Several Million Demand Elasticities*

Jordan Rosenthal-Kay
University of Chicago

James Traina
Federal Reserve Bank of San Francisco

Uyen Tran
Arizona State University

22 March 2024

Abstract

A firm's residual demand elasticity takes into account both the household's demand elasticity and the equilibrium response of competitors, and therefore measures its market power. We measure over 9 million of these residual demand elasticities for over 100,000 products in different regions and years using retail scanner data. We find the distribution of these elasticities is stationary over time, suggesting any conclusions that markups are rising in retail markets must be driven by assumptions on conduct. We document substantial spatial heterogeneity in residual demand elasticity estimates, implying the toughness of competition varies considerably more across markets than across time.

JEL Codes: L11, L13, D43, E23, R32
Keywords: Markups, Market Power, Demand Estimation, Secular Trends, Regional Pricing

*Rosenthal-Kay: jrosenthalkay@uchicago.edu. Traina: traina@uchicago.edu. Tran: utran@asu.edu. Earlier versions of this paper were circulated under the title "7 million demand elasticities." This draft includes a few million more. We thank Chris Liao for excellent research assistance. 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 researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

Demand elasticities are a key input into the analysis of markups and welfare. The price elasticity of residual demand faced by a firm measures how its demand changes as it changes its price, anticipating the equilibrium responses of competitors.[1] In a perfectly competitive setting, holding fixed competitors' prices, a marginal increase in price would cause the total evaporation of demand. In an imperfectly competitive setting, changes in price result in finite changes in demand. The residual demand elasticity is thus a measure of market power.

In this paper, we estimate residual demand elasticities for nearly every product in nearly every market-year in the Nielsen Retail Measurement Services (RMS) retail scanner data, made available through the Kilts Center for Marketing at the University of Chicago Booth School of Business. The RMS database records weekly prices and quantities for over 4 million UPCs across 200,000 brands at roughly 40,000 stores per year throughout the entire contiguous US. The data covers total sales worth well over $100 billion per year and accounts for half of spending at grocery and drugstores and a third at mass merchandisers. We measure the price elasticity of residual demand for over 100,000 unique products in 200 spatially segmented markets in each year from 2006 to 2020. Altogether, we estimate over 9 million demand elasticities.[2]

Our main finding is the distribution of these demand elasticities is stationary; we do not find evidence of increased market power over this period. We investigate what drives the stationarity of residual demand elasticity distribution by using standard time-series decompositions. Entry and exit dynamics preserve the stationarity of the distribution: entrants face inelastic residual demand curves, while exiters shift the composition towards products facing more elastic demand curves. We find a small rise in the sales-weighted mean residual demand elasticity over this period, driven by sales volume shifting towards products with more market power, and not changes in the residual demand curves faced by firms.

Our results offer complementary but contradictory evidence to work finding increasingly inelastic demand in the retail scanner data. Relative to others, our work uses a much larger set of product categories, and maintains different and arguably weaker assumptions on the nature of competition. For example, in the first work in this vein, Brand (2021) estimates the distribution of own-price elasticities for nine product-categories: Fruit drinks, soup, cookies, pizza, ice cream, entrees, yogurt, fruit, and light beer. Döpper et al. (2021) extends this work to a larger set of product categories to

[1] Residual demand elasticities are not household demand elasticities. For example, Baumol (1983)'s theory of contestable markets implies residual demand may be elastic even as household demand is inelastic because of the threat of entry. The same contrast occurs with Bertrand duopoly over homogenous products.

[2] Most products do not appear most market-years, hence why we don't estimate 300 million demand elasticities!

reach similar conclusions. Both use the approach of Berry et al. (1995), which specifies a model of competition in which increasing market shares without decreasing prices can drive rising markups.

We find large spatial heterogeneity in residual demand elasticities within products across markets. Our results of significant spatial variation with little temporal variation is consistent with the work of Anderson et al. (2020), who use cost data to infer retailer margins and markups across space and time. We find residual demand is more elastic in larger markets—consistent with tough competition in big cities, as validated by an entirely different approach in Franco (2024). Local concentration, not national concentration, is associated with less vigorous competition in product markets. As local product market concentration has been falling over this period (Rossi-Hansberg et al., 2020), these results lend added scrutiny to the idea that markups may be rising.

Over the last few decades, the US economy has experienced upward trends in national product-market concentration and corporate profits (Berry et al., 2019). The impact of these trends on economy-wide welfare, and therefore the efficacy of antitrust policy, depends on the level of competition and the extent to which firms can raise prices above marginal costs. In short, these trends provoke the question: is market power rising?

Our approach blends theory and data to answer this question in a scalable and transparent way. We closely follow Baker and Bresnahan (1988), which advocates controlling for competitors' costs to isolate a firm's ability to raise prices above marginal cost: its residual demand elasticity. By doing so, we quantify a compound parameter that incorporates both household price sensitivity and equilibrium competitive effects. Our approach needn't and cannot pinpoint the specific sources of market power—whether changing costs, technology, competition, or preferences. We exploit this generality to credibly estimate whether market power is rising for a wide variety of products, each plausibly subject to idiosyncratic sources of market power. We find market power is not increasing: the distribution of residual demand estimates remains stationary.

The literature in industrial organization largely uses demand system modeling to infer markups (Berry et al., 1995; Brand, 2021; Döpper et al., 2021; Grieco et al., 2023).[3] The demand approach calls for researchers to model own-price and cross-price demand elasticities based on market shares, product characteristics, and consumer demographics. Researchers then impose pricing and conduct assumptions to invert firm first-order conditions and recover markups. A rigid structural approach is useful for disentangling the sources of market power and estimating counterfactuals, but it is also costly. In particular, researchers have to assume a time-invariant model of competition such as static

[3] Other lines of work include direct price-cost margins (Anderson et al., 2020; Ganapati and McKibbin, 2023) and the production function approach (Hall, 1988; De Loecker and Warzynski, 2012; De Loecker et al., 2020; Hall, 2018).

Bertrand pricing, which rules out changes in competition as a source of market power. Identifying assumptions such as exogenous product characteristics and no threat of potential entry are often especially questionable for long-run studies.[4]

Our approach follows the large-scale estimation of demand parameters in the industrial organization, trade, and urban literature. For example, Hitsch et al. (2021) and DellaVigna and Gentzkow (2019) estimate demand elasticities at the brand-store level in the scanner data to understand price dispersion within and across establishments. Broda and Weinstein (2006) estimate elasticities of substitution using trade data to construct ideal price indices. Handbury and Weinstein (2010) and Handbury (ming) construct regional price indices for goods in the scanner data. We depart from this literature by studying the residual demand elasticity estimates themselves, rather than use them as inputs to construct markups or price indices. Instead, we closely follow the approach of Tran (2021), who estimates the effect of broadband exposure on brand-county-year demand elasticities.

