Appendix

A Additional Details About the Data

A.1 Data Construction

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

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

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

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

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

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

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

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

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

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

Figure A.1: Private Label Shares Over Time

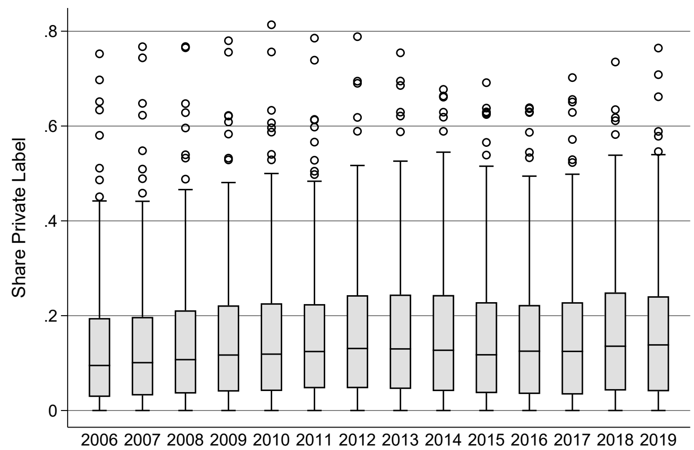

Notes: Figure provides the distribution of private label shares across our 133 baseline categories in each year.

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

A.2 Market Size Calculations

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

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

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

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

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

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

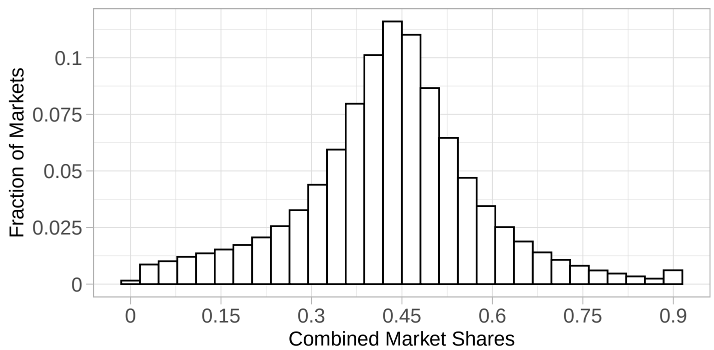

Notes: This figure shows the distribution of market shares of inside goods. Observations are at the chain-region-year-quarter level and reflect the sum of the market shares of all inside goods in a market at a given point in time.

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

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

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

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

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

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

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

B Estimation Details

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

B.1 First Step

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

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

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

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

B.2 Second Step

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

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

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

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

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

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

B.3 Computation Notes

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

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

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

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

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

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

C Validation for Selection of Product Categories

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

C.1 Marginal Cost Estimates

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

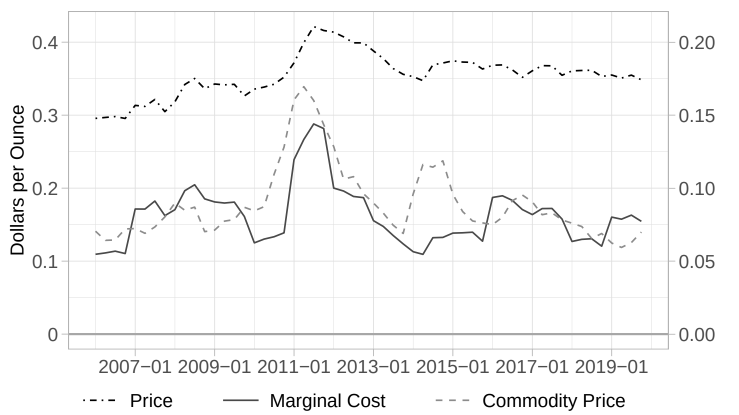

Notes: This figure plots the time series of quantity-weighted prices and marginal costs (solid line) for ground/whole bean coffee. Prices are observed and marginal costs are recovered from the profit-maximization conditions. Also shown is the commodity price index for coffee (dashed gray line), which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

