markups, we find that average markups increase from 0.39 to 0.50 between 2006 and 2019, an increase of more than 25 percent.² 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 ini-

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). So long as marginal cost does not exceed price, it can take values from zero to one.

tial markups; however, in percentage terms, the changes that we estimate are similar for high- and low-markup products. Thus, we interpret our results as suggesting that the full distribution of markups may be shifting upward over time.

By construction, 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, real prices increase by seven percent on average but, by 2019, average real prices are only two percent higher than in 2006. Although price increases partially account for rising markups initially, by the latter years of the sample, marginal cost reductions account for nearly all of the aggregate markup trend. We estimate that the average consumer price sensitivity has declined by about 25 percent from 2006 to 2019, which can explain why declines in marginal costs have not been passed on to consumers in the form of lower prices.

We exploit the panel structure of our data to explore potential mechanisms in greater detail. Controlling for product and time fixed effects, we find that products with larger increases in markups tend to have (i) greater reductions in consumer price sensitivity, (ii) greater marginal cost reductions, and (iii) larger increases in market concentration. Although each of these effects are statistically significant, the first two account for substantially more of the variation in markups. In an initial attempt to explain these changes at a deeper level, we compute firm-level statistics and (for the publicly-traded firms) match to Compustat data on marketing and R&D expenditures. We find that the decreases in price sensitivity are correlated with greater marketing expenditures. The reductions in marginal cost do not correlate with either variable at statistically significant levels. Finally, neither the changing price sensitivity nor the changing marginal costs appear to correlate with a measure of product proliferation.

In our final analyses, we explore changes in consumer surplus and welfare over time. Our findings indicate that—despite the increase in markups—consumer surplus per capita has increased by about 16 percent during our sample period. The increase in consumer surplus is likely due to changing preferences, particularly lower price sensitivity. Nonetheless, changes in markups have been costly for consumers. Our results suggest that if markups had not changed since 2006, consumer surplus would be almost 13 percent higher in 2019. If firms would price at marginal costs, consumer surplus would increase substantially, and total welfare would increase by more than 7 percent.

The evolution of consumer surplus varies considerably across the income distribution. While consumers with income 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 as they had in 2006. These results suggest that changes in markups and consumer preferences had important distributional consequences.

Our research contributes to a growing empirical literature on the evolution of markups in

recent decades. We obtain results that are broadly consistent with De Loecker et al. (2020), although our methodology is different and complementary, as we detail in the next section. As we recover markups from demand elasticities, from a methodological standpoint our research is more similar to Ganapati (2021b) and Grieco et al. (2021). Of these, the former examines wholesalers and finds that an increase in scale economies has led to higher markups; the latter examines automobile manufacturers and finds that markups fall due to increasing competition. Relative to these papers, we consider a broader class of products and introduce the possibility that markups could change due to changes in consumer preferences. Even closer is Brand (2021), which recovers markups in nine of the categories we consider. Brand also finds evidence of rising markups, and attributes them to an increase in product variety.

The paper proceeds as follows: In Section 2, we discuss our approach for recovering markups, specify the model of demand and supply, and describe the estimator and our identification strategy. We discuss the data set in Section 3 and illustrate our empirical approach using a case study for the ready-to-eat (RTE) cereals industry. Section 4 describes the evolution of markups over time and studies possible determinants of market power. In Section 5, we calculate consumer surplus and welfare over time for different scenarios. Section 6 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 the 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 (Lerner, 1934), one commonly-used measure of markups. Knowledge of the demand elasticity identifies the Lerner index. With data on price, one also can recover marginal cost, additive markups (i.e., $P - C'$), and multiplicative markups (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, including two recent contributions on the evolution of markups (Ganapati, 2021b; Grieco et al., 2021). A significant 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. (2020) demonstrates that markups may not be identified if revenue is used as a proxy for output.³ 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).

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).⁴ 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. That said, there may be some product categories for which our assumptions are inappropriate, due to collusive pricing (e.g., Miller and Weinberg, 2017; Miller et al., 2021) or inter-temporal price discrimination (Hendel and Nevo, 2006a,b), for example. Our strategy to mitigate any such misspecification bias is to aggregate results across many product categories, which should be effective under the assumption that Bertrand competition usually provides a reasonable characterization of pricing behavior. Implemented at scale, this allows us to explore how markups have evolved, the reasons for any such changes, and the consequences for consumers and firms.

³See also Doraszelski and Jaumandreu (2019).

⁴One 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

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

where x_j is a vector of observable product characteristics that includes a constant, 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 an unobserved demand shock, 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

$$\begin{bmatrix} \beta^* \\ \alpha^* \end{bmatrix} = \begin{bmatrix} \beta \\ \alpha \end{bmatrix} + \Pi D_i + \Sigma v_i, \quad v_i \sim N(0, I_{K+1}) \quad (3)$$

where D_i contains the observed demographics, v_i contains the unobserved consumer demographics, and (Π, Σ) are matrices of parameters. We assume that Σ is diagonal. Putting equations (2) and (3) together, and allowing β to be absorbed by the product fixed effects, the structural parameters to be estimated are $\theta = (\alpha, \Pi, \Sigma)$. We restrict the final element of Σ to zero, as we find that this parameter is difficult to identify in practice. Thus, the heterogeneity in the consumers' price coefficients is due to the observed demographic variables.

With demographics normalized to have a mean of zero, an average consumer has a “willingness-to-pay” for any product j that equals

$$\overline{WTP}_j = \frac{1}{\alpha} (\xi_{jr} + \xi_{cr} + \xi_t + \Delta \xi_{jcrt}). \quad (4)$$

We interpret \overline{WTP}_j as a measure of the product's perceived quality, relative to the perceived quality provided by the outside good. As we implicitly normalize the quality of the outside good to zero in every year, improvements in the outside good—which includes shopping through online retailers for example—could manifest as declines in the qualities of the inside goods.

The quantity demanded is given by $q_{jcrt}(p_{rct}; \theta) = s_{jcrt}(p_{rct}; \theta) M_{rct}$, where $s(\cdot)$ is the market share, p_{rct} is a vector of prices, and M_{rct} is the “market size” of the chain-market-period,

a measure of potential demand. We refer readers to (Nevo, 2000, 2001) for equations that characterize market shares and the demand elasticities. We use market sizes that scale with the population of the region and the number of stores operated by the retail chain within the region, as we detail in Appendix B. Qualitatively similar results obtain with market sizes that are proportional to the maximum observed sales within a retail chain and region.

Our specification accommodates vertical differentiation among the inside goods, allowing higher quality (and thus more expensive) products to attract relatively price-insensitive consumers. It also incorporates heterogeneity in the utility that consumers receive from the inside goods, so different product categories can appeal to different consumers. This is implemented by placing a random coefficient on the constant in x_j . Thus, we allow the data to determine the extent of substitution between the inside and outside goods.⁵ In principle, an arbitrary number of product characteristics could be incorporated to better capture the horizontal differentiation that exists within the categories. We do not pursue this because the characteristics would have to be extracted from the Nielsen product descriptions or obtained from auxiliary datasets, and both options are difficult to implement at scale. We view the loss of flexibility as mostly inconsequential given our focus on markups. An improved understanding of cross-elasticities would be more valuable for research about mergers or other changes in product ownership.

2.3 Supply Model

We model supply as profit-maximizing firms that set prices each quarter. Consumer products are produced by manufacturers and sold through retail chains. For the most part, these manufacturers and retailers have long-standing relationships that involve contractual payments, advertising, periodic sales, and other joint efforts. With this in mind, our baseline model assumes a form of joint profit maximization in which the retail prices for each manufacturer’s products are set to maximize revenue less the combined costs of the manufacturer and retailer, conditional on the retail prices of other manufacturers. We do not view this assumption as critical however. For example, the markups we recover are isomorphic to the manufacturer markups that would obtain with cost-plus retail pricing (Miller et al., 2021).⁶

The first order conditions for profit maximization can be expressed in matrix notation as

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

⁵An 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; Miller and Weinberg, 2017). With the RCNL model, the speed of estimation slows dramatically for higher values of the nesting parameter, making the model inappropriate for our application.