Overall, our work adds to the growing body of literature on secular trends in markups using a demand approach. Our results offer new insights into how a key measure of market power – residual demand elasticities – vary over time and space in the United States. By understanding these patterns, researchers can build better models of economic behavior, and policymakers and practitioners can make more informed pricing and strategy decisions.

The rest of the paper is organized as follows. We open the paper in Section 2 linking our theoretical framework with our empirical strategy. In Section 3, we discuss the RMS data and our demand model. In Section 4, we present our evidence that demand elasticities have remained stable over time. Section 5 follows this evidence by showing they vary significantly across regions. Finally, in Section 6, we conclude with a summary of our findings and suggestions for future research.

2 From theory to estimation

We present a simple, general environment in which we define the concept of a 'residual demand elasticity.' We closely follow Baker and Bresnahan (1988), though a more modern treatment of the same environment (including a proof of equilibrium existence) is found in Amiti et al. (2019).

In our environment, there are N single-product firms each producing a differentiated good in a

[4] The approach is also data-hungry and computationally costly. For example, estimation is notoriously numerically difficult, and implementation procedures can converge to local optima (Knittel and Metaxoglou, 2014), often forcing researchers to impose common parameters across products and markets.

single market. The demand system for goods in this market is,

$$Q_i = X_i(P_1, ..., P_N, M)$$

where X_i aggregates Marshallian demand across all consumers purchasing goods in the market for i. We assume this demand system is invertible. The term M is a (potentially vector-valued) index of demand shifters, like income. P_i is the price of good i, and Q_i is the amount of good i purchased at prevailing prices. The own and cross-price demand elasticities of this system are given by ϵ_{ii} and ϵ_{ij}.

Residual demand is the demand faced by firm i, accounting for the supply responses of competitors. We now outline all firms $j \neq i$'s supply decisions. We assume all firms in the market solve the static profit-maximization problem,

$$\max_{P_j} \; P_j Q_j - C_j(Q_j, W_j), \quad Q_j = X_j(P_1, ..., P_N, M)$$

where W_j are firm-j specific cost shifters (like factor prices). The statement of this problem assumes firms' have some market power. The first order condition to this problem can be written as marginal cost equals *perceived* marginal revenue.

$$\frac{\partial C_j(Q_j, W_j)}{\partial Q_j} = P_j + Q_j \left(\frac{dQ_j}{dP_j} \right)^{-1}.$$

Marginal revenue is "perceived" in that it depends on firms' beliefs regarding their competitors' reactions. The nature of competition provides structure to the term $\frac{dQ_j}{dP_j}$, in that it incorporates the extent to which firm j internalizes its competitors' price-responses to changes in its own price, $\widetilde{\frac{dP_k}{dP_j}}$.[5] These are the firm's conjectures, indicated by the tilde.

Holding fixed P_i, the system of first-order equations, coupled with each firms conjectures, implicitly determines each firm $j \neq i$'s pricing strategy,

$$P_j = E_j(P_i, W_{-i}, M),$$

where E_j represents the equilibrium relationship between firm i's price and all cost and demand shifters to firm j's pricing decision. E_j is defined implicitly as the solution to all $j \neq i$ pricing

[5] That is,

$$\frac{dQ_j}{dP_j} = \frac{\partial Q_j}{\partial P_j} + \sum_{k \neq j} \frac{\partial Q_j}{\partial P_k} \widetilde{\frac{dP_k}{dP_j}}$$

decisions, holding fixed P_i. The *residual demand function* is,

$$Q_i = X_i\big(P_i, E_{-i}(P_i, W_{-i}, M), M\big),$$

while the *residual demand elasticity* $\beta_i = \frac{dQ_i}{dP_i}\frac{P_i}{Q_i}$ is,

$$\beta_i = \epsilon_{ii} + \sum_{j\neq i} \epsilon_{ij}\rho_{ji}$$

where $\rho_{ji} = \frac{d\log P_j}{d\log P_i}$. The term β_i measures firm i's market power, in that it demonstrates the capacity for firm i to raise prices above marginal costs, considering the supply responses of its competitors. The term β_i does not directly translate to markups. Firm i determines prices by setting its perceived marginal revenue to its marginal cost,

$$P_i\left(1 + \left[\epsilon_{ii} + \sum_{j\neq i}\epsilon_{ij}\tilde{\rho}_{ji}\right]^{-1}\right) = \frac{\partial C_i(Q_i, W_i)}{\partial Q_i}.$$

In this formulation, estimates of β_i are only informative of the markup if $\tilde{\rho}_{ji} = \rho_{ji}$. This condition means that firms' perceptions of the elasticity of their competitors' reaction functions with respect to their own price equals the realized elasticities in equilibrium. This occurs under several cases. Under monopolistic competition, ρ_{ji} and their conjectures $\tilde{\rho}_{ji}$ are all zero. Under Stackleberg competition, if firm i is the leader, it internalizes its effect on the best response of its competitors. Finally, in a consistent conjectures equilibrium (Bresnahan, 1981), firms' conjectures of the slope of their competitors' reaction functions must be equal to the realized equilibrium slope. Despite having the flavor of rational expectations, "rational" or "consistent" conjectures equilibria have undesirable properties, like inconsistent off-equilibrium behavior or multiplicty of equilibria.[6]

We nonetheless view residual demand elasticities as informative of a firm's market power. They inform us exactly how much price-setters can manipulate demand at the observed equilibrium. Translating residual demand elasticities to markups requires some assumptions on the nature of competition, which is no different than recovering markups by inverting firms' first-order conditions, exactly as is done in the demand-estimation literature.

[6] Conjectural equilibria have fallen out of favor in the industrial organizational literature for several reasons. First, arbitrary conjectures can be supported as a conjectural equilibrium, and even restricting to rational (i.e., consistent) conjectures, there may be multiplicity of equilibria. Moreover, consistent conjectures equilibria cannot be interpreted as "Nash in supply correspondences" because agents' conjectures will not align with profitable deviations around a equilibrium perturbations (Makowski, 1987); in short, they will have the "wrong idea" about their competitors' reaction functions (Lindh, 1992).