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

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

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

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

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

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

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

C.2 Elasticity Estimates in the Literature

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

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

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

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

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

D Alternative Model Specifications

D.1 Random Coefficients Logit versus Logit Demand

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

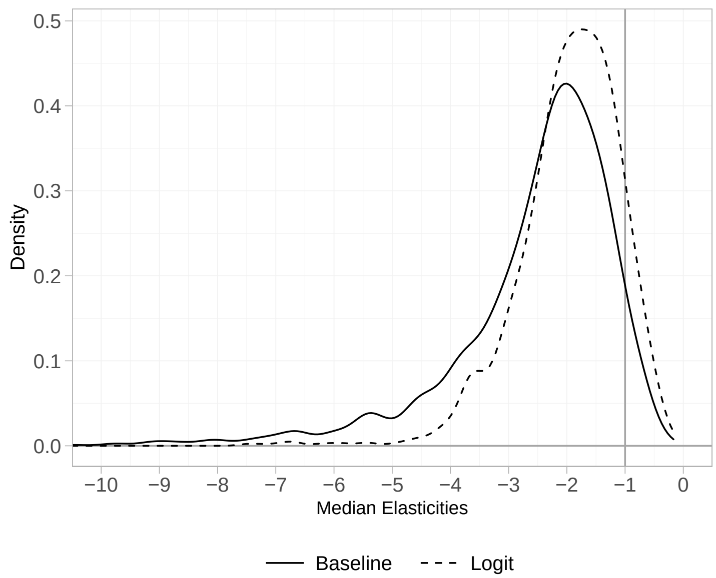

Notes: This figure plots the density of the median own-price elasticity by category and year. The solid black line shows the density of median elasticities using our baseline specification. The dashed line shows the density of median elasticities from a logit specification without random coefficients.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

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

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

D.2 Incorporating Additional Product Characteristics

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

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

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

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

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

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

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

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

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

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

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

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

Table D.1: Estimation Results for RTE Cereals

Panel A: Point Estimates and Standard Errors

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

Panel B: Other Statistics

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

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

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

Panel A: Point Estimates and Standard Errors

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

Panel B: Other Statistics

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

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

Figure D.2: Implied Elasticities Under Alternative Identification Restrictions

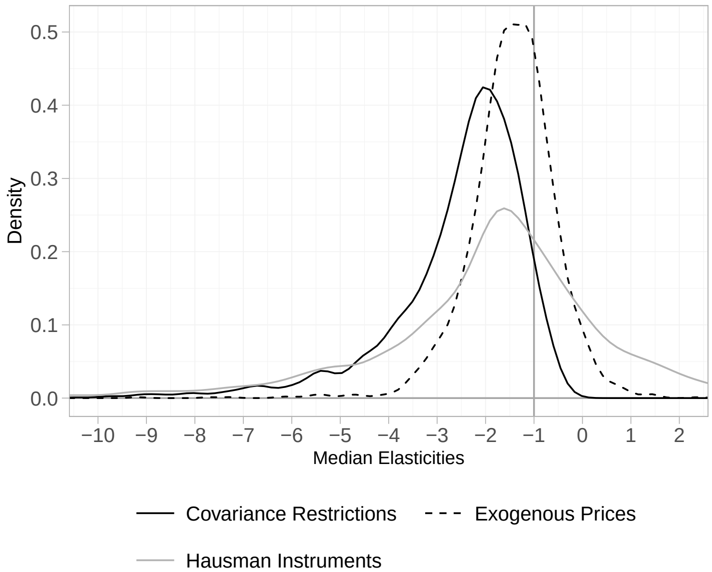

Notes: This figure plots the density of the median own-price elasticity by category and year under different identification assumptions. The solid black line shows the density of implied elasticities using covariance restrictions. The dashed line shows the density of implied elasticities assuming exogenous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

D.3 Alternative Identification Strategies

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

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

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

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

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

D.4 Supply Model

Figure D.3: Changes in Demand Over Time

(a) Elasticity Trend (b) Price Sensitivity Trend

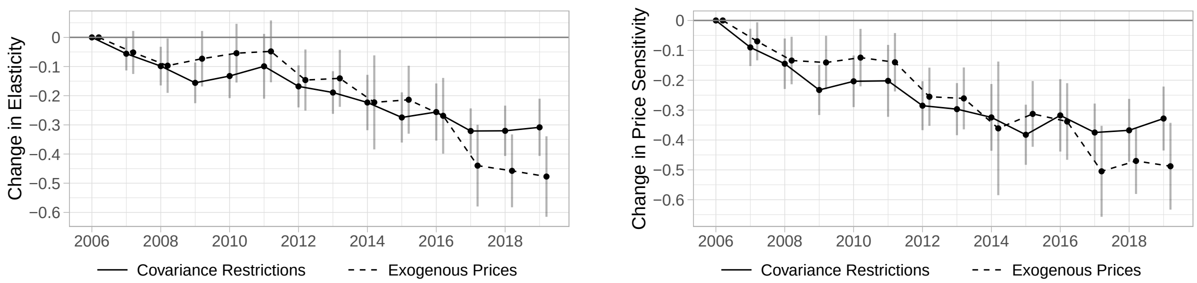

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of the log absolute value of the own-price elasticity (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed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

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

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

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

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

E Robustness Checks

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

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

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

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

E.1 Markup Measure

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

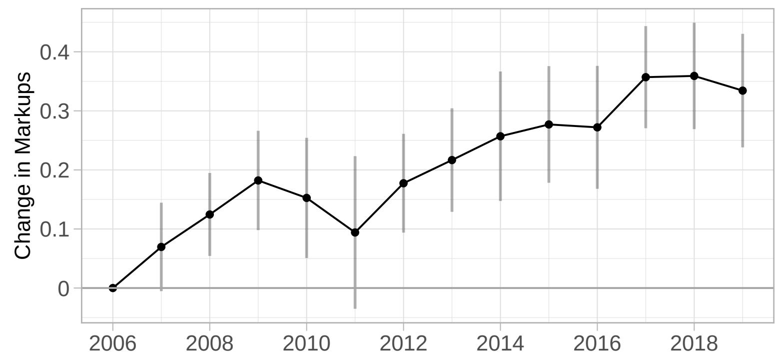

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category. Markups are defined as price over marginal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