⁶We plan to consider alternative supply-side arrangements in future versions of this manuscript, two-part tariffs, and a model of profit-maximizing linear pricing with double marginalization. We do not envision modeling competition between different retail chains, beyond the inclusion of the outside good.

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. We assume that marginal costs are constant in output. Thus, marginal cost can be decomposed 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 cost shock that includes the changes in materials costs (e.g., due to factor prices), packaging costs, and distribution costs that affect products and chains differentially.

2.4 Estimation and Identification

We estimate the model separately for each category and year using GMM. Thus, we allow θ and the fixed effects to vary across categories and years. The GMM estimator for θ is:

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

where W is a weighting matrix, $g^{CR}(\theta)$ implements a covariance restriction between unobserved demand and cost shocks, and $g^{MM}(\theta)$ 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. There are a number of important details that arise with the estimation of random coefficients logit models. However, as our implementation is fairly standard, we focus discussion on the key assumptions that identify the structural parameters, and defer the discussion of estimation details to Appendix A.

The first moment—the covariance restriction—is that the demand and cost shocks are uncorrelated in expectation. The role of the moment in estimation is to address the endogeneity of prices in the demand system, i.e, that firms adjust prices in response to unobserved demand shocks. We construct the empirical analog of the moment condition by averaging across the covariances that exist in each chain-region-quarter combination:

$$g^{CR}(\theta) = \frac{1}{T} \sum_{c,r,t} \Delta\xi_{crt}(\theta)' \Delta\eta_{crt}(\theta) \quad (8)$$

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 combinations for a given year.

As the application of covariance restrictions in demand estimation is somewhat novel, we provide a discussion to build intuition. The key insight is that oligopoly models of price competition and constant marginal costs can be recast in terms of residual supply and demand schedules, with equilibrium being the prices and quantities that equate the schedules for each

product (MacKay and Miller, 2019). The slopes of the schedules are determined by the price parameter. All else equal, a price parameter that is more negative corresponds to flatter inverse demand schedules (i.e., more price sensitive consumers) and flatter inverse supply schedules (i.e., less market power). Uncorrelated shifts in these curves generate more variation in quantity than price. With a price parameter that is less negative, by contrast, uncorrelated shifts generate more variation in price than quantity. Thus, the covariance restriction allows the relative variation in price and quantity in the data to identify the price parameter.

Resolving price endogeneity in this manner is particularly helpful in scaling estimation over so many different product categories. The main alternatives are to construct instruments based on the product characteristics of competing products (e.g., Berry et al., 1995; Gandhi and Houde, 2020), based on the cost of inputs (e.g., Nevo, 2001), or based on mergers or other idiosyncratic events that affect prices (e.g., Miller and Weinberg, 2017). None of these strategies are easy to employ at scale. Further, care must be taken to confirm that the instruments satisfy the first-stage relevance condition in each estimation sample. By contrast, the covariance restriction does not require auxiliary data and a first-stage relevance condition does not exist, as the variation in the endogenous price and quantity data is exploited in estimation.

The validity of either approach—covariance restrictions or instruments—hinges on an orthogonality condition involving one or more unobserved structural error terms. We view our assumption as reasonable for most consumer products. We address the most obvious threats to validity with fixed effects. For example, higher quality products may be more expensive to produce, but with product fixed effects this does not confound identification. Similarly, time fixed effects address trends in quality and costs over time. The covariance restriction would be violated if marginal costs vary with output, but the assumption of constant marginal costs is reasonable for most consumer products, and it is commonly maintained in the literature (Nevo, 2001; Villas-Boas, 2007; Chevalier et al., 2003; Hendel and Nevo, 2013; Miller and Weinberg, 2017; Backus et al., 2021). Finally, we compare the elasticities generated under the covariance restriction to those reported elsewhere in the literature as an ex post check.⁷

The second set of moments—the micro-moments—serve to identify the demographic parameters (Π, Σ) . We construct the micro-moments at the level of the product and demographic. The element corresponding to product j and demographic k is

$$g_{jk}^{MM}(\theta) = \frac{1}{T_j} \sum_{c,r,t} \left(\frac{\sum_i \omega_i s_{ijcrt}(\theta) D_{ik}}{\sum_i \omega_i s_{jcrt}(\theta)} - \mathcal{M}_{jrk} \right) \quad (9)$$

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

⁷See Section 3.3.

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).⁸ In our baseline specification, we use two observed demographic variables and at most 21 products, so there can be up to 42 micro-moments.

3 Data and Sample Selection

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, again at the UPC code, store, and week level, 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 in the data, and that group together UPC codes that consumer might reasonably view as substitutes. Within these categories, we define products at the brand level, which consolidates the thousands of UPC codes into a more manageable set. Each UPC code is associated with a “unit” that characterizes its volume (e.g., liters), mass (e.g., ounces), or count (e.g., six-pack). We weight the UPC codes by unit when aggregating to the brand level, and measure price using the ratio of revenue to equivalent unit sales, following standard practice (e.g., Nevo, 2001; Miller and Weinberg, 2017; Backus et al., 2021).⁹ We rank products within each category by their total revenue and include up to largest 20 in the estimation sample; typically this includes a private label product. All remaining products are collapsed into a single composite “fringe” product that we assume is priced by an independent firm.

Our baseline sample comprises 127 product categories. We obtain these categories by first identifying categories within the top 200 by revenue, and then applying a screen based on observed price dispersion to avoid categories with highly dissimilar products.¹⁰ Table 1 provides examples of product categories along with their average annual revenue and the number of products that appear in our estimation sample. The largest three categories are cigarettes, carbonated soft drinks, and refrigerated milk, which respectively generate \$5.4, \$5.3, and \$4.3

⁸We 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.

⁹In 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.

¹⁰Our ranking of categories is based on total revenue over 2006–2016. The screen is described in detail in Section 3.2.

Table 1: Product Categories in the Scanner Data

Rank	Name	Revenue	Brands	Rank	Name	Revenue	Brands
1	Cigarettes	5376	20	20	Ground/Whole Bean Coffee	1755	17
2	Soft Drinks-Carbonated	5276	19	30	Precut Fresh Salad Mix	1344	18
3	Milk-Refrigerated	4307	18	40	Frozen Poultry Entrees	1140	17
4	Fresh Bread	3327	19	60	Butter	803	16
5	Ready To Eat Cereal	3225	19	80	Liquid Creamers	636	13
6	Soft Drinks-Low Calorie	3062	19	100	Baby Accessory	544	18
7	Wine-Dom. Dry Table	2999	18	120	Snacks - Pretzel	473	16
8	Water-Bottled	2995	19	140	Fresh Tomatoes	403	15
9	Toilet Tissue	2880	15	160	Nutritional Products	350	13
10	Light Beer	2559	19	200	Frozen Desserts	275	15

Notes: This table shows a sample of product categories in the data sorted by revenue. *Revenue* is measured in average yearly sales in millions of nominal US \$ between 2006 and 2016. *Brands* measures the median of the number of brands within categories excluding fringe brands and private labels.

billion in average annual revenue. The smallest category we consider is frozen desserts, which generates \$275 million in average annual revenue.

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). We include food stores, drug stores, and mass merchandisers. Finally, we aggregate the week-level data up to the level of quarters, following Miller and Weinberg (2017), which reduces the overall size of the estimation samples and might help mitigate any bias due to forward-looking consumer behavior and the periodic sales that are observed in some consumer product categories (e.g., Hendel and Nevo, 2006a,b). 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.

To support estimation, we generate consumer-specific demographic draws by sampling 2,000 consumers from the Consumer Panel Data for each region and year.¹¹ 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 equation (3). We de-mean 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. For the micro-moments that appear in equation

¹¹By 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 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).

