Uniform Pricing in US Retail Chains

Stefano DellaVigna, *UC Berkeley and NBER**
Matthew Gentzkow, *Stanford University and NBER*

September 16, 2017

Abstract

We show that most US grocery, drug, and mass-merchandise chains charge nearly-uniform prices across stores, despite wide variation in consumer demographics and the level of competition. Estimating a model of consumer demand reveals substantial within-chain variation in price elasticities and suggests that chains sacrifice 3-10 percent of variable profits relative to a benchmark of flexible prices. In contrast, differences in average prices between chains broadly conform to the predictions of the model. We show that the uniform pricing we document dampens the overall response of prices to local economic shocks, shifts the incidence of taxes and intra-national trade costs, and significantly increases the prices paid by poorer households relative to the rich. We discuss fixed costs of managerial decision making, tacit collusion, and fairness concerns as possible explanations for near-uniform pricing.

*E-mail: sdellavi@econ.berkeley.edu, gentzkow@stanford.edu. We thank Nicholas Bloom, Liran Einav, Benjamin Handel, Ali Hortacsu, Kei Kawai, Peter Rossi, Stephen Seiler, Steven Tadelis, Sofia Villas-Boas and seminar participants at Stanford GSB, UC Berkeley, UCLA, the University of Bonn, and the University of Chicago Booth, and at the 2017 Berkeley-Paris conference in Organizational Economics for helpful comments. We thank Angie Acquatella, Sahil Chinoy, Bryan Chu, Johannes Hermle, Ammar Mahran, Akshay Rao, Sebastian Schaube, Avner Shlain, Patricia Sun, and Brian Wheaton for outstanding research assistance.

1 Introduction

Recent research across several domains highlights the importance of retail price adjustment to local shocks. Beraja, Hurst, and Ospina (2016) and Stroebel and Vavra (2014) show that local retail prices increase in response to positive shocks to consumer demand, and argue that such price responses have important implications for understanding business cycles. Atkin and Donaldson (2015) show that retail prices are higher in more remote areas due to intra-national trade costs, and that consumers in these areas benefit less from globalization as a result. Jaravel (2016) shows that local prices have fallen more in high-income areas, possibly due to higher rates of product innovation, and that this has significantly exacerbated rising inequality. In interpreting the data, authors in these areas typically start from models in which local prices are set optimally in response to local costs and demand.

In this paper, we show that most large US grocery, drugstore, and mass-merchandise chains in fact set uniform or nearly-uniform prices across their stores. This fact echoes uniform pricing "puzzles" in other domains, such as movie tickets (Orbach and Einav, 2007), sports tickets (Zhu, 2014), rental cars (Cho and Rust, 2010), and online music (Shiller and Waldfogel, 2011), but is distinct in that prices are held fixed across separate markets, rather than across multiple goods sold in a single market. We show that limiting price discrimination in this way costs firms significant short-term profits. We then show that the result of nearly-uniform pricing is a significant dampening of price adjustment, and that this has important implications for the pass-through of local shocks, the incidence of trade costs, and the extent of inequality.

Our main analysis is based on store-level scanner data for 9,415 grocery stores, 9,977 drugstores, and 3,288 mass-merchandise stores from the Nielsen-Kilts retail panel. While we observe no cases in which a chain charges identical prices for all products across stores, we find that the variation in prices within chains is much smaller than variation of prices between stores in different chains. Our benchmark measure of price distance, the average absolute difference in quarterly log prices across pairs of stores, equals 0.03 log points for stores within a chain, but 0.12 log points for stores across different chains. A second measure capturing the high-frequency similarity in prices paints a similar picture: the correlation of weekly prices is 0.84 for stores within a chain, but 0.13 for stores across chains. The coordination of prices within a chain occurs despite the fact that consumer demographics vary widely. For example, consumer income per capita ranges from $22,450 at the 25th-percentile store to $33,450 at the 75th-percentile store.When we restrict to comparing pairs of stores with substantially different consumer income (in the top third versus in the bottom third) we obtain very similar results. This within-chain price uniformity is also not explained by the need to coordinate price in order to issue similar coupons within a geographic area: the result are similar comparing pairs of stores in different DMAs. Finally, this pattern holds similarly for high-selling products and less popular products, for high unit-price (high-quality) products and average-price products, for brand-name products and for generics.

The extent of uniformity does vary across chains. Of the 64 food store chains in our data, 56 charge prices that we characterize as nearly uniform. Six chains vary prices across large regions, but charge nearly uniform prices within regions, in a pattern that we call *zone pricing*. And two chains vary prices more substantially at the store level. The four drugstore chains and five mass-merchandise chains in our sample generally practice zone pricing too. Chains with a high degree of similarity in average prices across stores also exhibit high correlation across stores in the timing and extent of sales.