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

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

E.2 Product and Retailer Composition

Figure E.2: Balanced Panel

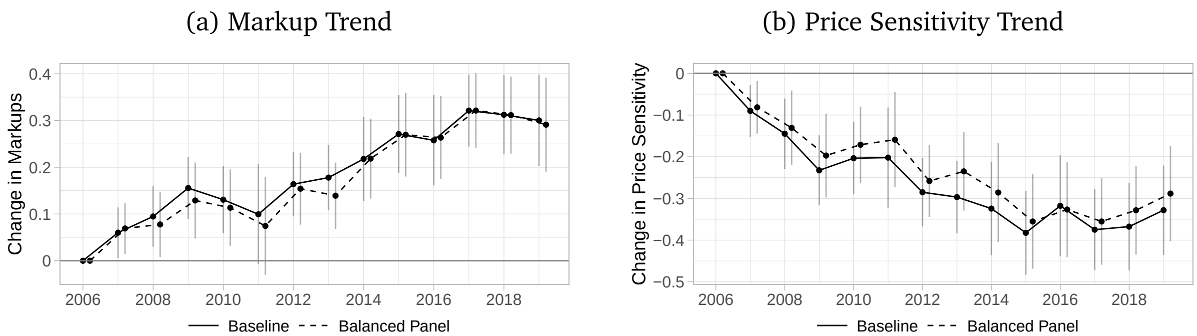

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.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

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

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

E.3 Sample of Categories

Figure E.3: Markups Over Time: Alternative Samples

(a) Markup Trend (b) Price Sensitivity Trend

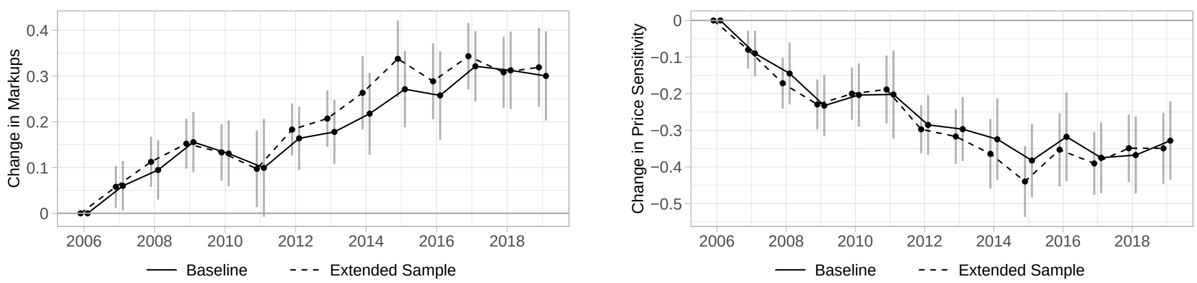

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.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

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

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

E.4 Sample of Retailers

Figure E.4: Including Additional Retail Chains

(a) Markup Trend (b) Price Sensitivity Trend

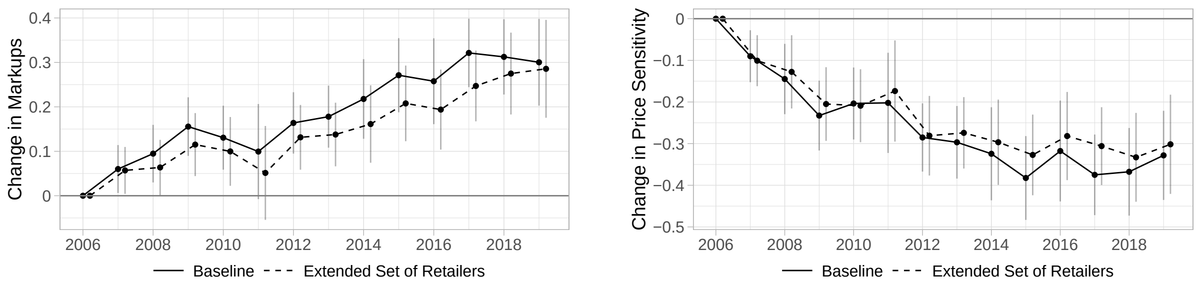

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-time level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.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