(9), we obtain the average values of the observed demographics for each product, region, and year, again applying the projection weights.

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. Since Compustat only covers publicly traded firms, this information is available for about half of the observations in our sample. With the estimation samples in hand, we deflate prices and incomes using the Consumer Price Index such that they are in real dollars as of the first quarter of 2010.¹²

3.2 Data Preparation

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 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, especially non-food 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

¹²We 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.

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.

As a proxy for within-category product heterogeneity, we use the within-category distribution of prices to identify and screen out categories for which our model may not be sufficiently flexible. Specifically, we remove categories in which the 99th percentile of prices is greater than five times the median price. Based on this criterion, we retain 133 of the top 200 product categories by revenue. We then remove categories that have an estimated median Lerner index above 2 in any year. This ex post screen filters out an additional 6 categories, leaving a baseline sample of 127 categories.

Our main results are similar if we use different procedures to select product categories. In the Appendix, we present results for two alternative screens without first filtering on price dispersion. First, we only impose the ex post screen to remove categories with a median Lerner index greater than 2 in any year. This retains 165 of the top 200 categories. Second, we consider a less conservative screen that removes categories only when the mean Lerner index exceeds 5 in any year. This second alternative retains 174 categories. Appendix Figure C.1 plots the markup trends over time under the different alternatives.

Also worthy of discussion are the compositional changes that occur in the Nielsen data as retail stores (and entire retail chains) enter and exit the sample. Such churn appears to be inconsequential over 2006–2016, but significant changes do occur over 2017–2019. Our specification insulates estimates from churn to some extent because we estimate the model by year, include chain \times region fixed effects in the demand and marginal cost equations, and allow the market sizes to scale with the number of stores operated by the retain chain. Still, compositional changes could affect markups if they change the price sensitivity of the typical consumer represented in the data.¹³ Thus, we are most confident in the markup trends we document over 2006–2016 though, for the most part, these trends simply continue through 2017–2019.

3.3 Case Study in Ready-to-Eat (RTE) Cereal

For our baseline analyses, we estimate the model for 1,778 distinct category-year combinations, with the 127 categories we select and the 14 years of Nielsen data. Before summarizing results, we focus specifically on the estimates for one product category, RTE Cereal, which has been previously studied (e.g., Nevo, 2001; Michel and Weiergraeber, 2018; Backus et al., 2021), in order to provide something of an illustrative example. Backus et al. (2021) provide details on the RTE cereal brands and the firms that supply the market over our sample period.

¹³Our specification of the demand model applies the same price parameters to every store in the same category, so our estimates could be interpreted as average effects.

Table 2 provides parameter estimates and standard errors (panel A), along with the number of observations, revenue-weighted median own-price elasticity of demand, and revenue-weighted median Lerner index (panel B). Most of the coefficients take the expected sign in most years. The price parameters in particular are negative and reasonably stable except for in the earliest and latest years. Consumers that have a higher household income and a child present in the household tend to be less price sensitive and more likely to purchase RTE cereal. The median price elasticities we obtain are reasonably close to the estimate of 2.67 provided in Backus et al. (2021). Aggregating across years, our estimates imply a median Lerner index of 0.50 and a median additive markup (i.e., $p - c$) of 0.10, which again are close to what is obtained in Backus et al. (2021).¹⁴ Thus, we conclude that our methodology can obtain reasonable results that are consistent with those of studies that make use of specific institutional details to a greater degree (and are therefore difficult to be implemented at scale).

It also is worth noting that our methodology allows markups to change over time for a number of reasons, including changes in product ownership, changes in products' qualities and marginal costs, and changes in consumer preferences as captured by the demand estimates. With RTE Cereals, the net effect of these mechanisms appears to be a Lerner index that declines slightly over the sample period, providing a clean illustration that our approach does not enforce rising markups. Our ultimate interest is in the evolution of markups across the many different categories in our estimation samples, and we turn to that exercise next.

3.4 Assessing Covariance Restrictions

To assess the impacts of using covariance restrictions to pin down mean price coefficients, we compare the estimated own-price elasticities to those obtained when using Hausman instruments (Hausman, 1996; Nevo, 2001). We construct the Hausman instruments as the average price for the brand in other markets in the same period. These instruments are valid if cost shocks are correlated across markets, e.g., due to shared manufacturing or distribution facilities, but demand shocks are uncorrelated across regions. Similar instruments have been widely used in the demand estimation literature since they can be constructed from standard market-level data sets. Hausman instruments may not be valid when demand is linked across regions (e.g., due to advertising) or if pricing strategies yield little cross-market variation (e.g., Berry and Haile, 2016). As a benchmark, we also consider the elasticities that would be obtained when assuming prices are exogenous.

We find that the covariance restrictions approach consistently yields more elastic demand, i.e., more negative price coefficients, compared to estimates that assume prices are exogenous. By comparison, Hausman instruments obtain a mix of substantial positive and substantial neg-

¹⁴The histograms presented in Figure 4 of the Backus et al. (2021) draft dated January 11, 2021 suggest a median additive markup per serving size of about 0.10. Our estimate of the additive markup is per ounce, but the average serving size is about an ounce, so the estimates appear to be similar.

ative shifts in the coefficients, depending on the year.

To illustrate this, Figure C.2 in the appendix plots the density of median own-price elasticities by category and year for a sample of 50 categories from 2006 through 2009. The green line shows the density of elasticities assuming that prices are exogenous, and the red line shows the elasticities using covariance restrictions. The distribution of elasticities is shifted downward using covariance restrictions, consistent with systematic price endogeneity arising from firms increasing prices in response to demand shocks.


By comparison, the density of elasticities using the Hausman instruments has more observations in both tails. When assuming that prices are exogenous, 23 percent of the estimates obtain median elasticities less than one in magnitude. Using Hausman instruments, this increases to 39 percent of the observations. When using covariance restrictions, only 4 percent of observations have inelastic demand.

Table 2: Estimation Results for RTE Cereals

Panel A: Point Estimates and Standard Errors														
	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Price	-18.430 (0.020)	-8.387 (0.008)	-11.962 (0.016)	-9.642 (0.008)	-10.272 (0.008)	-8.703 (0.005)	-10.420 (0.006)	-11.377 (0.006)	-12.234 (0.007)	-11.282 (0.013)	-13.622 (0.009)	-18.527 (0.023)	-13.680 (0.014)	-17.191 (0.036)
<i>Demographic Interactions</i>														
Income×Price	0.678 (0.001)	1.328 (0.001)	1.155 (0.001)	0.589 (0.001)	0.315 (0.001)	0.731 (0.001)	0.797 (0.001)	1.250 (0.001)	0.852 (0.001)	0.639 (0.001)	0.680 (0.001)	0.862 (0.001)	0.513 (0.001)	0.316 (0.001)
Income×Constant	0.128 (0.001)	0.209 (0.002)	0.395 (0.002)	0.215 (0.001)	0.293 (0.001)	-0.011 (0.000)	-0.073 (0.000)	-0.106 (0.000)	-0.050 (0.000)	-0.033 (0.001)	0.040 (0.000)	0.145 (0.001)	0.229 (0.001)	0.345 (0.003)
Children×Price	-0.439 (0.002)	-1.433 (0.002)	-0.731 (0.002)	1.140 (0.001)	1.652 (0.001)	2.833 (0.001)	3.321 (0.002)	2.389 (0.003)	2.326 (0.002)	2.399 (0.002)	2.932 (0.002)	2.610 (0.003)	2.416 (0.002)	2.196 (0.002)
Children×Constant	6.533 (0.024)	4.636 (0.015)	5.536 (0.017)	2.208 (0.007)	3.563 (0.007)	0.825 (0.004)	0.566 (0.000)	0.680 (0.005)	0.527 (0.002)	0.785 (0.008)	2.535 (0.004)	2.985 (0.012)	4.941 (0.013)	5.632 (0.044)
<i>Random Coefficient</i>														
N(0,1)×Constant	5.156 (0.022)	3.752 (0.015)	5.002 (0.018)	2.265 (0.011)	4.436 (0.009)	0.555 (0.014)	0.004 (0.024)	0.238 (0.038)	0.240 (0.022)	1.370 (0.020)	5.227 (0.007)	6.426 (0.024)	9.596 (0.026)	11.095 (0.083)
Panel B: Other Statistics													2018	2019
Observations	15,441	16,336	16,604	16,791	17,241	17,329	16,444	16,213	16,443	15,829	15,487	12,654	18,850	17,805
Median Own Elasticity	3.413	1.587	2.368	1.928	1.964	1.659	2.094	2.262	2.396	2.130	2.502	3.357	2.373	3.016
Median Lerner	0.338	0.727	0.492	0.587	0.596	0.657	0.515	0.473	0.446	0.515	0.467	0.350	0.492	0.390

Notes: This table summarizes the results of estimation for the RTE Cereals category for each year in the sample. Panel A provides the parameters and the standard errors, which are clustered at the region level. Panel B provides the number of product-chain-region-quarter observations, the revenue-weighted median own price elasticity of demand, and the revenue-weighted median Lerner index.