2.1 Estimation

Our goal is to estimate β_i. However, a simple linear regression of log Q_i on log P_i will be invalid due to the simultaneous equation bias common in demand estimation settings. For firm i, the equilibrium system of equations is,

$$Q_i = X_i\big(P_i, E_{-i}(P_i, W_{-i}, M), M\big)$$
$$P_i = \Big(1 + 1/\tilde{\beta}_i\Big)^{-1} \frac{\partial C_i(Q_i, W_i)}{\partial Q_i}.$$

This system of equations shows that P_i depends on Q_i through its effect on marginal cost and perceived marginal revenue (the source of simultaneous equation bias), but also suggests that W_i is a valid instrument for P_i: it is relevant, as it shifts P_i monotonically through the optimal pricing equation and it is excluded from the residual demand equation. However, W_i does not immediately satisfy independence; it may be correlated with the error term. To see this, we log linearize the residual demand equation,

$$d \log Q_i = \beta_i d \log P_i + u_i$$

where the error term is,

$$u_i = \sum_{j \neq i} \epsilon_{ij} \left(\frac{\partial \log E_j}{\partial \log W_j} d \log W_j + \frac{\partial \log E_j}{\partial \log M} d \log M \right) + \frac{\partial \log X_i}{\partial \log M} d \log M$$

The instrument, $\log W_i$, may be correlated with u_i if cost shocks are correlated across firms within a market, $\mathbb{E}[W_i W_j] \neq 0$.[7] Consequently, a regression of $\log Q_i$ on $\log P_i$, using $\log W_i$ as an instrument, and controlling for $\log W_j$, recovers an estimate of β_i. Once we control for competitors' costs, competitors' residual price variation, which affects demand for i, is either idiosyncratic or due to strategic responses from changes in P_i. Said differently, changes in Q_i associated with exogenous shifts in P_i come from either movements down the demand curve (ϵ_{ii}), or through competitors' supply responses ($\epsilon_{ij} \rho_{ji}$). This estimation strategy thus isolates movements along the residual demand curve.

[7] It is also problematic if W_i are correlated with market-level demand shifters M. We handle this in practice by using cost variation uncorrelated with local market characteristics; see Section 3.2.

3 Estimating residual demand elasticities

3.1 Data

Our analysis uses the Nielsen Retail Measurement Services (RMS) retail scanner data from 2006 to 2020, available through the Chicago Booth Kilts Center for Marketing. RMS has information on UPC-store-week prices and quantities covering over 60,000 stores from about 200 retailers across all states in the contiguous US. Nielsen aggregates transaction-level sales and quantities over each week. Consequently, we observe quantity-weighted average prices, and we miss prices which have zero transactions.

The dataset covers a range of consumer packaged goods commonly available in brick-and-mortar retail stores. The data are big, with annual sales over $100 billion and total file sizes approaching 10 TB. It accounts for over 50% of spending at grocery and drugstores, and over 30% at mass merchandisers. Beraja et al. (2019) validates the representativeness of the RMS price data by showing its aggregated values closely track BLS price indices.

The dataset includes 1,100 product categories (modules) across 110 product groups. We exclude categories that aren't consistently recorded from 2006 to 2020, and focus on the largest categories representing 99% of total sales. This process narrows the dataset to 800 categories across 90 groups. Categories are granular: The top five by sales are cigarettes, refrigerated dairy milk, ready-to-eat cereal, fresh bread, and carbonated soft drinks, which collectively account for 10% of our total sales.

From these categories, we remove UPCs identified by Nielsen's documentation as inconsistent over time. Specifically, we exclude UPCs with changes in core attributes like size or category, and those measured in units not standard to the category (e.g., in "units" when every other product is measured in "ounces"). This filter makes sure prices and quantities are comparable across competitors. We standardize prices and quantities to a base unit, such as 1 cigarette or 1 ounce of milk, regardless of how they're packed or packaged, such as in 20 packs or 128 ounce gallons.

Following best practices in industrial organization and quantitative marketing, we aggregate UPCs into brands, considering each brand as a distinct product. Before introducing filters required to estimate demand elasticities, our dataset has over 4 million UPCs and over 200,000 brands.

We define our geographic market areas by media markets using Nielsen's Designated Market Areas (DMAs), based on television and other shared media coverage. In any year, Nielsen samples between 30,000 and 50,000 stores across about 200 DMAs. Metropolitan areas such as Chicago, Philadelphia, and Phoenix have their own DMA, while other smaller areas roll up into larger areas, like Racine into Milwaukee.

Stores belong to around 200 different retail chains. These retailers can change based on acquisitions. We focus on stores that have maintained a consistent retailer identity, allowing for changes at the corporate parent level that don't affect the store's brand on the street. We also exclude liquor stores from our study, given their unique and often complex regulatory and pricing environment.

Given the internal data organization by Kilts, we have thousands of structured files that record observations across brands, retailers, markets, and weeks, segmented by category and year. Given that these panels sometimes have missing data, and we don't want to infer variation from observations jumping in and out of our models, we include only those panels where over 80% of the weeks have data. To fill in occasional gaps (at most 10 weeks per year), we carry forward the last known nonmissing value.

In our analysis, we estimate residual demand elasticities at the product($\times$market$\times$year) level instead of the firm level. Implicit in this methodological choice is that individual products-markets are the primary unit of pricing, instead of firms which manage portfolios of multiple products that are sold in multiple markets. This choice is intentional; our theoretical framework accounts for internal competition effects in the same way as it accounts for external ones. Underlying this assumption is that markets are spatially and temporally segmented, and that firms do not internalize cross-market nor cross-year substitution effects. Moreover, our estimates average over product retailers' market power, which plausibly may vary across vendors within a market-year, but is likely to be small if competition is tough across retailers of the same product within a given market-year.

Uniquely in the literature, our dataset accounts for private label products, sometimes called "store brands." Nielsen masks the UPCs and brand identifiers of these products to preserve the confidentiality of reporting retailers. While we don't know the exact identity of a private-label product, we do know sales of products in a particular category at a particular retailer. We group these together under one brand. For example, if we saw Italian and French roast coffee sold at retailer Tralymart, we would consider both under the Tralymart coffee brand. Keeping private-label products is important because they constitute a growing share of total retail sales Hitsch et al. (2021), and they plausibly face more elastic residual demand.

3.2 Estimation

Our estimating equation mirrors the log first-order approximation to the residual demand curve, so that we identify the residual demand elasticity facing the firm,

$$q_{jmrt} = \beta_{jmy} p_{jmrt} + \theta_{jmy} w_{jmrt} + \alpha_{jmy} + \gamma_{rky} + \tau_{kt} + \varepsilon_{jmrt}, \tag{1}$$

where q_{jmrt} and p_{jmrt} denote log quantities and prices of brand j in market m at retailer r in week t, and β_{jmy} is the brand-market-year residual demand elasticity. ε_{jmrt} is the structural error term, capturing unobserved factors that affect residual demand.