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

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

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

E.5 Sample of Geographic Markets

Figure E.5: Smaller and Larger Selection of DMAs

(a) Markup Trend (b) Price Sensitivity Trend

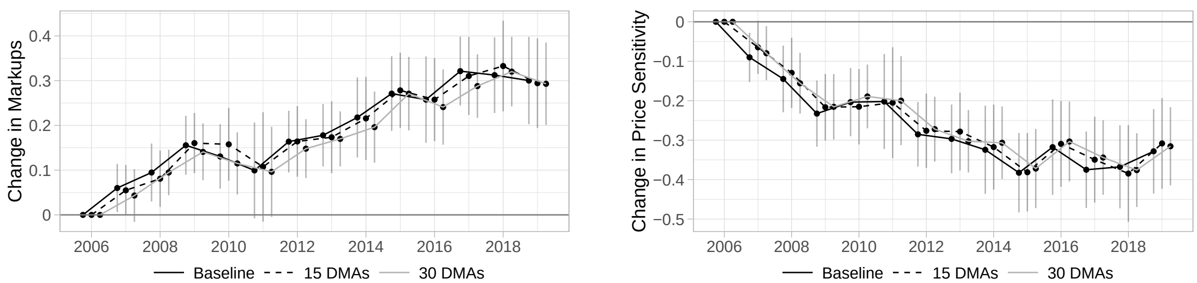

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) on year dummies controlling for year and unit-of-observation fixed effects. The year 2006 is the base category. The baseline estimates are plotted with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

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

E.6 Sample of Brands

Figure E.6: Smaller number of top brands

(a) Markup Trend

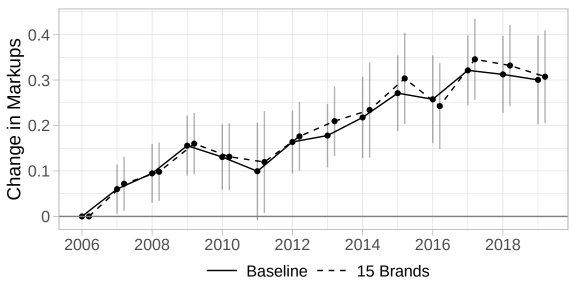

(b) Price Sensitivity Trend

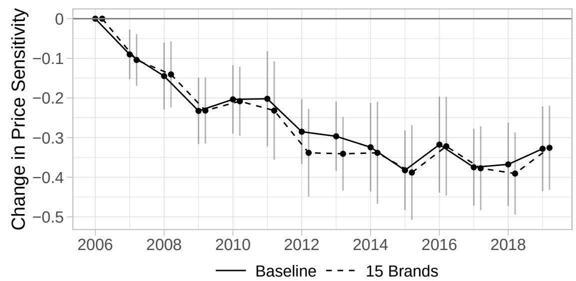

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.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

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

E.7 Time Aggregation

Figure E.7: Alternative Aggregation of Time Periods

(a) Markup Trend

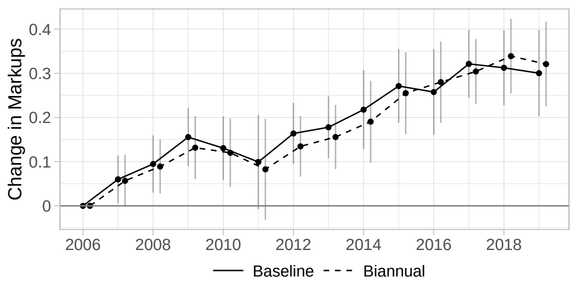

(b) Price Sensitivity Trend

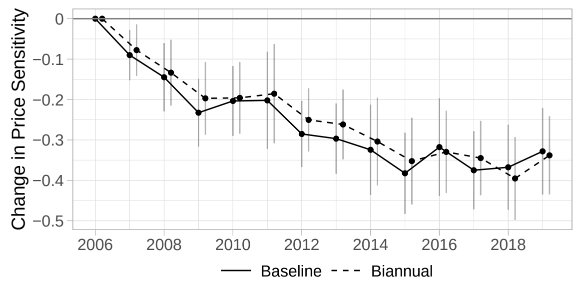

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) at the product-chain-DMA-time level on year dummies controlling for product-chain-DMA and time period fixed effects. The year 2006 is the base category.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

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

E.8 Market Definition

Figure E.8: Alternative Market Definition

(a) Markup Trend (b) Price Sensitivity Trend

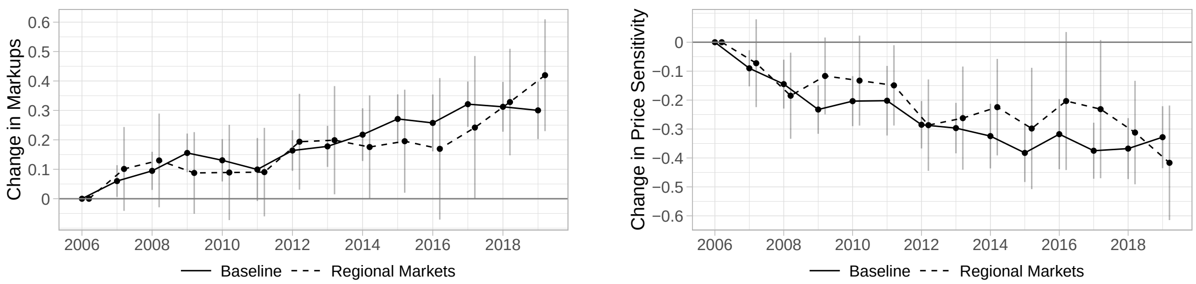

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panel (a)) and price sensitivity (panel (b)) on year dummies controlling for year and unit-of-observation fixed effects. The year 2006 is the base category. In the baseline specification, the unit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