Figure 1: Markups Over Time Across Product Categories



Notes: This figure plots the mean of within-category median markups over time. Markups are defined by the Lerner index, $(p - c)/p$ and are estimated separately by product category and year.

4 The Evolution of Markups in Consumer Products

In this section, we document the evolution of markups across a large set of 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 between variation within and across products and to decompose the evolution of markups into changes in prices and marginal costs.


4.1 Aggregate Trends

Figure 1 reports the development of markups over time. We compute the median within each product category and year and then take the mean across all categories. The figure indicates a substantial increase in markups over time. The median Lerner index has increased from 0.39 in 2006 to over 0.50 towards the end of our sample period. The average growth rate in the category-level median markup is 1.8 percent per year.

Next, we analyze how the distribution of markups 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 2. 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.

Figure 2: 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.

4.2 Decomposition of Effects

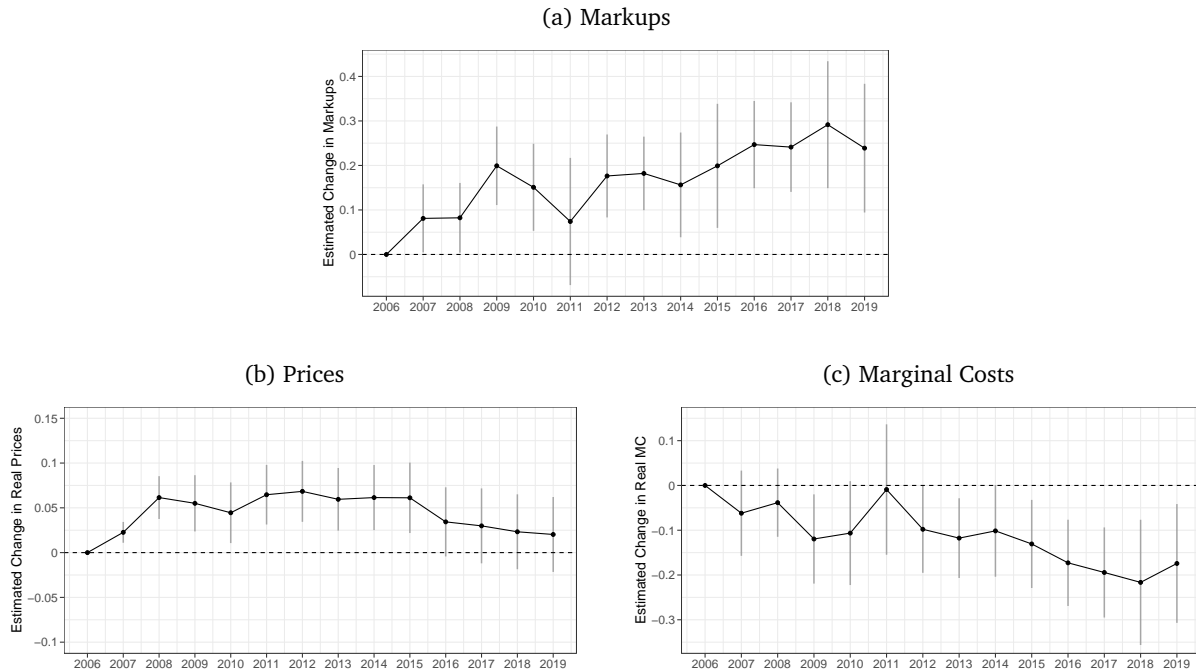
Aggregate trends in markups may be explained by firms charging higher markups or by consumers shifting to higher-markup products. To evaluate these channels, we analyze the change in markups at the product level, where our unit of observation is a unique product-chain-DMA-quarter-year combination. We regress the log of the Lerner index on quarter, year, and product-chain-DMA fixed effects, using revenues as weights.¹⁵

The results of this regression are documented in panel (a) of Figure 3. The figure displays point estimates and 95 percent confidence intervals for the year-specific coefficients. The estimates indicate an increase in product-level markups of more than 25 percent between 2006 and 2019. The estimated annual growth rate in product-level markups is 1.7 percent per year. Due to the inclusion of product-chain-DMA fixed effects, year dummies only capture 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, in our sample, aggregate markups 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). Figure C.3 in the Appendix shows a plot of the corresponding yearly coefficients.

Table C.1 in the Appendix documents full results of the regression that corresponds to panel (a) of Figure 3 alongside alternative specifications in which we replace year dummies with a linear time trend and consider dropping product-chain-DMA fixed effects or replacing them with category fixed effects. We obtain qualitatively similar results across these specifications, in which we estimate an average yearly increase in average markups between 1.2 and 1.7 percent. We estimate larger changes when controlling for product-level fixed effects, indicating

¹⁵Revenue weights capture an average effect across a quantity-weighted bundle of products when measuring changes to log markups.

Figure 3: 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.

that the within-product changes in markups are greater than the aggregate (revenue-weighted) changes in markups. Though these differences are not significant, this suggests that some of the product-level increase in markups may be offset by the introduction of lower-markup products over time.¹⁶


Using our detailed data on prices and our demand estimates, we are able to decompose the increase in markups into changes in prices and marginal costs. For this purpose, we regress log prices and log marginal costs on product-DMA-retailer fixed effects and year dummies. 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 3. Panel (b) shows that real prices increased at the beginning of our sample period, but declined in later years. From 2006 to 2012, the average real price for products in our sample increased by 7 percent, but by 2019, real prices were only 2 percent higher than in 2006.

Panel (c) of the figure reports the yearly coefficients for log marginal costs. We estimate that marginal costs decline by 1.3 percent per year on average. In 2017–2019, marginal costs are roughly 20 log points lower than in 2006. Thus, though higher real prices account for part

¹⁶Table C.2 in the Appendix shows results using non-weighted regressions. The results are similar, with slightly lower estimates toward the end of the sample. These differences are not statistically significant.

Figure 4: Changes in Elasticity, Price Sensitivity, and Quality



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

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 fully passed on to consumers.¹⁷

Why might lower costs not lead to lower prices? One potential explanation is that demand is less responsive to prices. To investigate this potential explanation, we evaluate changes in consumer preferences over time. First, we 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 4. 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 production costs, quality, and competition. To shed additional light on the changes in markups we observe, we isolate two features of consumer preferences and examine trends in these parameters. First, we look at consumer price sensitivity, which

¹⁷Figure C.4 in the Appendix documents results using nominal, i.e., non-deflated, prices and marginal costs and shows that nominal marginal costs are relatively constant over time.

reflects the mean price coefficient within a category year. Specifically, we define the price sensitivity as the log absolute value of the mean price coefficient. Thus, a higher value of the price sensitivity parameter indicates that consumers are more sensitive to prices, independent of changes to underlying demographics or other changes in preferences. Second, we examine changes in perceived product quality. We measure quality as the mean value consumers place on products within a category, and we standardize the measure separately within each category.¹⁸

We repeat the product-level regressions of price sensitivity and quality using the fixed effects as above, and we display the estimated coefficients in panels (b) and (c) of Figure 4. The time trend shows that consumer price sensitivity is decreasing over time, consistent with the declining magnitudes of elasticities. 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.