To better understand the incentives chains face to vary prices, we examine the response of prices to a key determinant of consumer willingness to pay and price elasticity, the average per-capita income for consumers shopping at a store. We show that, within a chain, there is a very limited (if clearly statistically significant) response of prices to income: prices increase by 0.72 (s.e. 0.12) percent for each $10,000 increase in consumer income.

In contrast, we document a price response that is an order of magnitude larger to income differences *across* chains. Compare two chains, one of which operates stores with an average consumer income of $25,000, while the second chain operates stores with average income of $35,000. Do they charge similar prices? We estimate that, across food chains, a higher consumer income of $10,000 leads to an average price increase of 4.48 (s.e. 1.01) percent, an effect an order of magnitude larger than the within-chain pricing effect. We further consider a second between comparison, the comparison of pricing across states (but still within a chain). Across states, a higher income by $10,000 in a state is associated with 2.16 (s.e. 0.33) percent higher prices, an effect of a similar order of magnitude (if half the size) of the between-chain effect. This zone pricing effect is largest for the drug stores we consider and smallest for the mass merchandise stores.

Before we turn to quantifying these magnitudes in terms of a simple pricing model, we address a puzzle within the puzzle. Why do firms that charge largely uniform prices still respond to local income in their posted price, but with a very small response? We document that this response is most likely spurious and due to two compositional effects. The price in the Nielsen data (and most similar data sets) is computed as the ratio of weekly revenue to weekly units sold, computed Sunday to Saturday. But not all retailers change prices on the same day of week (Saturday), and furthermore not all consumers pay the same price, as some consumers do not have loyalty cards and thus always pay the full price. Both confounds contribute to a spurious income-price slope, as consumers in lower income areas are more likely to wait for sales and to use loyalty cards. Indeed, using the data set from a major grocer in Gopinath, Gourinchas, Hsieh, and Li (2011), we show that removing this two confounds leads to a completely flat price-income relationship. These results suggest that the standard practice of using the weekly ratio of revenue to units sold as a measure of price can bias the estimated price-income relationship both in the cross section and over time in response to shocks.

We then turn to provide a benchmark for the optimal price response by estimating a simple constant-elasticity model of local consumer demand. The model fits the data well. In particular, we document that the relationship between weekly log quantity and weekly log price is very close to linear, consistent with the constant-elasticity assumption. We also show that the store-level estimate of elasticity, $\hat{\eta}_s$, is both statistically quite precise and closely predicted by store-level measures of demographics and competition, and in particular by the income measure used above. Like demographics and competition, these elasticities vary widely within chains, ranging from -2.28 at the 10th percentile to -2.98 at the 90th percentile in food stores, -1.94 to -2.65 in drugstores, and -2.92 to -3.67 in mass-merchandise stores. Our model implies that the ratio of the optimal price to marginal cost for a store with elasticity η_s is $\eta_s/\left(1+\eta_s\right)$. Assuming no variation in marginal costs across stores, this implies that prices at stores with elasticities in the 90th percentile should be 18 percent higher than stores with elasticities in the 10th percentile in food stores, 29 percent higher in drugstores, and 11 percent higher in mass-merchandise stores, whereas observed prices are on average only 0.4 percent higher in food stores, 0.8 percent higher in drugstores, and 0.4 percent higher in mass-merchandise stores..

To more precisely characterize the response of prices to demand conditions, we regress log price on the $log\left[\eta_s/\left(1+\eta_s\right)\right]$ term suggested by the model, instrumenting for elasticity with income. Under the model the coefficient in this regression should be 1. The results show that the food chains respond to their average demographics roughly as the model would predict, with a coefficient of 0.94 (s.e. 0.22), compared to the model prediction of 1. The zone pricing response across the different zones is also sizable, at 0.35 (s.e. 0.19), if smaller than the statistical benchmark. In comparison, instead, the within-chain price response is an order of magnitude smaller, at 0.09 (s.e. 0.03), and in fact, even this small response is likely biased upward by the compositional biases outlined above. These results are robust to using a variety of alternative products (e.g, brand-name versus generics) to derive the price measures, and to using different instruments for the store-level elasticity. For the drugstores and the mass merchandise stores, we cannot do between-chain comparisons given that there are only 5 and 5 chains, respectively, in the sample. The other results are qualitatively parallel, with larger price response to elasticity. For the drug stores, the between-state zone pricing pattern is actually consistent with the model at 0.86 (s.e. 0.27). The within-chain response is, like for the food stores, still much smaller than the response predicted by the model, at 0.29 (s.e. 0.04).

Our model suggests that the average food chain could increase variable profits by 1.2 percent and total profits by 9.7 percent were they to price flexibly compared to if they set uniform prices corresponding to elasticity faced by the average store.