We use fixed effects to subsume the intercept of the residual demand curve, which will vary across product markets. The term α_{jmy} accounts for brand-market-year effects, absorbing residual demand variation because of local market shocks and unobserved product quality variation that we allow to differ for each market-year pair. Brand-market-year is also typical of market share measurement, so α_{jmy} removes all low-frequency residual demand shocks from, e.g., market concentration. γ_{rky} captures retailer-category-year effects, designed to absorb the influence of retailer-wide unobserved advertising efforts or promotional campaigns. Retailers often undertake marketing strategies that can alter the demand for brands across different categories within their stores. Finally, τ_{kt} accounts for category-week effects, absorbing residual demand variation because of seasonality and time-varying demand components associated with prices.

To address the further potential endogeneity of prices in our model, we build an instrument following DellaVigna and Gentzkow (2019) that draws inspiration from the methodologies of Hausman (1996) and Waldfogel (2003). Per Baker and Bresnahan (1988), brand-specific cost shifters identify the elasticity of the residual demand curve, β_{jmy} when competitors' costs are held fixed. We instrument log prices with the log average price of the same brand-retailer-week, but calculated over other markets. Conditional on our fixed effects and competitor costs, the assumption needed to estimate β_{jmy} consistently is that changes in the pricing decisions that apply to the entire retail chain are orthogonal to local competitive conditions for a brand. This exclusion restriction is consistent with evidence of uniform pricing across stores despite substantial variation in demand elasticities, as documented in DellaVigna and Gentzkow (2019) and Hitsch et al. (2021). Under this standard assumption, the interpretation of our estimates are that $\hat{\beta}_{jmy}$ represents the own-price short-run elasticity of residual demand facing the seller of a given brand in each market-year.

To identify the influence of price changes on residual demand, we not only use a brand's own costs as an instrument to tackle the potential endogeneity problem inherent in prices, but we also include a model-consistent proxy for competitor costs w_{jmrt}. This variable measures the weighted averages of costs for competing brands in the same category-market-week. The costs underlying this average are the same costs we use for those brands as instruments. The coefficient associated with this control, θ_{jmy}, quantifies the effect of competitors' costs on a brand's residual demand.

Following the structure of the raw Nielsen data, we process files by category and year. Within these files, we require at least 52 observations for each brand-retailer-market to make sure we're not basing our residual demand estimates on tiny samples. This requirement might imply data for every week of the year or, for instance, data from two retailers, each contributing 26 weeks. Because private-label products have unique retailers, they must have a full year's data to meet our analysis threshold. After estimation, we apply basic inclusion criteria for our further analysis: we only keep elasticity estimates that fall between -100 and -0.1 and have standard errors between 0.001 and 100. Estimates must also have t-statistics less than -2, ensuring they are statistically significantly different from 0 at the 5% level. This filter cuts outliers that may be unreliable or have an undue influence on the results.

We estimate 9 million residual demand elasticities for 100,000 brands over 200 markets from 2006 to 2020. Our dataset is only 3% of the theoretical maximum of 300 million brand-market-year combinations because not every brand is available in every market in every year.

3.3 Empirical Bayes Adjustment for Sampling Error

Estimates $\hat{\beta}_{jmy}$ from the instrumental variable, fixed effect regression are noisy. We shrink overly elastic or inelastic but noisy estimates using an empirical Bayes procedure. Our procedure uses information from the elasticity distribution within a category-year as a prior to update each individual brand-market-year elasticity estimate. Empirical Bayes models are computationally tractable approximations to hierarchical Bayes models.

Using empirical Bayes corrections for noisy estimates are common in literatures with large-scale parameter estimation. For example, hierarchical models have been used to better measure teacher value-added (Kane and Staiger, 2008; Chetty et al., 2014), hospital effects for clinical outcomes (Chandra et al., 2016), and demand elasticities using scanner data (DellaVigna and Gentzkow, 2019; Hitsch et al., 2021; Brand, 2021). Brand (2021) argues that because of growth in sales and the number of products, earlier demand elasticity estimates may be less precise than those estimated in more recent years. Failure to account for changing estimate precision over time and space may cause the researcher to mischaracterize elasticity trends.

Our underlying hierarchical model assumes that our estimates are noisily distributed around their true values $\hat{\beta}_{jmy} \mid \beta_{jmy} \sim N(\beta_{jmy}, \sigma_{jmy})$. Moreover, we assume that $\beta_{jmy} \sim N(\mu_{ky}, \sigma_{ky})$; the

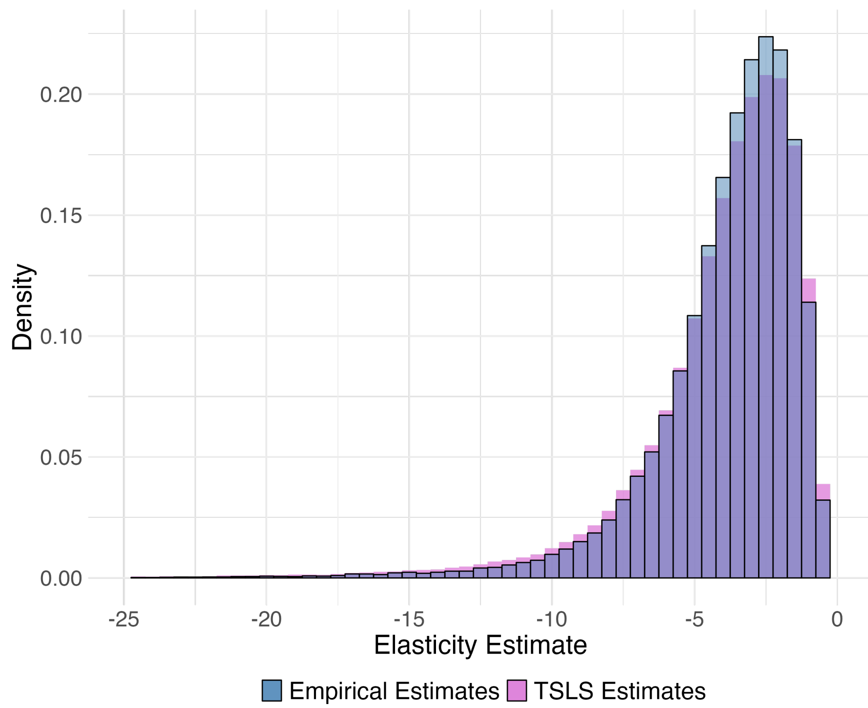

Figure 1: Distribution of TSLS and empirical Bayes (EB) estimates of the distribution of product-market-year demand elasticities, $\hat{\beta}_{jdy}$.

true elasticities within each product-year are normally distributed around a product-year mean and variance. In short, we assume a normal likelihood and normal prior. Following Morris (1983), we estimate μ_{ky} and σ_{ky} with an inverse-variance weighted regression of our first stage β_{jmy} estimates on a constant. Our posterior estimates shrink noisy outlying elasticities towards the product-year mean, reducing noise while retaining information across elasticities.