On the other hand, we find that quality declined over our sample, as shown in panel (c). In our model, lower quality yields lower shares and more elastic demand. Thus, we find that quality moves in the opposite direction from the estimated trends in elasticities and markups. It is important to note that quality in our model is defined relative to the outside option, i.e., outside of the retailer-DMA-category. The declines in quality we observe may be due to increases in perceived utility from substitute retail channels.

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, despite reductions in perceived quality.

4.3 Panel Data Analysis

Markups are determined in equilibrium by consumer preferences, production costs, and competition. To evaluate the role of these demand and supply channels, we perform a more detailed analysis that exploits the panel structure of our estimates across products and over time. We evaluate how simultaneous changes in preferences, production costs, and competition correlate with changes in markups. We then examine potential mechanisms that drive these changes. Specifically, we regress log markups observed for each product-chain-DMA-quarter-year on our estimated preference parameters, marginal costs, and concentration measures. For preferences, we examine price sensitivity, defined (as above) as the log absolute value of the mean price coefficient. We use marginal costs to capture changes in production technology. We also relate markups to perceived product quality derived from our demand model, following equation (4).

¹⁸Specifically, we first obtain the product-specific mean utility, and then divide by the absolute value of the mean price coefficient to obtain a value in dollar terms, following equation (4). We then standardize this measure using the mean and standard deviation by category across all years.

Table 3: Dependent Variable: Log Markup

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price Sensitivity	-0.801*** (0.023)			-0.555*** (0.037)	-0.559*** (0.036)	-0.563*** (0.036)	-0.555*** (0.037)	-0.563*** (0.036)
Marginal Cost (Standardized)		-0.687*** (0.041)		-0.429*** (0.037)	-0.424*** (0.037)	-0.422*** (0.037)	-0.429*** (0.037)	-0.422*** (0.037)
Quality (Standardized)			-0.266*** (0.032)	0.003 (0.008)	0.001 (0.007)	0.001 (0.007)	0.003 (0.008)	0.001 (0.007)
Income (Log)				0.058*** (0.016)	0.074*** (0.015)	0.075*** (0.014)	0.058*** (0.015)	0.074*** (0.014)
Children at Home				-0.058 (0.042)	-0.051 (0.042)	-0.046 (0.043)	-0.053 (0.040)	-0.045 (0.041)
Brand HHI					0.436*** (0.046)			-0.033 (0.100)
Parent HHI						0.524*** (0.061)		0.550*** (0.118)
Retailer HHI							0.038 (0.032)	0.003 (0.033)
Brand-Category-DMA-Retailer FEs	X	X	X	X	X	X	X	X
Time Period FEs				X	X	X	X	X
Observations	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053
R^2 (Within)	0.677	0.625	0.104	0.854	0.856	0.857	0.854	0.857

Standard errors in parentheses

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

We standardize marginal costs and quality separately by product category.¹⁹ For concentration, we use the average HHI across DMAs in each year, and we calculate HHI by brand, parent company, and retailer. Parent HHI takes into account multi-brand ownership, whereas brand HHI does not. We measure HHI on a 0 to 1 scale.

Results are displayed in Table 3. Each regression includes product-market (i.e., brand \times category \times retailer \times DMA) fixed effects, and the coefficients reflect within-product changes over time. To account for changes in demographics, we include demographic controls for income and the presence of young children in columns (4)-(8).²⁰ Standard errors are clustered at the product category level. Column (1) indicates that changes in consumer preferences, measured by price sensitivity, are highly predictive of changes in markups. The price coefficient explains 68 percent of the within-product changes in markups over time and is economically significant. A 10 percent decrease in price sensitivity is associated with approximately a 8 percent increase in markups.

Column (2) indicates that lower marginal costs are strongly associated with increased markups. Marginal costs alone can explain 63 percent of the within-product changes in markups. On the other hand, column (3) suggests that increases in markups are associated with lower perceived product quality. However, the R^2 in column (3) is only 0.1, indicating that changes in quality explain less of the within-product variation in markups than price sensitivity or marginal costs.

¹⁹We use standardized marginal costs instead of log marginal costs to include the small fraction of observations where we estimate marginal costs to be negative.

²⁰Demographic variables add little explanatory power to the regression.

In column (4), we combine all three measures and add demographic controls for log income and the presence of young children. We also include year-quarter fixed effects, so that the coefficients now capture deviation from the time series. 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 quality are correlated with increasing markups in the time series, products with greater increases in quality do not realize differential changes in markups. We find that increases in income are 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 , which now also accounts for time period fixed effects, is 0.85.

We add our measures of concentration to specifications (5)-(8). Columns (5)-(7) demonstrate that brand and parent-level concentration are positively correlated with markups, after controlling for price sensitivity, marginal costs and quality. Retailer concentration is not correlated with higher markups. Upstream concentration at the parent level has a stronger relationship than upstream concentration at the brand level. When including all three measures of concentration, only parent HHI is positive and significant, as shown in column (8). The coefficient on brand HHI actually becomes negative.

These regressions indicate that changes in markups are most strongly associated with changes in consumer preferences and lower marginal costs. Changes in concentration are also associated with changes in markups, though our results suggest that the primary channel is through changes in ownership of brands, rather than the increasing dominance of individual brands or retailers. The coefficient of 0.550 in column (8) 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 1.1 percent increase in markups. The relationships between markups and changes in concentration at the brand or retailer level are much weaker. We obtain similar results if we instead run regressions at the product category level. We report these results in Table C.3 in the Appendix.

We next examine the potential drivers of consumer price sensitivity, marginal costs, and quality over time. 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.

To explore these relationships, we aggregate our data to the brand-category-year level. We first regress own-price elasticities on the above measures. Each regression uses brand-category-parent fixed effects, in addition to year dummies, so the coefficients reflect time-series variation within each product that departs from the aggregate trend. We cluster standard errors at the category level.

Table 4 reports the results. Column (1) shows that greater marketing expenses are correlated with less elastic demand. Column (2) indicates that greater R&D expenditures are also

Table 4: Potential Mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Log Abs. Elasticity	Log Abs. Elasticity	Log Abs. Elasticity	Log Abs. Elasticity	Price Sensitivity	Marginal Cost	Quality
Log Marketing Spend	-0.071** (0.030)			-0.045 (0.034)	-0.106*** (0.037)	0.025 (0.040)	0.049* (0.029)
Log R&D		-0.109*** (0.036)		-0.086** (0.039)	-0.068* (0.040)	-0.043 (0.040)	0.004 (0.033)
Log Num. UPCs			-0.010 (0.020)	-0.019 (0.025)	-0.027 (0.026)	0.025 (0.029)	0.169*** (0.028)
Brand-Category FEs	X	X	X	X	X	X	X
Time Period FEs	X	X	X	X	X	X	X
Observations	13,524	13,448	33,599	11,257	11,259	11,259	11,259
R ²	0.697	0.686	0.667	0.704	0.919	0.758	0.572

Standard errors in parentheses

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

correlated with less elastic demand. Column (3) shows that product variety also has a negative correlation with the (absolute) elasticity, though it is smaller and not statistically significant. Column (4) includes all three explanatory variables in one regression. Controlling for the other factors, the magnitudes for marketing and R&D shrink modestly toward zero. R&D remains statistically significant ($p < 0.05$).

In columns (5) through (7), we investigate each of these firm-level variables on the channels that can drive lower markups. We repeat the above regression while replacing the dependent variable with price sensitivity, marginal costs, and quality. Column (5) shows that increased marketing expenditure is associated with a reduction in price sensitivity. The coefficient is over 50 percent larger than the coefficient in column (1) relating marketing to the price elasticity.