We consider a number of potential threats to the validity of our model predictions. First, our baseline model abstracts from variation in marginal costs across stores. Stroebel and Vavra (2014) present a range of evidence suggesting that such variation is likely to be small. Moreover, to the extent the marginal costs do vary, we would expect them to be positively correlated with income, meaning that our model would understate the variation in true optimal prices. Marginal cost variation would thus deepen rather than resolve the uniform pricing puzzle.

Second, our baseline estimates assume that short-run week-to-week elasticities are equal to long-run elasticities. The literature suggests that the long-run elasticities relevant to the store's problem could in fact be smaller (due to consumer stockpiling) or larger (due to search). We repeat our analysis using prices and quantities aggregated to the quarterly level and find that the qualitative results are unchanged. We also show that the results are similar for storable and non-storable products.

Third, our main analysis treats demand as separable across products. This dramatically simplifies our analysis, but it is clearly unrealistic. Cross-product substitution could lead us to over-state the relevant elasticities to the extent that consumers substitute among products, or under-state them to the extent that consumers substitute on the store-choice margin as in Thommasen et al. (2017). To address this concern, we present a version of our analysis where we estimate elasticities at the product category level rather than the individual product level and show that the results remain similar.

Finally, our results assume chains are free to vary prices across stores if they want to. In reality, prices and promotions are often determined jointly by retailers and manufacturers (Anderson et al., 2017). We show that our results are similar for store-brand and national-brand products, suggesting that constraints imposed by manufacturers are unlikely to be a key driver of our results.

What, then, explains the observed results? Menu costs (Mankiw, 1985) are unlikely to provide a convincing explanation, since stores change prices frequently over time. Another possible explanation is that committing to uniform or zone pricing benefits chains by allowing them to soften price competition. Dobson and Waterson (2008) present a model of this phenomenon, and Adams and Williams (2017) find mixed support for it using data from the hardware industry. In our context, though, we show that the extent of price uniformity is about the same among firms that operate stores mostly without competitors, for which the tacit collusion story should have less bite. We also discuss the possibility that the price uniformity may be due to a constraint posed by the advertising of coupons. This would force price uniformity within a DMA, the relevant advertising zone, but not between DMAs. Yet, we find no evidence of zone pricing at the DMA level, as opposed to at the state level, as we discussed above. Thus, these three explanations do not appear to help with the results.

Fairness concerns on the part of consumers may be more plausible. Such concerns are often cited as a possible explanation for uniform pricing across products such as movie tickets sold by a single seller (e.g., Orbach and Einav, 2007), and we show examples below of grocery chains citing fairness as a concern. However, several facts seem to us to argue against fairness as a main driver: few consumers would know if a grocery chain charged different prices at stores in different states, it is not obvious how fairness concerns would explain the zone pricing we see for several chains in our sample, and chains in other markets such as gasoline do vary prices substantially without provoking any fairness outcry.

An explanation we see more support for is managerial decision-making costs (Bloom and Van Reenen, 2007)[1]. Implementing more flexible pricing policies may impose costs such as up-front managerial effort in policy design or investments in more sophisticated information technology. Inertial managers may also perceive a cost in deviating from the traditional pricing approach in the industry. A stylized model of such costs would require the chain to pay fixed costs at the chain and/or store level to implement flexible pricing. We find some limited support for these costs in the data. We find no evidence of deviations from (near) uniform pricing for stores with more extreme elasticity (within a chain), suggesting that store-level fixed costs are not a key driver. There is, however, a weak positive association across chains between measures of the loss from uniform pricing—influenced by the number of stores and the variation in the demand elasticities they face—and the extent of uniform pricing, consistent with chain-level fixed costs playing a role.

In the final section of the paper, we turn to the implications of uniform pricing for the broader economy. We first show that uniform pricing exacerbates inequality, increasing prices posted to consumers in the poorest decile of zip codes by 9.4 percent relative to the prices posted to consumers in the richest decile for food stores, 8.0 percent in drugstores, and 3.9 percent in mass-merchandise stores. We then show that uniform pricing substantially dampens the response of prices to local demand shocks. This significantly shifts the incidence of these shocks – for example, exacerbating the negative effects of the great recession on markets with larger declines in housing values. We next show that it changes the incidence of trade costs, benefiting more remote areas that otherwise would pay significantly higher prices.