Figure 1 displays the histogram of estimated product-market-year demand elasticities before and after the Empirical Bayes' correction. The distribution of TSLS estimates is left-skew and centered below zero, though there is some positive mass, and has very long tails. The EB distribution is by construction shifted left and has considerable mass around zero, due to a nontrivial amount of precisely estimated positive demand elasticities. The median of the sales-weighted EB estimates is -3.29, though the distribution is left skew and its mean is -3.88.

4 Stable Secular Trends

The demand elasticity time-series Figure 2 displays the time series for the middle 80% of the distribution of demand elasticities, weighted by their sales share. The median elasticity of about -3.3 stays almost constant over time, becoming slightly more inelastic after 2017. Almost all the

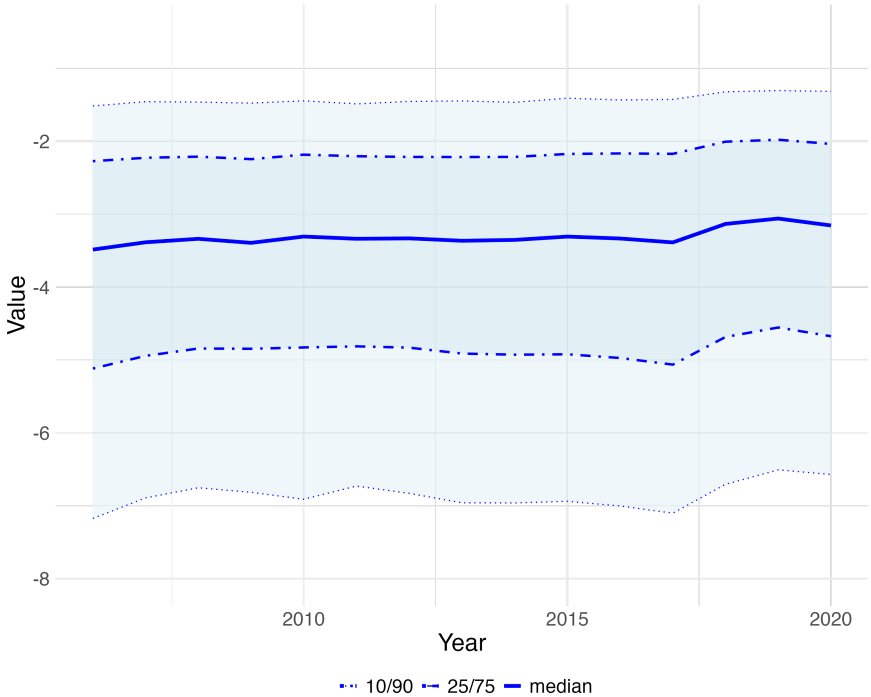

Figure 2: The time series of the distribution of demand elasticities estimates, weighted by revenue. The light-blue colored bands indicate percentiles of the distribution. The thick dashed line is the medium, while the thin dashed lines represent the 25th and 75th percentiles of the distribution.

other percentiles of the distribution remain constant in time as well.

Less than 4% of the sales-weighted elasticities in our data are measured *above* -1 ($\approx 4.5\%$ unweighted), indicating *inelastic* demand. These estimates are inconsistent with monopoly pricing, wherein the profit-maximizing monopolist always prices along the elastic portion of the demand curve.

In the following sections, we investigate the sources of the stationarity of the residual demand elasticity estimates time series.

Product-specific trends We estimate product-region time trends through a linear regression,

$$\log(-\hat{\beta}_{jmy}) = \delta_{jm} \times y + \xi_{jm} + \nu_{jmy}, \tag{2}$$

where ξ_{jm} is a product-market specific intercept. The transformation of the dependent variable means that the unique product-region coefficient δ_{jd} has the interpretation of the average annual percent change in the demand elasticity. Moreover, we flip the sign on the elasticities, so that

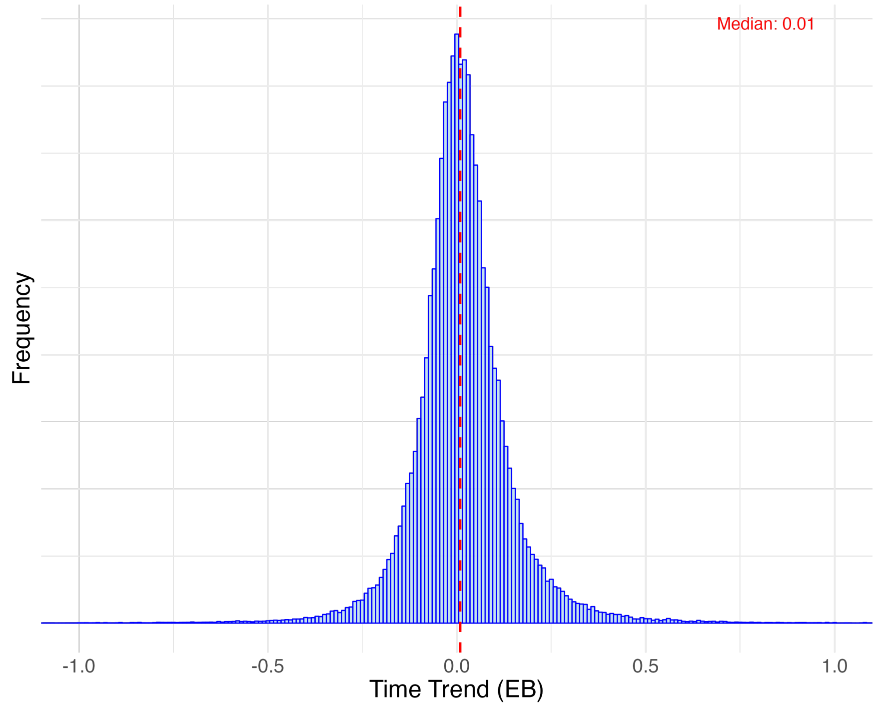

Figure 3: The distribution of individual time trends. The dashed line represents the median of the distribution.

$\delta_{jd} < 0$ indicates demand becoming *more inelastic*.

We estimate product-region unique time trends by first-differencing (2) and recovering the fixed effects, $\hat{\delta}_{jd}$. This ameliorates the need to estimate an extraordinary number of fixed effects and slopes jointly. Figure 3 displays the middle 98% of the distribution weighted by sales volume.