We do not find any significant relationship between these explanatory factors and marginal costs, as shown in column (6). We do find directional relationships that seem plausible, as the coefficients on marketing and product variety are positive.

Column (7) indicates that product variety is positively and significantly correlated with perceived quality. We view this as an encouraging sanity check of our empirical model, which estimates everything at the brand level. Thus, consistent with Brand (2021), we find that increases in product variety are associated with increases in perceived quality. However, we do not find that changes in quality 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).

Overall, these results suggest that firm expenditures on marketing and R&D may lead to increased markups over time, potentially through a reduction in consumer price sensitivity.

5 Markups, Welfare, and Consumer Surplus

In this section, we analyze how consumer surplus and total welfare in the market for consumer products have changed over time. We also calculate consumer and total welfare for various

Table 5: Consumer Surplus, Producer Surplus, and Welfare Per Capita

Specification	Year	CS	PS	Welfare	% Change CS
Baseline	2006	629	226	855	0.0
Price = MC	2006	863	0	863	37.2
Prices Scaled to 2019 Price Levels	2006	607	227	835	-3.4
Markups Scaled to 2019 Markup Levels	2006	550	212	761	-12.6
Baseline	2019	732	292	1024	0.0
Price = MC	2019	1102	0	1102	50.6
Prices Scaled to 2006 Price Levels	2019	763	269	1032	4.3
Markups Scaled to 2006 Markup Levels	2019	829	224	1054	13.4

Notes: This table reports consumer surplus (CS), producer surplus (PS), and welfare 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 scenarios, to estimate the deadweight loss from (changes in) market power.

The expected value of consumer surplus (CS) in our model is (Small and Rosen, 1981):

$$CS = -\frac{1}{N} \sum_i \frac{1}{\alpha_i} \ln \left(\sum_j \exp(v_{ij}) \right) \quad (10)$$


where $v_{ij} = x'_{ij}\beta_i^* + \alpha_i^* p_{jcr} + \xi_{jr} + \xi_{cr} + \xi_t + \Delta\xi_{jcr}$ and N denotes the number of consumers.²¹ Consumer surplus is the additional surplus relative to outside options and can be interpreted as the added value from the presence of the focal products. Producer surplus (PS) reflects profits and is measured as price less marginal costs multiplied with quantities: $PS = (p - c)q$; welfare (W) is the sum of PS and CS .

Table 5 shows the values of consumer and producer surplus per capita for the first and the last year of our sample using observed prices (“Baseline”) and prices under different counterfactual scenarios. To compute counterfactual values, we hold fixed estimated preference parameters and marginal costs, and we simulate choices using different prices. First, we consider the counterfactual where prices equal marginal costs. Second, we scale all prices by the average realized price change for all products in the same module from one year to another (e.g., from 2006 to 2019). Third, we scale all markups by the average realized markup change for all products in that module 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. The last column shows changes in consumer surplus relative to the baseline scenario of observed prices, estimated markups and expected utility in each year.

The results indicate that consumer surplus per capita has increased by about 16 percent between 2006 and 2019, from \$629 to \$732 per capita. Since average prices have not declined and perceived quality has not increased, the increase in consumer surplus is likely due to lower

²¹In 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.

Figure 5: 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 module fixed effects, separately for different quartiles of the income distribution.

price sensitivity, i.e., the fact that consumers receive lower disutility from any given price in 2019. Along with higher markups, producer surplus has increased over the period, from \$226 to \$292 per capita. Overall, we estimate that total welfare has increased by approximately 20 percent. Roughly sixty percent of the welfare gains have accrued to consumers. In 2006, roughly three fourths of surplus went to consumers, indicating that the gains have shifted toward firms over time.

Markups are costly for consumers. If firms would set prices equal to marginal costs, consumer surplus would be substantially higher in both 2006 and 2019, as shown by the second specification for each year. Our estimates suggest that markups in 2006 produced a modest deadweight loss, reducing aggregate welfare from \$863 to \$855 per capita (1 percent). In 2019, markups resulted in a much greater deadweight loss, reducing welfare by more than 7 percent.

Despite the overall increase in consumer surplus between 2006 and 2019, consumers suffered from rising markups. Based on 2019 preferences, consumer surplus in 2019 would be 13 percent higher if markups were scaled down to match average 2006 levels. If prices were scaled down to match average 2006 levels, consumer surplus in 2019 would be more than 4 percent higher than in the baseline scenario. Counterfactual markup increases are more costly than price increases because marginal costs have been declining over time.

Next, we analyze how the change in consumer surplus over time varies by income class. For this purpose, we calculate the log of consumer surplus per purchasing decision separately by each quartile of the income distribution and for each module-year. We relate these values to module and year fixed effects and document the coefficients across years in Figure 5. The results indicate that the overall increase in consumer surplus per capita between 2006 and 2019 is mainly driven by consumers with relatively high income and takes place during the second half of our sample period. In contrast, the lowest quartile of the income distribution suffers

from significant losses in consumer welfare until 2014 and reaches a level similar to the initial value in 2006 towards the end of the sample period. In Figure C.5 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 gain substantially over time, while losses in consumer welfare are limited to lower income groups. These findings suggest that changes in market power and consumer preferences over time have important distributional consequences.

6 Conclusion

This paper analyzes the evolution of market power in consumer products in the US 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 more than 25 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 25 percent less price sensitive over the sample period. 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.

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.
- 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. HAILE (2016): “Identification in differentiated products markets,” *Annual Review of Economics*, 8, 27–52.
- 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*.
- BOND, S., A. HASHEMI, G. KAPLAN, AND P. ZOCH (2020): “Some Unpleasant Markup Arithmetic: Production Function Elasticities and Their Estimation from Production Data,” Tech. rep., National Bureau of Economic Research.
- 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*.
- 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.
- 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,” Tech. rep., National Bureau of Economic Research.

- 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.
- 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.
- 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.
- GRIECO, P. L. E., C. MURRY, AND A. YURUKOGLU (2021): “The Evolution of Market Power in the US Automobile Industry,” Working Paper.
- 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.
- 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 (2006a): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74, 1637–1673.
- (2006b): “Sales and Inventory Behavior,” *RAND Journal of Economics*, 37, 543–561.

- (2013): “Intertemporal Price Discrimination in Storable Goods Markets,” *American Economic Review*, 103, 2722–2751.
- LERNER, A. P. (1934): “The Concept of Monopoly and the Measurement of Monopoly Power,” *Review of Economic Studies*, 1.
- MACKAY, A. AND N. H. MILLER (2019): “Estimating Models of Supply and Demand: Instruments and Covariance Restrictions,” Available at SSRN: <https://ssrn.com/abstract=3025845>.
- MICHEL, C. AND S. WEIERGRAEBER (2018): “Estimating Industry Conduct in Differentiated Products Markets,” *Working Paper*.
- MILLER, N. H. AND M. C. WEINBERG (2017): “Understanding the Price Effects of the Miller-Coors Joint Venture,” *Econometrica*, 85, 1763–1791.
- MILLER, N. H., M. C. WEINBERG, AND G. SHEU (2021): “Oligopolistic Price Leadership and Mergers: The United States Beer Industry,” *American Economic Review*, forthcoming.
- NEVO, A. (2000): “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.
- RAVAL, D. (2020): “Testing the Production Approach to Markup Estimation,” *Working Paper*.
- ROSSE, J. N. (1970): “Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology,” *Econometrica*, 38, 256.
- 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 Details on Estimation

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 $\theta = (\alpha, \Pi, \Sigma)$. We first estimate $\theta_2 = (\Pi, \Sigma)$ 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, 2019). 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 (e.g. Berry et al., 1995; Nevo, 2001; Miller and Weinberg, 2017; Backus et al., 2021). However, our approach has substantial computational benefits due to a complication that arises with covariance restrictions, as we explain below. In the appendix, we focus on the implementation of the estimation steps, and then discuss computational burden and the resources employed.