We are not the first to document uniform pricing policies in retailing. Reports from UK regulators show that roughly half of UK supermarket chains charge uniform prices across stores (Competition Commission, 2003, 2005) as do the main UK electronics retailers (MMC, 1997a,b). Retailers such as IKEA are known to honor online prices in their brick and mortar stores (Dobson and Waterson, 2008). Cavallo, Nieman, and Rigobon (2014) show that Apple, IKEA, H&M, and Zara charge nearly uniform prices across the Euro zone in their online stores, though they charge different (real) prices across countries with different currencies. Early studies of the Dominicks chain in the Chicago market including Hoch, Kim, Montgomery, and Rossi (1995), Montgomery (1997), and Chintagunta, Dubé, and Singh (2003) show that Dominicks varies prices across several zones defined "almost entirely by

[1] A different version of this explanation is that managers are simply unaware of the income differences across their stores, or that they lack the information to recognize its implications for optimal prices. This seems unlikely to us. There is also no evidence of firm learning, as the observed patterns are at least as strong in the most recent three years of data than in the first three years.

the extent of local competition," holding prices constant across stores within a zone, and estimate the potential gains to more flexible pricing.

Recent work by Adams and Williams (2017) is particularly closely related. These authors show that the Home Depot and Lowe's US hardware chains use a zone pricing strategy, with different degrees of price flexibility for different products. They then estimate a structural model of demand and oligopoly pricing for a single product, drywall, and use it to evaluate how profits would change under more flexible pricing for this product. Their analysis differs from ours in focusing on cost differences across stores as a source of variation in optimal prices—a factor that turns out to be important in their setting—but ruling out differences in price sensitivity of consumers across markets—a factor that plays a key role in ours.

To the best of our knowledge, our study is the first to highlight the extent of uniform pricing across a large set of US grocery, drug, and mass-merchandise chains, to compare the observed prices in such chains to a benchmark of optimal pricing outside of the Chicago market, and to address the broader economic implications of dampened response to local demand and cost differences.

Our paper relates to a broad range of work studying the extent and implications of local retail price variation. Examples beyond those cited above include Broda and Weinstein (2008), Gopinath, Gourinchas, Hsieh, and Li (2011), Coibion, Gorodnichenko, and Hong (2015), Fitzgerald and Nicolini (2014), Handbury and Weinstein (2015), and Kaplan and Menzio (2015). Our work also speaks to the wider literature tracing out the implications of retail firms' price setting strategies for macroeconomic outcomes. This includes influential early work using scanner data by Bils and Klenow (2004) and Nakamura and Steinsson (2008), as well as recent contributions such as Anderson et al. (2017).

More broadly, our paper also relates to the work in behavioral industrial organization (for a review, see Heidhues and Koszegi, 2018). Most of the work in this area in the last 15 years has focused on firms optimally responding to behavioral consumers (DellaVigna and Malmendier, 2004; Gabaix and Laibson, 2006). Our paper is part of a smaller literature which considers instead *behavioral firms*, that is, instances and ways in which firms deviate substantially from profit maximization (Goldfarb and Xiao, 2011; Romer, 2006; Massey and Thaler, 2013; Hortaçsu and Puller, 2008; Hortaçsu, Luco, Puller, and Zhu, 2017; Ellison, Snyder, and Zhang, 2017). Among the patterns documented so far, there appears to be wide variation in the firm's ability to maximized profits, itself tied to managerial ability (Bloom and Van Reenen, 2007). Firms also appear to respond well to some variables, but largely neglect other determinants of profitability (Hanna, Mullainathan, and Schwartzstein, 2014).

2 Data

The main data source is the Nielsen data at the Kilts center. The data consists of two components: the retail scanner panel (RMS) and the consumer panel (Homescan, HMS). The retailer scanner panel records the average weekly revenue and quantity sold for over 35,000 stores in the US over the 2006-2014 period, covering about a million different UPCs (unique product identifiers). We use this data set to extract the information on weekly price and quantity. We also use some information from the consumer panel which is based on following the purchase of 60,000+ consumers across different stores. We present the main information in this Data section, with additional detail in the Appendix.

Store Sample. We focus the analysis on grocery stores ("food stores" in the Nielsen categorization), drugstores, and mass-merchandise stores. As Panel A of Table 1 shows, this initial sample includes 38,539 stores for a total average yearly revenue (as recorded in the RMS data) of $224 billion.[2]

For our analysis, the definition of chain is important, since we focus on comparisons of prices set within a chain, as well as between chains. The Nielsen data has two chain identifiers: a 'parent_code' variable and a 'retailer_code' variable. A company with a given 'parent_code' may have 2-3 different 'retailer_code's for its different stores, presumably capturing different store formats, or independent chains (or subchains) operated under common ownership.[3] Sometimes, a single 'retailer_code' appears under multiple 'parent_codes' even though, according to Nielsen, each 'retailer_code' should belong to a single 'parent_code.'[4] In light of this, we define a chain as a unique combination of a 'parent_code' and 'retailer_code', to ensure that we are focusing on a single decision-making unit. Further, as we detail below, we introduce additional restrictions to ensure proper identification of the chains.

First, we introduce restrictions at the store level. We exclude stores that switch chains over time to ensure that we assign a store correctly to its chain.[5] We also exclude stores that are in the sample for fewer than 104 weeks (i.e., two years worth of data), to ensure proper estimation of price elasticities.[6] We further eliminate a small number of stores without any consumer purchases in the Homescan data, introduced below. This reduces the sample to 22,985 stores in 113 chains.