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

E.9 Market Size

Figure E.9: Smaller and Larger Market Size

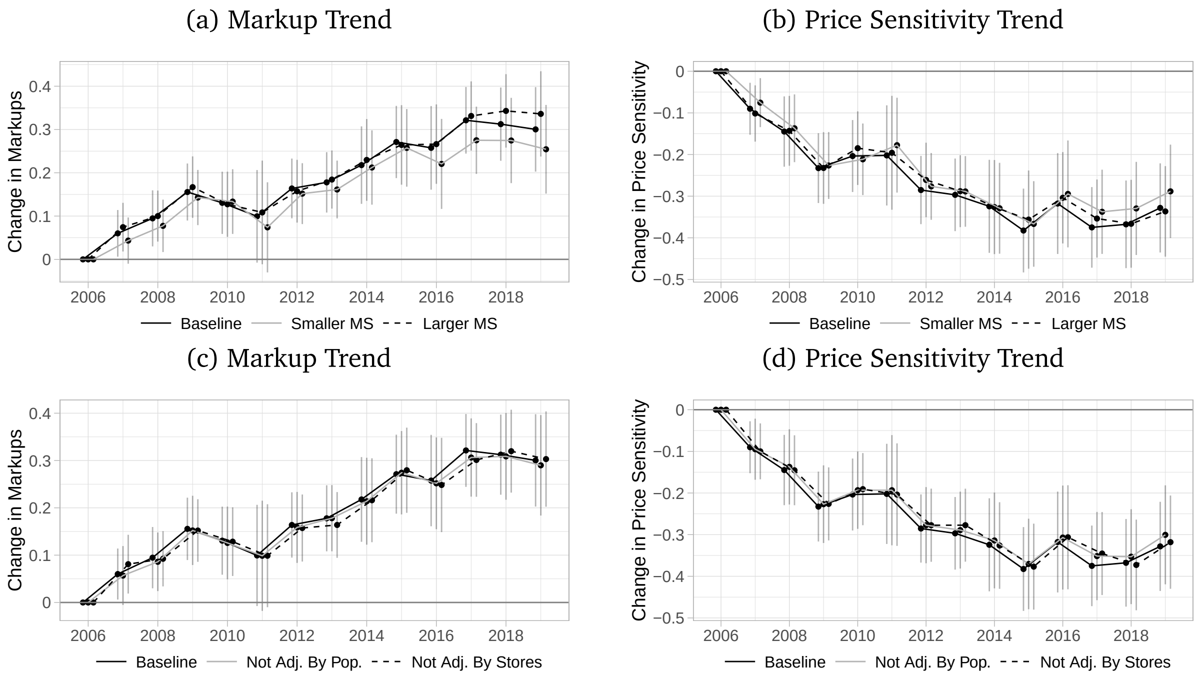

Notes: This figure shows coefficients and 95 percent confidence intervals from a regression of log markups (panels (a) and (c)) and price sensitivity (panels (b) and (d)) at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year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

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

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

F Exploring Alternative Mechanisms

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

F.1 Consumer Spending Across Retail Channels

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

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

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

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

Table F.1: Share of Revenue by Retail Channel

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

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

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

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

F.2 Analysis of Potential Shifts within Data Sample

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

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

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

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

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

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

F.3 Category-Level Variation

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

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

Table F.2: Potential Mechanisms

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

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

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

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

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

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

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

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

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

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

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

G Additional Figures and Tables

Figure G.1: Changes in the Distribution of Markups

(a) Absolute Change (b) Relative Change

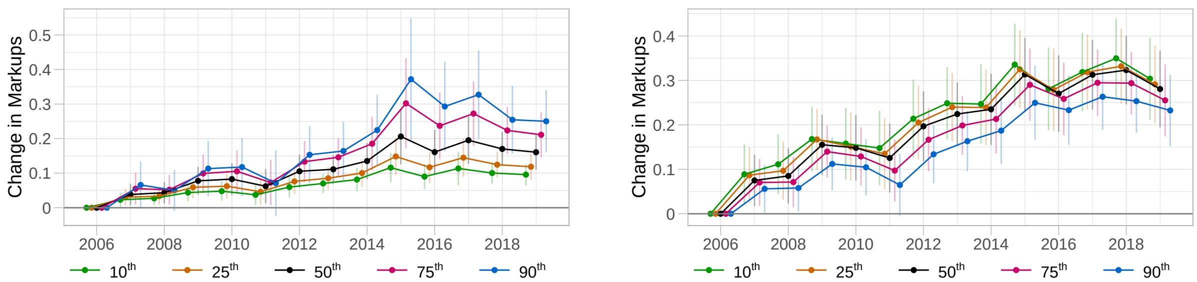

Notes: This figure shows coefficients and 95 percent confidence intervals of regressions of percentiles of the markup distribution at the product category level on year dummies using the year 2006 as the base category. In panel (a), outcomes are percentiles of the level of the Lerner index, $(p - c)/p$,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