A.1 First Step

In the first estimation step, we use the micro-moments to pin down the “nonlinear” parameters, i.e., $\theta_2 = (\Pi, \Sigma)$. 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.

To determine the weighting matrix, W , that appears in the GMM objective function (see equation (7)), we start with an initial candidate parameter vector of $\theta_2^{(0)} = \vec{0}$, such that the model collapses to a logit. In this special case, the own-price elasticity of demand for product j is given by:

$$\epsilon_j = \alpha p_j (1 - s_j) \tag{A.1}$$

We fix the price parameter at the level, $\alpha = \tilde{\alpha}$, such that equation (A.1) holds, on average across the observations in the estimation sample, with unit demand elasticity ($\epsilon_j = -1$). For any candidate θ_2 , there is a unique vector of the mean product valuations that align the predicted and observed shares. In the special case of $\theta_2^{(0)}$ the mean valuations have a closed-form solution:

$$\delta_{jcr}(\theta_2^{(0)}) \equiv \log(s_{jcr}) - \log(s_{0cr}) \tag{A.2}$$

With $\delta_{jcr}(\theta_2^{(0)})$ and $\tilde{\alpha}$, we calculate the micro-moments, $g^{MM}(\theta_2^{(0)})$, from equation (9).

We construct the weighting matrix as follows. We set the first element—the weight placed on the covariance restrictions—at 1e-8. The remaining diagonal elements weight the micro-moments according to their contribution, such that the total sums to 100. Specifically, elements

$k = 2, \dots, K$ along the diagonal element are given by

$$w_k = 100 \frac{\left(g_{k-1}^{MM}(\theta_2^{(0)})\right)^2}{\sum_n \left(g_n^{MM}(\theta_2^{(0)})\right)^2}$$

We set all non-diagonal elements in the weighting matrix to zero.

We then proceed to estimate θ_2 based on equation (7), with nearly all of the weight being placed on the micro-moments, and with the price parameter still fixed at $\tilde{\alpha}$. For each candidate θ_2 , we recover the mean valuations, $\delta_{jcr}(\theta_2)$, using the contraction mapping of Berry et al. (1995) with a numerical tolerance of 1e-9. We then calculate the micro-moments with $\delta_{jcr}(\theta_2)$ and $\tilde{\alpha}$. It is still necessary to evaluate the covariance of unobserved shocks, even though the covariance restriction receives almost zero weight in the first-step objective function. To do so, we recover $\Delta\xi_{jcr}(\theta_2)$ as the residual from the OLS regression of $(\delta_{jcr}(\theta_2) - \tilde{\alpha}p_{jcr})$ on the fixed effects. We also obtain marginal costs from using equation (5), looping over the chain-region-quarter combinations, and then recover $\Delta\eta_{jcr}(\theta_2)$ 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 θ_2 , and repeat to convergence.

A.2 Second Step

In the second step, we hold fixed the estimated nonlinear parameters, scale up the covariance restriction contribution, and choose the price parameter that minimizes the objective. In other words, we estimate α taking as given the estimates of θ_2 obtained in the first step. In doing so, we change the first diagonal element of the weighting matrix from 1e-8 to 100, so that the covariance restriction receives the same weight as all of the micro-moments combined. This effectively estimates the α based on the covariance restriction alone because micro-moments do not identify the mean price parameter (Berry and Haile, 2020).²² We evaluate moments using the same procedure outlined for the first step estimation, and iterate to convergence constraining 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 are typically *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, 2019). Care must then be taken to ensure that the estimator converges to the smaller value. Figure C.6 illustrates this in the context of RTE cereals. Each panel traces out the contribution of the covariance restriction to the objective function for different values of α . In 2006 and 2018, 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

²²We confirm in *ex post* checks that the value of α matters little for the micro-moments.

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.²³

We proceed by selecting starting values of $\alpha^{(0)} = \phi \tilde{\alpha}$ where $\tilde{\alpha}$ is as defined for the first step of estimation, 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. However, it is computationally-intensive and also can break down for unreasonable candidate θ_2 parameters (which are inevitably considered in estimation). Thus, our two-step approach to estimation conveys both speed and numerical stability, both of which are important given the scale of the empirical exercise.

A.3 Computation Notes

Our code builds on the BLPEstimatorR package for R (Brunner et al., 2020).²⁴ In early experiments, we replicated our results for some categories using the PyBLP package for Python (Conlon and Gortmaker, 2020).²⁵ We ultimately selected the R package because the micro-moments implementation was substantially faster. 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 12GB 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.

B Data Details

B.1 Market Size

Recall from Section 2.2 that the quantity demanded in our model is given by $q_{jct}(p_{rct}; \theta) = s_{jct}(p_{rct}; \theta)M_{rct}$, where $s(\cdot)$ is the market share, p_{rct} is a vector of prices, and M_{rct} 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

²³The 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.

²⁴<https://github.com/cran/BLPEstimatorR>, last accessed March 26, 2021

²⁵<https://github.com/jeffgortmaker/pyblp>, last accessed March 26, 2021.

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 for each product category:

1. Obtain a “base” value by multiplying the population (at the region-year level) with the number of stores (at the retail 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: $Q_{crqy} = \sum_j q_{jcrqy}$.
3. Calculate $\gamma_{cr} = mean_{q,y} \left(\frac{Q_{crqy}}{BASE_{crqy}} \right)$ 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 quantity-to-base ratios that are less than 5 percent of the mean ratio, which helps avoid extraordinary small inside good market shares.
4. We obtain the market size by scaling the base value according to

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

which generates markets sizes for each retail chain, region, quarter, and year, such that the combined share of the inside goods is around 0.45, on average.

5. For a small minority of cases (<5 percent of categories), this procedure generates a combined share of the inside goods that exceeds 0.90 in some period, which is high enough that we encounter numerical problems in estimation. For any category in which this occurs, we repeat the steps above using the alternative conversion factor $\gamma_{cr} = max_{q,y} \left(\frac{Q_{crqy}}{BASE_{crqy}} \right)$, which leads to market sizes that are more workable in practice.


B.2 Other Details

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, in an attempt 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 implement this screen before culling 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, we note that we define food categories as belonging to the product departments “Dry Grocery,” “Frozen Foods,” “Dairy,” “Deli,” “Packaged Meat,” “Fresh Produce,” and “Alcoholic Beverages.” Non-food categories belong to the product departments “Health and Beauty Care,” “Non-food Grocery,” and “General Merchandise.”


C Appendix Figures and Tables

Figure C.1: Markups Over Time: Alternative Screens




Notes: This figure plots the mean of within-category median markups over time. Markups are defined by the Lerner index, $(p - c)/p$ and are estimated separately by product category and year. The solid line represents our baseline set of 127 categories. The dashed line represents all 165 categories that never exceeded a median markup of 2 in any year. The dash-dotted line represents 174 categories that never exceeded a mean markup of 5 in any year.

Figure C.2: Implied Elasticities Under Alternative Identification Restrictions




Notes: This figure plots the density of the median own-price elasticity by category and year. The sample of estimates includes 50 categories from 2006 through 2010. The red line shows the density of implied elasticities using covariance restrictions. The green line shows the density of implied elasticities assuming exogenous prices. The blue line shows the density of implied elasticities using Hausman instruments. The vertical dashed line indicates an elasticity of -1 .

Figure C.3: 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).