We then introduce restrictions at the chain level. To start with, we require that the chains are present in the sample for at least 8 of the 9 years; this eliminates a few chains with typically only a small number of stores each with inconsistent presence in the data. Next, we eliminate chains for which we are not sufficiently confident about the grouping of stores. A first concern occurs when the same retailer_code identifier appears for stores with different parent_codes. It is unclear whether the use of the same retailer_code in this case indicates that these stores belong to one chain, or perhaps

[2] This number understates the total revenue for these stores since some products, like drugs sold on prescription, are not recorded in the data.

[3] Nielsen assigns these codes to mask the identities of chains, which are not to be identified.

[4] This may occur because retailer_codes are assigned based on what HMS panelists report, while parent_codes are automatically assigned by Nielsen to stores from the same reporting unit. Hence, while parent_codes should be reliable, retailer_codes are subject to misreporting.

[5] In some cases, we can validate the ownership change with significant observed changes in prices in the switching stores. However, some of the pricing changes occur up to 2 years in advance, or 2 years after, the change, suggesting a possible inaccurate record of the timing of ownership changes. Hence our conservative rule of excluding such switches.

[6] This sample restriction is especially important since the price elasticities are computed including controls for 52 week-of-year indicators, requiring multiple observations per week-of-year.

they belong to a subchain that changed owner, or something else. Thus, for each retailer_code, we only keep the parent_code associated with the majority of its stores, and then further exclude cases in which this retailer_code-parent_code combination accounts for less than 80% of the stores with a given retailer_code. A second concern is for chains in which a number of stores switch chain, given that this may indicate a change in ownership of the entire chain. We thus exclude chains in which 60% or more of stores belonging to the retailer_code-parent_code change either parent_code or retailer_code in our sample.

This leaves us with our main sample of 22,680 stores from 73 different chains, covering a total of $191 billion of average yearly revenue. These include 9,415 stores from 64 food store chains ($136 billion average yearly revenue), 9,977 stores from 4 drugstore chains ($21 billion), and 3,288 stores from 5 mass-merchandise chains ($34 billion). These numbers underscore the difference betwen the food stores on the one hand and the drug and mass merchandise stores on the other hand. The food store industry is less concentrated, with dozens of different chains: the median food chain (Panel C of Table 1) has 66 stores, and has locations in 4 DMAs and 2.5 states. Drugstore and mass-merchandise chains are significantly larger and span significantly more states, as Panels D and E of Table 1 show. Indeed, 99 percent of the drug stores in the sample belong to just 2 chains. Thus, when we compare prices *across* chains (as opposed to *within* chains) we will focus just on food stores, since there are too few chains to compare in the drugstore or masss merchandise sector.

Taken together, these stores cover the entire continental US, as the map in Appendix Figure 1 displays. The number of stores and chains in the sample remains fairly constant between 2007 and 2013 (Online Appendix Table 1), with a somewhat smaller number of stores in 2006, 2007, and 2014.

Store Characteristics. To define the demographics of the stores, we use the Homescan data, which includes all shopping trips for the consumers in the Nielsen HMS panel. The median store has 21 Nielsen consumers ever purchasing at the store who make a total of 502 trips (Panel B of Table 1). We use demographic variables like income and education from the 2008-2012 5-year ACS for the 5-digit zip code of residence of the consumers shopping in each store, and then compute the weighted average across the consumers, weighting by the number of trips that they take to the store.[7]

As Panel B of Table 1 shows, the median store has an average per-capita income of $26,900, with sizable variation across stores; for example, the 75th percentile is at $33,750. The per-capita income at the store level is our main proxy for demand elasticity. As a secondary measure, we also use the share of the population over 25 years old with at least a bachelor degree, which is 17.8% at our median store.

Products. We focus most of our analysis on a subset of products that are available (commonly

[7]This demographic information is more accurate than the one that can be computed directly from the location of the store in the RMS data, since in this dataset the most precise geographic location given is the county or 3-digit zip code. Weighting by total dollar amount spent or using the unweighted average does not meaningfully change our imputed demographics.

sold) in many stores and chains. As we discuss below, sparsely-sold products introduce a potential bias in the price measure, and comparability of prices across stores and chains is a hallmark of our analysis. Specifically, for food stores we identify 10 product categories ('modules' in the Nielsen classification) which have high sales and cover both storable and non-storable products (Panel D of Table 1): canned soup, cat food, chocolate, coffee, cookies, soda, bleach, and toilet paper among the storable products, and yogurt and orange juice as the non-storable products.[8] These modules account for an average yearly revenue of $13.7 billion across the 9,415 stores in our food store sample, that is, 10.1% of total revenue in these stores.