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

(a) Prices (b) Marginal Costs

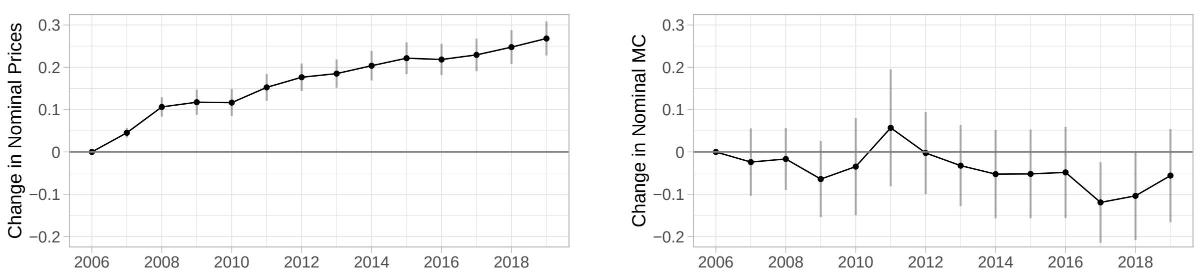

Notes: This figure shows coefficients and 95 percent confidence intervals of regressions of the log of nominal prices and marginal costs at the product-chain-DMA-quarter-year level on year dummies controlling for product-chain-DMA and quarter fixed effects. The year 2006 is the base category.