Figure C.4: Product-Level Changes in Nominal Prices and Marginal Costs



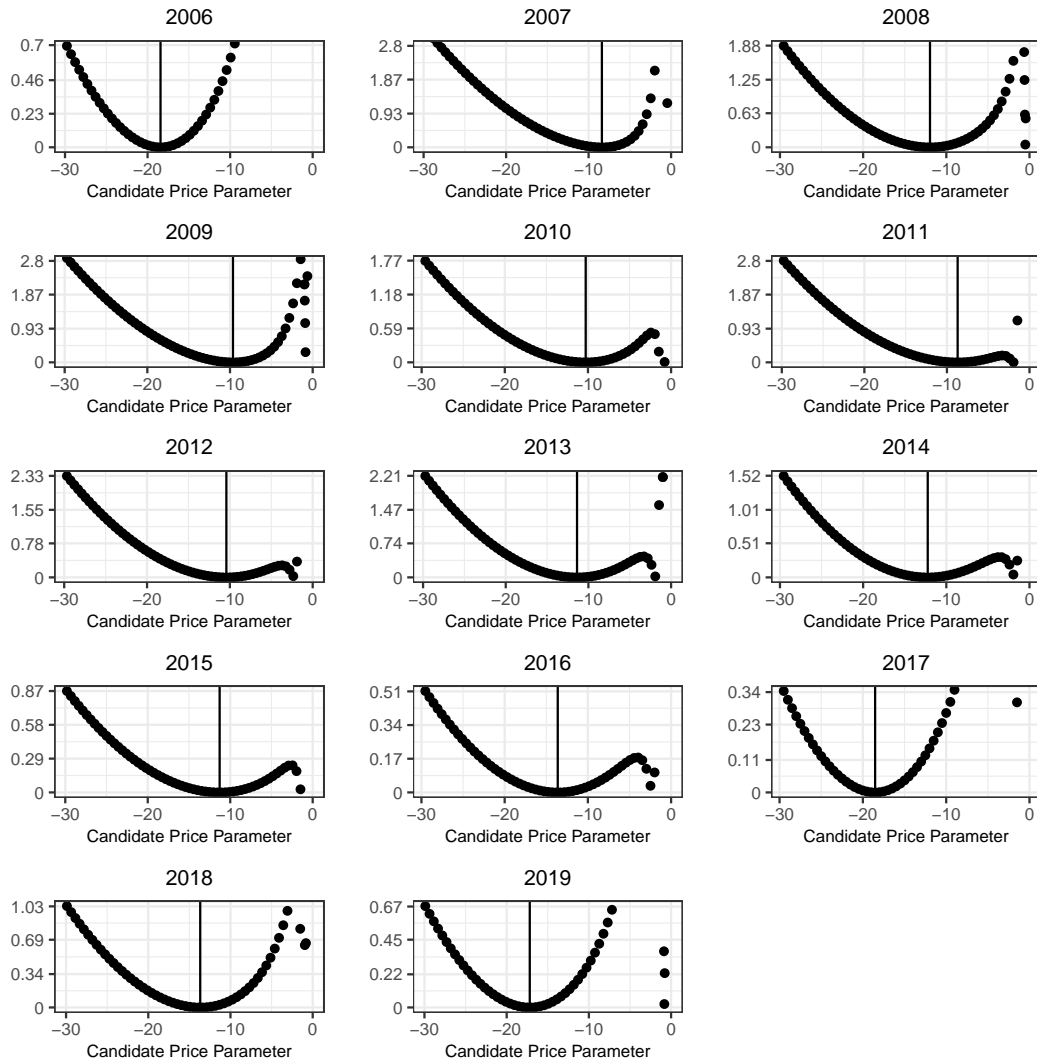
Notes: This figure shows coefficients and 95 percent confidence intervals of regressions of the log of nominal prices and marginal costs at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.

Figure C.5: 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 module fixed effects, separately for different deciles of the income distribution.

Figure C.6: Contribution of Covariance Restriction to Objective Function With RTE 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 C.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	0.012*** (0.004)		0.012*** (0.004)		0.017*** (0.004)	
Year=2007		0.078** (0.038)		0.080** (0.038)		0.081** (0.039)
Year=2008		0.081** (0.039)		0.083** (0.039)		0.082** (0.040)
Year=2009		0.204*** (0.044)		0.205*** (0.044)		0.199*** (0.045)
Year=2010		0.155*** (0.049)		0.155*** (0.049)		0.151*** (0.049)
Year=2011		0.072 (0.072)		0.074 (0.071)		0.074 (0.072)
Year=2012		0.173*** (0.047)		0.172*** (0.047)		0.176*** (0.047)
Year=2013		0.181*** (0.043)		0.177*** (0.042)		0.182*** (0.042)
Year=2014		0.148** (0.061)		0.145** (0.060)		0.156*** (0.059)
Year=2015		0.175** (0.072)		0.173** (0.070)		0.199*** (0.070)
Year=2016		0.211*** (0.050)		0.208*** (0.049)		0.247*** (0.050)
Year=2017		0.202*** (0.050)		0.194*** (0.049)		0.241*** (0.051)
Year=2018		0.241*** (0.073)		0.248*** (0.072)		0.292*** (0.072)
Year=2019		0.181** (0.072)		0.189*** (0.071)		0.239*** (0.073)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053
R ²	0.005	0.008	0.335	0.338	0.693	0.696

Notes: Dependent variable is the log of the Lerner index.

Standard errors are clustered at the product category level and reported in parentheses.

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

Table C.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	0.010** (0.004)		0.013*** (0.004)		0.015*** (0.004)	
Year=2007		0.090** (0.040)		0.089** (0.040)		0.097** (0.040)
Year=2008		0.089** (0.043)		0.087** (0.043)		0.103** (0.043)
Year=2009		0.203*** (0.045)		0.202*** (0.045)		0.222*** (0.046)
Year=2010		0.158*** (0.047)		0.158*** (0.047)		0.180*** (0.047)
Year=2011		0.069 (0.051)		0.074 (0.051)		0.098* (0.051)
Year=2012		0.173*** (0.048)		0.181*** (0.048)		0.209*** (0.048)
Year=2013		0.187*** (0.048)		0.195*** (0.046)		0.220*** (0.046)
Year=2014		0.158*** (0.053)		0.169*** (0.051)		0.196*** (0.051)
Year=2015		0.239*** (0.061)		0.258*** (0.059)		0.289*** (0.059)
Year=2016		0.225*** (0.051)		0.244*** (0.050)		0.276*** (0.051)
Year=2017		0.165*** (0.052)		0.180*** (0.051)		0.214*** (0.053)
Year=2018		0.192*** (0.068)		0.221*** (0.069)		0.252*** (0.070)
Year=2019		0.163*** (0.060)		0.193*** (0.060)		0.224*** (0.061)
Quarter FEs	X	X	X	X	X	X
Category, Retailer & DMA FEs			X	X		
Brand-Category-DMA-Retailer FEs					X	X
Observations	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053	13,436,053
R ²	0.003	0.007	0.332	0.335	0.683	0.687

Notes: Dependent variable is the log of the Lerner index.

Standard errors are clustered at the product category level and reported in parentheses.

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

Table C.3: Dependent Variable: Log Category Markup

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price Sensitivity	-0.702*** (0.021)			-0.499*** (0.031)	-0.504*** (0.029)	-0.506*** (0.029)	-0.500*** (0.030)	-0.507*** (0.029)
Marginal Cost (Standardized)		-0.329*** (0.013)		-0.157*** (0.014)	-0.156*** (0.014)	-0.152*** (0.014)	-0.157*** (0.015)	-0.153*** (0.015)
Quality (Standardized)			-0.297*** (0.014)	0.016 (0.010)	0.016* (0.010)	0.014 (0.010)	0.016* (0.010)	0.014 (0.010)
Brand HHI					0.388*** (0.116)			0.033 (0.127)
Parent HHI						0.662*** (0.174)		0.634*** (0.160)
Retailer HHI							0.135 (0.206)	0.112 (0.197)
Category FEs	X	X	X	X	X	X	X	X
Year FEs				X	X	X	X	X
Demographic Controls			X	X	X	X	X	X
Observations	1,778	1,778	1,778	1,778	1,778	1,778	1,778	1,778
R ² (Within)	0.829	0.714	0.581	0.899	0.901	0.903	0.899	0.903

Standard errors in parentheses

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