By store type, within each module, we identify a product (UPC code) with high sales and high availability across our sample of stores. For food stores, in 7 of the 10 modules, this product remains the same over the 9 years in our sample, while in 3 other modules the product changes at least once over the 9 years. In all cases, within a given year, the product is the same for 9,415 the stores in all 64 chains, making it possible to compare prices not only within chain, but also across chains.[9] Examples of products are a 12-can package of Coca-Cola, a single 14.5 oz can of Campbell's Cream of Mushroom soup, and a 59 oz. bottle of Simply Orange juice.[10]

For the drug and merchandise stores, we build the set of products in a comparable fashion as for grocery stores, but focus on a subset of the 10 modules, since some products are not commonly sold in drug and merchandise stores. For drugstores, only 2 of the 10 modules have a product with high availability (90 percent or higher) across all the stores: soda and chocolate. For merchandise stores we are able to include 5 modules: soda, chocolate, cookies, bleach, and toilet paper.

For our supplemental analysis we also construct measures of prices for a less commonly sold item (the 20th highest-availability product across stores), for a high-quality (defined as high-unit-price) product, the top-selling generic product within each chain, and a subset of generic products comparable across chains. We do this for food stores only because of availability issues.[11] We also build a price index, as we detail below.

Prices. The price measure for a given store s, product j, and week t, P_{sjt}, is the ratio of the weekly revenue to the weekly number of units sold. The price is thus not defined if no units are sold in a UPC-store-week. This is the reason for focusing on products that are sold frequently, to minimize the occurrence of cases with missing price.[12] We denote with P_{sjt} the price level and with p_{sjt} the (natural) log of the price.

[8] These modules have a large overlap with ones used in previous analyses, e.g., Hoch, Kim, Montgomery, and Rossi (1995)

[9] This requirement of between-chain comparison made it difficult to include more nonstorable products, as for example milk and egg UPCs will typically differ across geographic areas.

[10] Notice that to maximize comparability we do not aggregate across different flavors or varieties of the same product. For example, our 59 oz. Simply Orange OJ consists of the regular pulp-free version only.

[11] The 20th highest selling products and high-quality items are the same across chains, but the chain-specific generic products vary across chains. For the sample of generic products comparable across chains, see the Appendix for details.

[12] Missing prices are an issue in particular because they do not occur at random: for any given item, a week with no items sold (and thus no recorded price) is more likely to occur when the price is higher, not when an item is on sale.

There are two important and potentially problematic features of this price measure. First, this price measure includes sales; there is no separate record for the regular price. It is possible that, when a product is on sale, the sale price applies only to the consumers who use loyalty cards. In this case, the price captures an average of the (sale) price paid by the consumers with loyalty card, and the (regular) price paid by the other consumers. Second, Nielsen records the weekly sale on a Sunday-to-Saturday timing. Yet, the within-week timing of price changes in some retailers may not coincide with the timing in the Nielsen data. (Indeed, we document below the case of a retailer where it does not.) In this case, the Nielsen price for a given week represents the average of the price charged in two contiguous weeks, weighted by the number of units that were purchased in each (part) week. We return below to the potential biases introduced in the estimate of price setting by these features.

Table 2 displays key features of the products chosen. The average price varies from \$0.49 for cat food in food stores to \$8.70 for toilet paper in mass merchandise stores (Column 3). The products have at least one recorded sale in the large majority of store-weeks, for example in 99.7% of store-week-UPC observations for chocolate in food stores (Column 4). Cat food, coffee, and toilet paper have somewhat lower availability in food stores, as do most of the products sold in drugstores and mass-merchandise stores, but are still in the range around or above 95%. We also compute the average yearly revenue per store that these products generate, with the highest number associated with the soda product in food stores, \$34,100.

Price Indices. Our benchmark set of results on pricing focuses on the products in Table 2 because of two key advantages: (i) these items are available in nearly all stores, such that we can compare prices not only within chains, but also across chains; and (ii) the number of items sold in a given week is quite high, so it is rare to have an entire week with no sale (and thus no price recorded); this minimizes the occurrence of missing prices. A disadvantage of focusing on these products is that they represent a small, and potentially not representative, basket of the goods sold in grocery stores, drugstores, and mass-merchandise stores.

Thus, for the food stores we also build a price series built with a larger basket of products within the 10 modules we identified. The basket is constructed with an eye to minimizing the occurrence of missing prices. Within each module, we select all products (UPCs) that have at least 95% average availability in a given chain-year. For some modules such as soda and orange juice, products meeting this criterion cover 50-60% of the total module revenue, while for other modules with more horizontal differentiation, like chocolate or coffee, they cover just 15-20%. (Panel B in Online Appendix Table 1). Summing over the 10 modules, these products cover an average annual per-store revenue of \$670,000, summing to over \$6bn annually over the 9,415 stores in our sample. This revenue coverage is an order of magnitude larger than for our benchmark products in Table 2.