The median of the revenue-weighted distribution (dashed line) is slightly to the right of zero; more than half the estimated elasticities are trending more elastic in time. The has distribution of time-trends has substantial spread; the middle 50% of the distribution corresponds to of annual percent changes between -4.5% and 9.0%. This corresponds to considerable heterogeneity in the competitive environment faced by different product-regions over time. Product-regions at the 25th percentile of the trend distribution could, at most, face demand in 2020 that is 50% less elastic than that what they faced in 2006. At the 75th percentile, residual demand becomes almost four times more elastic over a fifteen year period. This suggests for product-regions facing increasing competitive pressures, competition quickly toughens.

Marshall-Edgeworth Decomposition We use a Marshall-Edgewoorth decomposition for the mean of the distribution to assess the underlying drivers of the time series' stationarity.

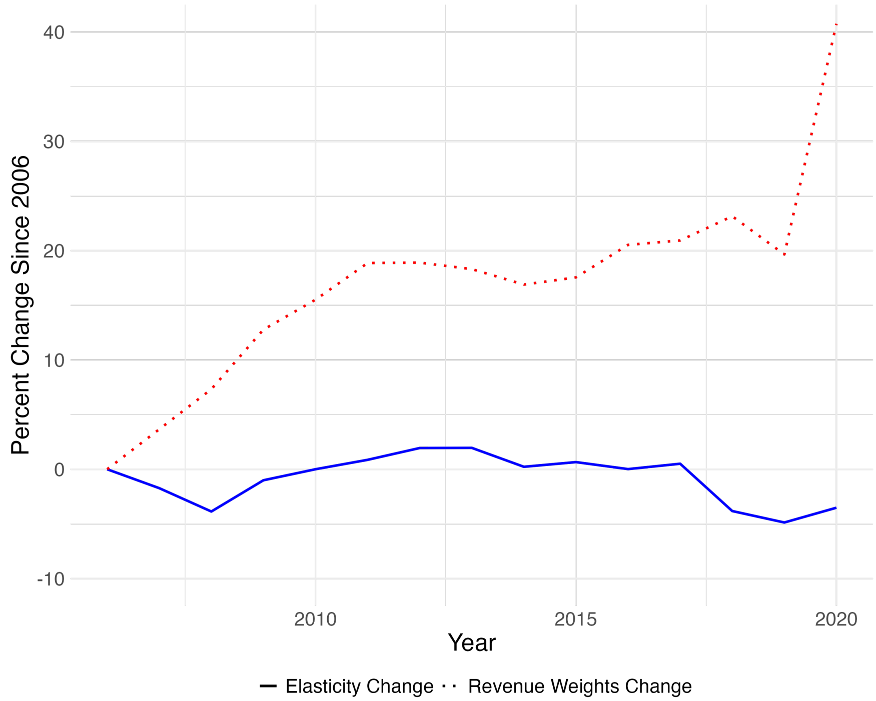

Figure 4: Accumulated Marshall-Edgeworth decomposition

For a given change in mean,

$$\Delta\bar{\beta}_{y,y-1} = \sum_i s_{i,y}\beta_{i,y} - \sum_i s_{i,y-1}\beta_{i,y-1}$$

where i indexes product-market j, m and $s_{i,y}$ are sales shares in product-market i. We use the decomposition,

$$\Delta\bar{\beta}_{y,y-1} = \underbrace{\sum_i \left(\frac{s_{i,y} + s_{i,y-1}}{2}\right)\Delta\beta_{i,y}}_{\text{Elasticity change}} + \underbrace{\sum_i \Delta s_{i,y}\left(\frac{\beta_{i,y} + \beta_{i,y-1}}{2}\right)}_{\text{Revenue weights change}}$$

The first term reflects how the mean elasticity changes holding revenue-weights $s_{i,y}$ at the mean value while the second term shows the component of the mean that shifts due to shifting demand, reflected by changes in revenue, holding the elasticities fixed at their mean.

We plot the accumulation of these terms relative to 2006 In Figure 4. Accumulating these terms means that an observation $x_y = \sum_{t=2006}^{y} \Delta\bar{\beta}_{t,t-1}$. Thus, each line represents the total shift in the mean β_{jmy}, relative to 2006, due to changes in the residual demand estimates and changes in sales.

The resulting figure shows that while the mean residual demand elasticity has remained unchanged

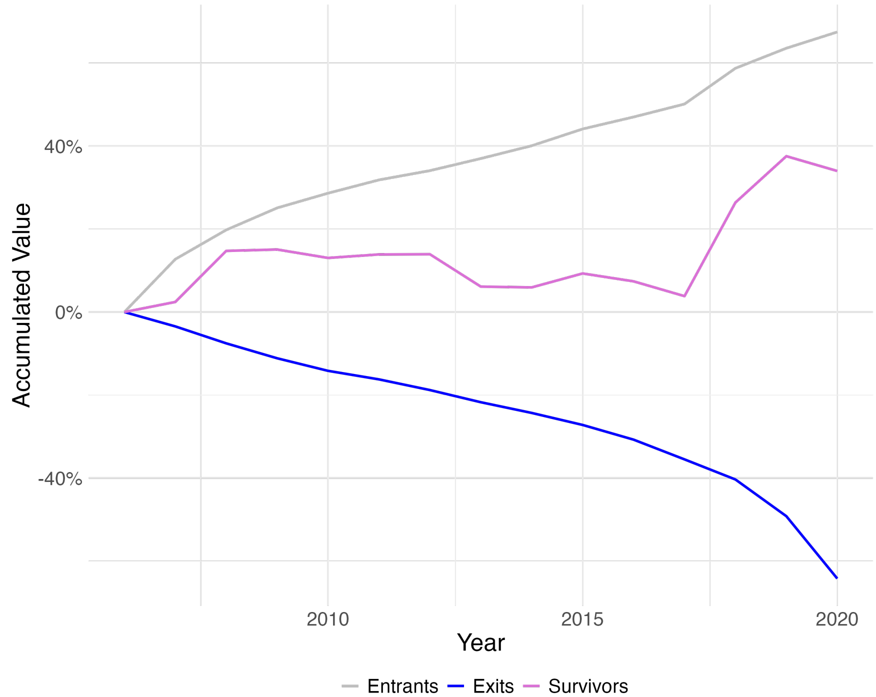

Figure 5: Accumulated Marshall-Edgeworth decomposition

over the period, consumer demand (i.e., sales) shifted towards goods facing more *inelastic* residual demand. Consequently, by 2019, the revenue-weighted mean residual demand elasticity had increased by approximately 20%. Were we to map our estimates to markups, this would suggest moderate increases in the average markup faced by households, entirely driven by demand shifting to higher-markup products.