Figure G.3: Changes in Demand Parameters

(a) Relative Quality

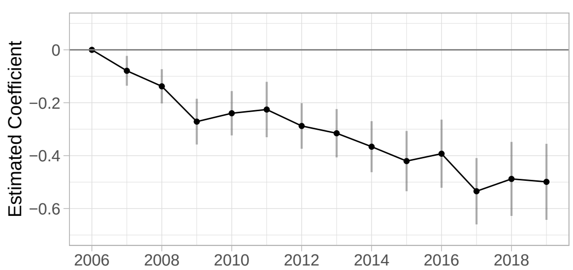

(b) Random Coefficient: $N(0,1)\times$ Intercept

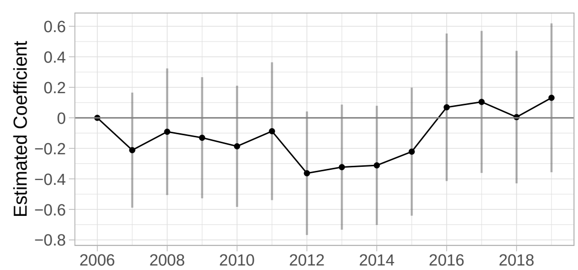

(c) Income $\times$ Price Coefficient

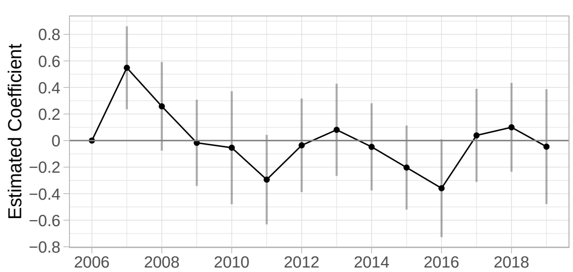

(d) Income $\times$ Intercept Coefficient

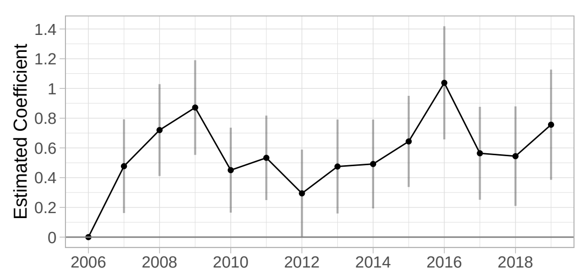

(e) Children $\times$ Price Coefficient

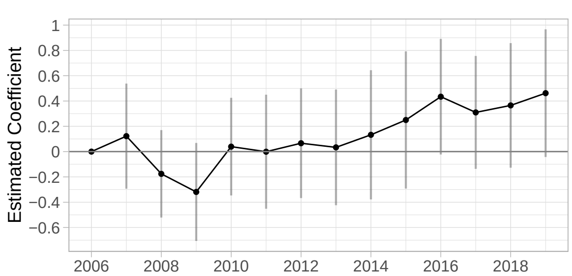

(f) Children $\times$ Intercept Coefficient

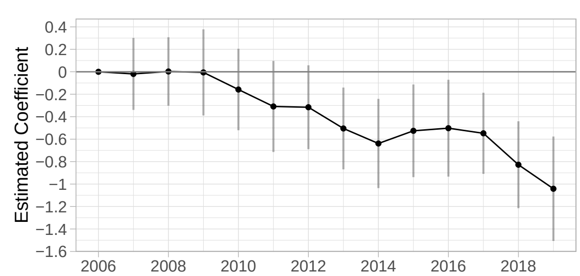

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of standardized demand parameters on year dummies controlling for product-chain-DMA and quarter fixed effects. Observations are at the product-chain-DMA-quarter-year level. The year 2006 is the base category.

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

(a) Market-Level vs. Own-Price Elasticities

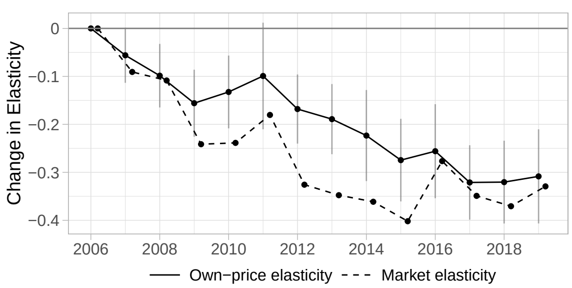

(b) Outside Diversion Ratio

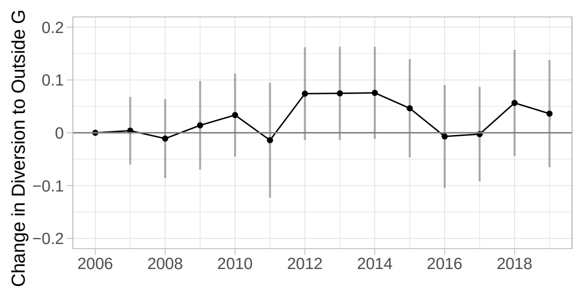

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of own-price and market-level elasticities (panel (a)) and diversion ratios (panel (b)) on year dummies controlling for product-chain-DMA and quarter fixed effects. Market-level elasticities are calculated as the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

Figure G.5: Changes in Market Concentration

(a) Parent HHI

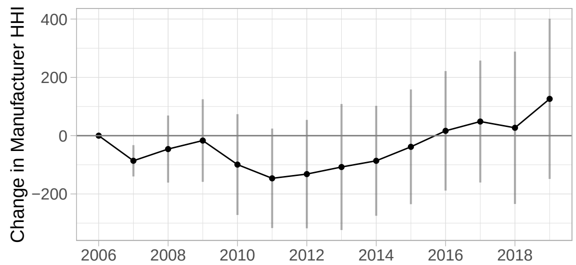

(b) Brand HHI

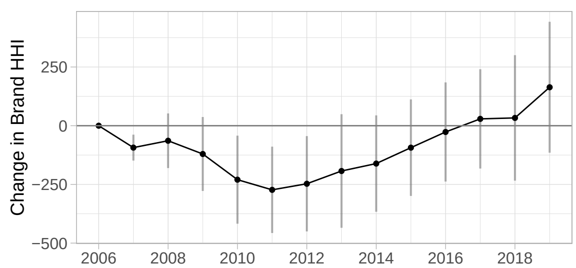

(c) Retailer HHI

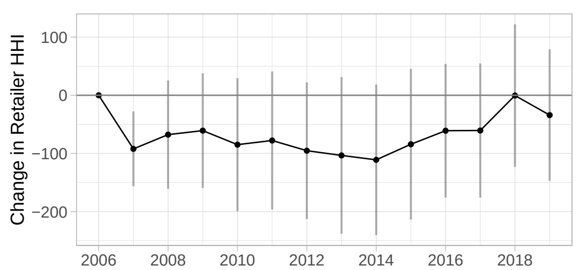

Notes: This figure shows coefficients and 95 percent confidence intervals of a regression of HHI measures on year dummies controlling for product-chain-DMA and quarter fixed effects, with 2006 as the base category. We measure HHI as the sum of squared market shares, where we first adjust market shares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

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

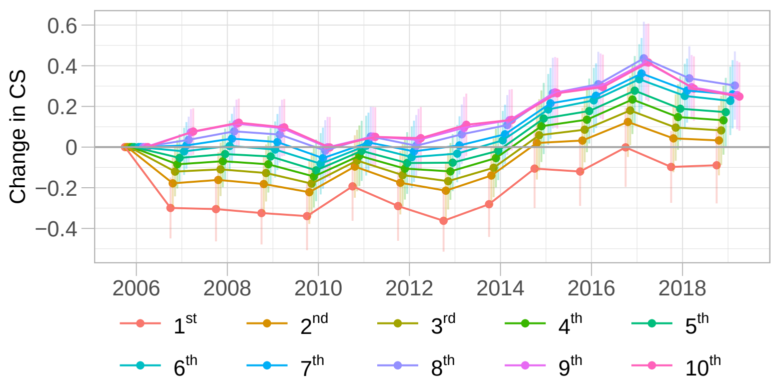

Notes: This figure reports coefficients and 95 percent confidence intervals of a regression of the log of consumer surplus by purchase on year dummies, controlling for category fixed effects, separately for different deciles of the income distribution.