To compute the price and quantity index for store s we start from the weekly log price p_{sjt} and weekly log quantity (units sold) $q_{sjt} = log(Q_{sjt})$. To compute the index variables $\bar{p}_{st}$ and $\bar{q}_{st}$, we average p_{sjt} and q_{sjt} across all products j included in the index basket for that module-chain-year. As weights, we use the total quantity sold for product j in a chain-year. If a product j has no sales in a particular store s and week t, in which case there is no recorded price, then product j is omitted in the computation of the index variables $\bar{p}_{st}$ and $\bar{q}_{st}$ for the particular store-week cell, and the other weights are scaled up for that store-week.[13]

Secondary Data Source. A drawback of the Nielsen RMS data, as we mentioned above, is that the price variable does not allow us to separate the sale price from the non-sale price, to control for the share of store card users, or to ensure that the weekly timing of prices as recorded by Nielsen matches the weekly timing of sales within a chain. In addition, the Nielsen data set has no information on costs.

To investigate these issues, we make use an an additional data set, also from scanner data, from a single major grocer ('parent_code', by Nielsen's definition). This data set, from Gopinath, Gourinchas, Hsieh, and Li (2011), contains data from 250 grocery stores belonging to twelve chains ('retailer_codes') located in the USA (as well as 75 stores in Canada) beginning in 2004 and ending in mid-2007. We select the largest retailer, which has 134 USA stores in the dataset. We match 133 of these 134 stores to stores in the Nielsen data set[14].

The data set reports the weekly revenue and quantity sold at the UPC level, just like in the Nielsen data set. However, unlike in the Nielsen data, the weekly timing at which revenue and quantity sold are reported corresponds exactly to the weekly timing with which the chain changes the prices. Furthermore, in addition to units sold and expenditure, this data set also contains "gross amount," which is weekly expenditure if all units were purchased at the non-sale price, wholesale prices, and adjusted gross profits for the UPC-week.

3 Descriptive evidence

3.1 Pricing Examples

We start by providing some examples of how pricing varies across stores and over time, within a chain. Figure 1a visualizes the pricing of a particular food chain (chain 128) for the orange juice product. Each of the 250 rows in the figure corresponds to a different store in the chain, with the stores sorted by local income. Each of the columns corresponds to one of the 52*9=468 weeks from

[13]We use the same weights for the price variable and the quantity variable so that, under the assumption that all products within a module have a constant-elasticity demand with the same elasticity η, we can recover the elasticity η regressing the index quantity $\bar{q}_{st}$on the index price$\bar{p}_{st}$. We use quantity weights so that our price index resembles a geometric modified Laspeyres Index, similar for example to Beraja, Hurst, and Ospina (2016) and to how the BLS builds category-level price indices. Note that our index is not exactly a Geometric Laspeyres Price Index because the weights are not week 1 weights but instead the average quantities sold in year y.

[14]See Appendix Section A.1.7 for details on how we matched stores.

January 2006 to December 2014.[15] For each store-week combination, we color-code the price in that store-week observation. More precisely, the price variable for store s, product j, and week t is a demeaned log price, $log(P_{sjt}) - log(\overline{P_j})$, where $\overline{P}_j$ is the average price for product j across weeks t in a year and across all stores s in all chains. The color coding, as indicated at the bottom of the figure, displays higher prices as darker.[16]

This visualization allows one to compare the price variation across stores to the price variation over time. Moving across columns within a row, one sees regular-price weeks followed by sale weeks at relatively consistent frequencies. The variation in the price from week to week is often as large as 30 log points, that is, sales of over 30 percent are not uncommon.

Moving across rows within a column, one can instead assess the phenomenon we focus on, the within-chain price variation. As the figure shows, there is no visible price variation across stores within a given week. In particular, there is no evidence of higher prices charged in the stores in higher income areas. This is noteworthy as this chain operates in areas with diverse per-capita income ranging from about \$13,000 to about \$50,000. Thus, this chain has an elaborate pricing scheme of sales over time, but essentially rigid prices across stores. We label this pattern *uniform pricing*.

Of course, this particular pattern may be peculiar to the displayed product. Figure 1b thus displays the same pattern, but for five different products: cat food, cookies, soda, chocolate, and yogurt. In order to display all five products, we display just 50 of the 250 stores shown in Figure 1a, with the same 50 stores shown for the 5 products, and still ordered by per-capita income within products. The white regions indicate weeks in which zero units of the product is sold in that store, and thus the price is missing.