Dynamic Olley-Pakes decomposition To assess whether shifting demand and changing demand elasticities are due to the composition of available products, we follow Melitz and Polanec (2015) and decompose the mean of the demand elasticity distribution into three terms, reflecting changes in the mean due to entry and exit.

Figure 5 reports the results of this exercise, again accumulating changes in each year relative to 2006. While entering products shift the mean demand elasticity increasingly inelastic, this is off offset by the fact that exiting firms shift the mean of residual demand elasticity towards more elastic estimates. The small rise in the mean of of the distribution over time is driven by surviving firms.

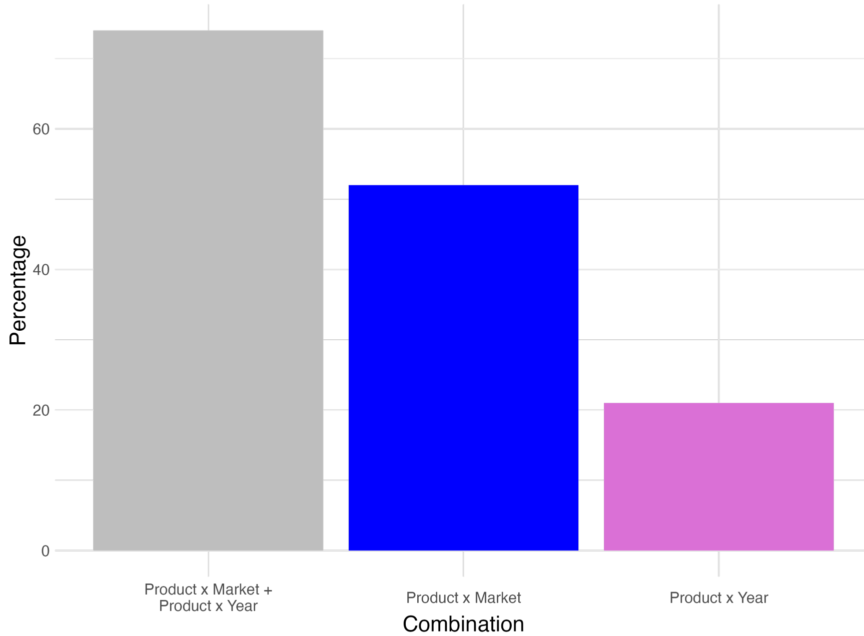

Figure 6: Partial R-squared statistics for product-market and product-year fixed effects and their sum.

5 Significant Spatial Variation

In this section, we report substantial variation in estimated elasticities across observations within a module-brand. To decompose the sources of variation, estimate,

$$\hat{\beta}_{jkmy} = \xi_{jm} + \xi_{jy} + u_{jkmy} \tag{3}$$

and report partial R-squared statistics for product$\times$market fixed effects ξ_{jm} and the product$\times$year fixed effects ξ_{jy}. This exercise can be interpreted as an anova decomposing the sources of variation within a product over spatially segmented markets and time.

Figure 6 reports the results of this exercise. First, these fixed effects account for 74% of the variation in product level residual demand estimates – only a quarter of the variation in residual demand within a product is due to market-year idiosyncrasies. Within a product, market fixed effects account for approximately 50% of the variation in estimated residual demand elasticities while year effects account for around 20%. Put differently, the there is more than twice as much variation across spatially-defined markets than across years within a product category.

Micro and macro correlates What explains the spatial and time-series variation in residual demand estimates within a product? We investigate this question by projecting our residual demand elasticity estimates onto product, local, and national observables, and use both cross-sectional and within-market time series variation to uncover correlates of residual demand elasticities.

In particular, we estimate,

$$\hat{\beta}_{jkmy} = \underbrace{X'_{jkmy}\theta_1}_{\text{product}\times\text{ market covariates}} + \underbrace{X'_{kmy}\theta_2}_{\text{local market covariates}} + \underbrace{X'_{jky}\theta_3}_{\text{nat'l market covariates}} + v_{jmy}, \tag{4}$$

where X contains observables that vary at the product, module, local market, and national market level.

Table 1 estimates equation (4) using both cross sectional and panel variation. For inference, we cluster our standard errors at the product level to account for correlated errors within product across markets and years. For product×market characteristics, we include a firm's market share, its entry/exit status, and if a 'survivor,' its tenure. For local market characteristics, we include the population and average income of consumers in the market (i.e., the total population and average income of households in a given DMA), as well as the market-product category HHI and the (log) total number of product with positive sales in a given market-product category-year. For 'national' covariates, we include a product's market share across all markets in its product-category, as well as the HHI of that product-category, assuming a national market.

Products with larger market shares face more inelastic residual demand, consistent with oligopoly theories in which market shares predict market power. Products that enter and exit a given product-region face more inelastic demand than survivors, for whom tenure is associated with toughening competition: on average, each decade in a market is associated with a small decline in the residual demand faced by a firm.

Products in larger and richer markets face more elastic residual demand curves, consistent with tougher competition in bigger markets, as in Franco (2024). Moreover, on average as the number of distinct products selling in a product category-region increases, residual demand becomes more elastic. This is also consistent with oligopoly theory in which entry toughens competition. Likewise, increases in local market HHI are associated with small increases in the residual demand elasticity, suggesting that market concentration is on average associated with less vigorous competition. However, 'national market characteristics' have the opposite sign: products with large national market shares face considerably more elastic residual demand curves, and products that operate in markets that appear concentrated at the national level actually face, on average, more elastic residual demand curves.

	(1)	(2)	(3)	(4)
Product-region characteristics				
Local market share	0.05			-0.27
	(0.01)			(0.01)
Entrant	0.10			0.11
	(0.01)			(0.01)
Exit	0.22			0.17
	(0.00)			(0.00)
Tenure/10	-0.04			0.01
	(0.01)			(0.01)
Local market characteristics				
Log population		-0.07		-0.06
		(0.00)		(0.00)
Log income		-0.00		-0.05
		(0.01)		(0.01)
Log N. products		-0.09		-0.12
		(0.01)		(0.01)
Local market HHI		0.02		0.15
		(0.01)		(0.01)
National market characteristics				
Nat'l share			-0.76	-0.65
			(0.09)	(0.09)
Nat'l market HHI/10			-0.03	-0.04
			(0.00)	(0.00)
N	9,086,894	8,393,112	9,086,894	8,393,112
Product FE	✓	✓	✓	✓

Table 1: Standard errors clustered at the product level in parentheses.