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

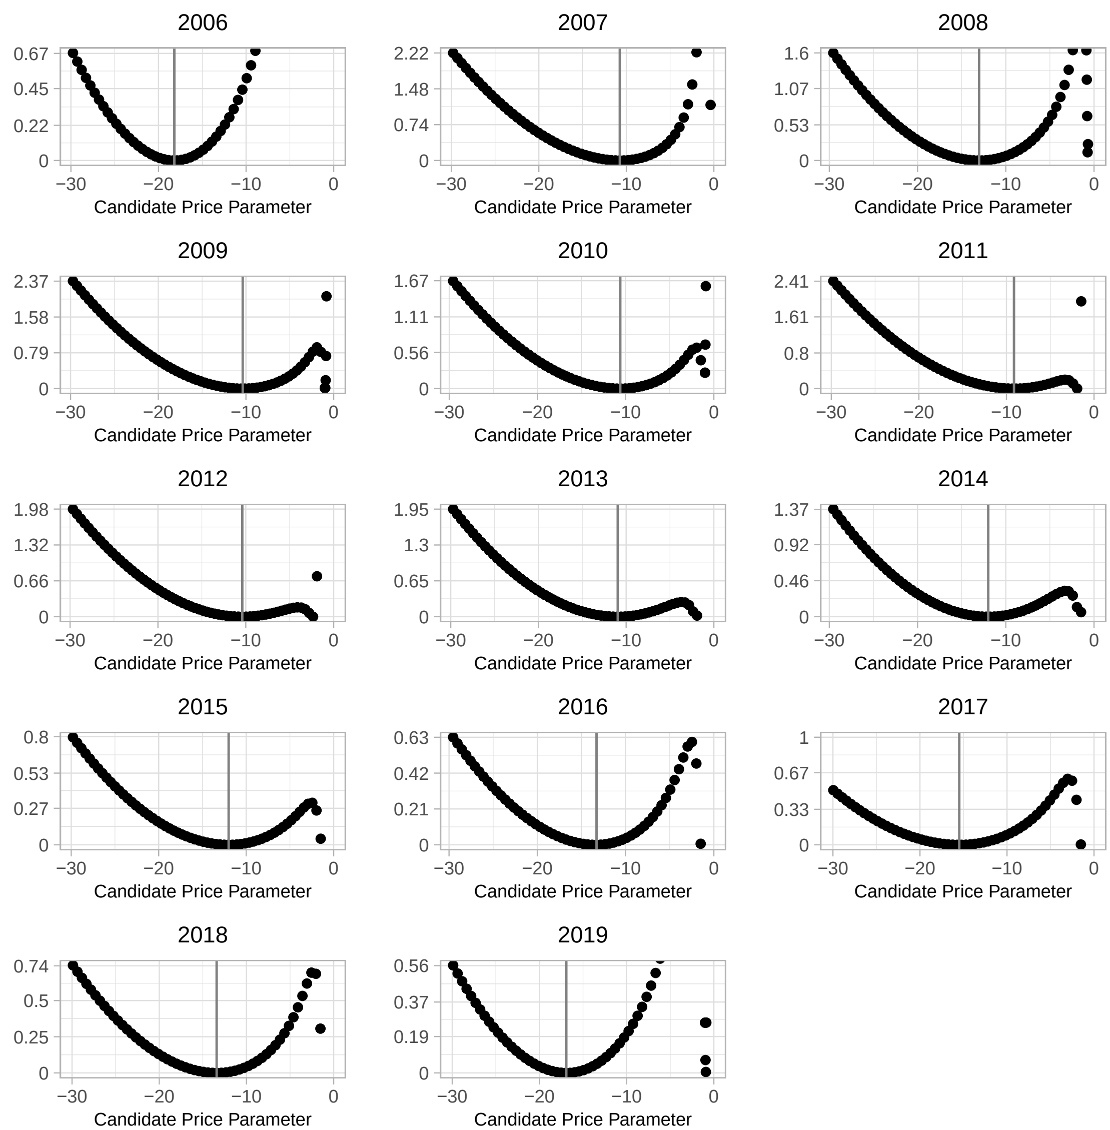

Notes: This figure plots the contribution of the covariance restriction to the objective function, scaled by ten thousand, for different candidate price parameters over the range $[-30, 0]$. Other parameters are held fixed at the levels obtained in the first step of estimation.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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