6 Conclusion

When demand faced by a firm is price-inelastic, firms are able to raise prices above marginal cost, potentially harming consumers and inducing allocative inefficiencies in factor markets. The extent to which this happening, and whether these wedges vary over time and space is an open question in the macroeconomics and industrial organization literature. We hope to shed new light on these questions with a transparent and empirically scalable approach. In this paper, we examine the residual demand elasticities faced by sellers of over 130,000 products in different markets using Nielsen Retail Measurement Services (RMS) scanner data from 2006 through 2020.

Our basic idea is to avoid using a structural model to recover markups, and instead directly examine firms' capacity to raise price over marginal cost. We do this by estimating changes in sales given exogenous changes in price, and controls for competitors' costs, for each unique product$\times$market$\times$year in our data. Our empirical strategy is motivated by theory developed in Baker and Bresnahan (1988), which provides us a framework to interpret our estimates as residual demand elasticities.

Our analysis yields two main findings. First, we find that the distribution of demand elasticities across all goods measured in the Nielsen data has remained stable over the 2006-2020 period. Additionally, we find that within each good-market, the distribution of time trends in demand elasticities is centered at zero and has limited spread. While the mean sales-weighted residual demand elasticity has trended slightly more inelastic over this period, this is primarily driven by changes in demand. Compositional changes from product entry and exit roughly cancel out.

Second, we observe substantial spatial heterogeneity within products across markets. Differences in market characteristics drive this variation: Larger markets have tougher competition, while concentration is associated with a less competitive environment locally, but not nationally.

Our findings are at odds with previous work that has suggested an increasing trend towards inelastic demand elasticities in the retail scanner data, particularly for certain product categories. Our results are more consistent with the theory of contestable markets. Overall, our study provides a comprehensive examination of demand elasticities using a large, multi-year dataset. Our findings have implications for the measurement of markups and the assessment of consumer welfare, particularly in light of the increasing concentration of product markets and rising corporate profits.

References

Amiti, M., O. Itskhoki, and J. Konings (2019). International shocks, variable markups, and domestic prices. *Review of Economic Studies 86*(6), 2356–2402.

Anderson, E., S. Rebelo, and A. Wong (2020). Markups across space and time. *Working Paper*.

Baker, J. and T. Bresnahan (1988). Estimating the residual demand curve facing a single firm. *International Journal of Industrial Organization 6*, 283–300.

Baumol, W. (1983). Contestable markets: an uprising in the theory of market structure. *American Economic Review 73*, 491–96.

Beraja, M., E. Hurst, and J. Ospina (2019). The aggregate implications of regional business cycles. *Econometrica 87*(6), 1789–1833.

Berry, S., M. Gaynor, and F. Scott Morton (2019). Do increasing markups matter? Lessons from empirical industrial organization. *Journal of Economic Perspectives 33*(3).

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4).

Brand, J. (2021). Differences in differentiation: Rising variety and markups in retail food stores. *Working Paper*.

Bresnahan, T. F. (1981). Duopoly models with consistent conjectures. *American Economic Review 71*(5), 934–945.

Broda, C. and D. Weinstein (2006). Globalization and the gains from variety. *Quarterly Journal of Economics 121*(2), 541–585.

Chandra, A., A. Finkelstein, A. Sacarny, and C. Syverson (2016). Health care exceptionalism? Performance and allocation in the US health care sector. *American Economic Review 106*(8).

Chetty, R., J. Friedman, and J. Rockoff (2014). Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review 104*(9).

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics 135*(2).

De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review 102*(6), 2437–2471.

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in US retail chains. *Quarterly Journal of Economics 134*, 2011–2084.

Döpper, H., A. MacKay, N. Miller, and J. Stiebale (2021). Rising markups and the role of consumer preferences. *Working Paper*.

Franco, S. (2024). Output market power and spatial misallocation. *Working Paper*.

Ganapati, S. and R. McKibbin (2023). Markups and fixed costs in generic and off-patent pharmaceutical markets. *Review of Economics and Statistics 105*(6), 1606–1614.

Grieco, P., C. Murry, and A. Yurukoglu (2023). The evolution of market power in the us automobile industry. *Quarterly Journal of Economics*, qjad047.

Hall, R. (1988). The relation between price and marginal cost in US industry. *Journal of Political Economy 96*(5), 921–947.

Hall, R. (2018). New evidence on the markup of prices over marginal costs and the role of megafirms in the US economy. Technical report, National Bureau of Economic Research.

Handbury, J. (Forthcoming). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Econometrica*.

Handbury, J. and D. Weinstein (2010). Goods prices and availability in cities. *Review of Economic Studies 82*(1), 258–296.

Hausman, J. (1996). Valuation of new goods under perfect and imperfect competition. *The Economics of New Goods*.

Hitsch, G., A. Hortaçsu, and X. Lin (2021). Prices and promotions in US retail markets. *Quantitative Marketing and Economics 19*(3), 289–368.

Kane, T. and D. Staiger (2008). Estimating teacher impacts on student achievement: An experimental evaluation. *NBER Working Paper 14607*.

Knittel, C. and K. Metaxoglou (2014). Estimation of random-coefficient demand models: Two empiricists' perspective. *Review of Economics and Statistics 96*(1).

Lindh, T. (1992). The inconsistency of consistent conjectures: Coming back to cournot. *Journal of Economic Behavior & Organization 18*(1), 69–90.

Makowski, L. (1987). Are rational conjectures' rational? *Journal of Industrial Economics*, 35–47.

Melitz, M. and S. Polanec (2015). Dynamic olley-pakes productivity decomposition with entry and exit. *RAND Journal of Economics 46*(2), 362–375.

Morris, C. (1983). Parametric empirical bayes inference: Theory and applications. *Journal of the American Statistical Association 78*(381), 47–55.

Rossi-Hansberg, E., P. Sarte, and N. Trachter (2020). Diverging trends in national and local concentration. *NBER Macroeconomics Annual 2020 35*, 115–150.

Tran, U. (2021). Demand elasticities and the rise of broadband: Evidence from over 500 product categories. *Working Paper*.

Waldfogel, J. (2003). Preference externalities: An empirical study of who benefits whom in differentiated-product markets. *RAND Journal of Economics 34*(3), 557–568.

Year	No. Unique Brands	Total Sales (millions USD)
2006	24,610	50733.11
2007	27,834	60277.87
2008	28,559	65420.67
2009	28,254	68558.52
2010	28,358	70524.2
2011	29,059	75537.26
2012	28,463	74974.56
2013	28,426	75481.03
2014	28,883	76553.5
2015	29,316	80115.02
2016	29,239	81740.25
2017	28,194	78601.84
2018	40,398	102188.8
2019	40,245	102554.2
2020	37,291	106569.3

Table A1: Nielsen coverage over time