Figure 1b shows within each product the same pattern which was visible for the orange juice product: substantial variation over time, with regular price weeks followed by sale weeks; and yet, no visible variation in price across the different stores, with essential uniform pricing across the different stores in any particular week. The figure also shows that the pattern of sales differs across the different products, with different sale cycles and different intensity of sales. That is, the chain follows a seemingly complex model of sales set product-by-product, which is then essentially applied to all the stores in the chain.

Figures 1a and 1b only display pricing patterns for one chain. Is this pattern representative? Chain 128 is indeed quite representative of the patterns of pricing for the majority of chains, with accentuated sales over time, applied quite uniformly across the stores. For example, Online Appendix Figures 1a and 1b display the equivalent of Figure 1b for two other chains, showing patterns echoing

[15] We drop the 53rd week in 2011.

[16] This unit of pricing is chosen to display not only variation in prices across stores and over time, but also the absolute level of prices in a chain. For example, a chain with darker colors corresponding to prices in the range 0.1-0.2 would indicate a chain that charges on average 10 to 20 log points higher prices than other chains for the orange juice product.

the ones for Chain 128.

While patterns like these are typical of the majority of chains, a few other chains follow a somewhat different pattern, which we label *zone pricing*. Figure 2a displays an example for chain 130, returning to the orange juice product. Figure 2a follows the same design as Figure 1a, except that we group stores approximately by geography by sorting them by 3-digit zip codes within states. This chain operates in 12 different states.

Unlike for chains with *uniform pricing*, Figure 2a shows prices that are essentially uniform within horizontal bands, but then differ for different bands. For example, stores in Georgia and Kentucky share the same pricing patterns, with little to no difference in prices across stores in these states. But the price differs quite a bit in Illinois and most of Indiana, with sales that differ in both timing and intensity.

The presence of different pricing zones is widely mentioned in the literature (for example, Hoch, Kim, Montgomery, and Rossi (1995) and Montgomery (1997)) and appears for other chains in the sample (Online Appendix Figures 1c-d). Still, for the majority of chains in our sample we do not find obvious pricing zones, with pricing that rather resembles the one of *uniform pricing* in Figures 1a-b.

Two chains do not conform easily into either of the two patterns above, displaying patterns of pricing as in Figure 3. While there still are clear patterns of similarity across stores in certain bands, other stores appear to follow different rules as far as price levels and sales. We call the pricing in such chains *individualized pricing.* The label should be taken with caution: it emphasizes the presence of differential pricing by store, but one should also keep in mind that there is still remarkable similarity in the pricing across stores, as Figure 3 shows.

3.2 Measures of Pricing Similarity

In the previous section, we visualized the pricing for some chains. To provide more systematic evidence, we introduce three measures of the extent of uniform pricing.

The first main measure is the *quarterly absolute log price difference* between two stores. For store s and product j, we compute the average unweighted weekly log price in quarter q in that store, $\overline{p}_{sjq}$. We then compute for each pair of stores s and s' the absolute difference in this average log price, and average it across the quarters in the data, and across the 10 products in our main sample: $a_{s,s'} = \sum_{q,j} |\overline{p}_{sjq} - \overline{p}_{s'jq}|/N_q N_j$, where N_q and N_j denote the number of non-missing quarters and products, respectively.

The second main measure is the *weekly correlation in prices* between two stores. To compute this correlation, we first demean the log price p_{sjt} at the store-year-product level to obtain $\widetilde{p}_{sjq}$. Then we compute the correlation between the demeaned prices between two stores $\varrho_{s,s'} = correl(\widetilde{p}_{sjt}, \widetilde{p}_{s'jt})$. The correlation is taken over all weeks t and over all products j which are non-missing in both store

s and store s'.

The two measures capture, by design, two orthogonal aspects of price similarity across stores. The absolute difference measure focuses on the similarity of the level of prices at the quarterly level. The weekly correlation measure, instead, focuses on the correlation of changes in price due to, say, coinciding sales timing. Thus, the measures can easily diverge. Two stores with the same timing and depth of sales, but different regular prices would have high correlation but at the same time also high differences in absolute prices. Conversely, two stores with similar average prices at the quarterly level, but different timing of sales would have low absolute differences, but low correlation at the weekly level.

As a third, auxiliary measure, we also record the *share of (nearly) identical prices*, defined as price differences smaller than 1 percent: $|P_{sjt} - P_{s'jt}|/((P_{sjt} + P_{s'jt})/2) < .01$. We compute the share of such observations across all products j and weeks t. We also record the absolute difference in weekly prices, constructed as for our first benchmark measure, but using differences in prices week-by-week (as opposed to quarterly averages).

Figures 4a-b display the distribution of the two main pricing similarity measures. Each observation is a pair (s, s') of stores. In blue we display the distribution of the variables for the approximately 380,000 pairs of stores[17] within the same chain (We use a maximum of 200 stores per chain to form within-chain pairs to avoid excessively overweighting the largest chains). Figure 4a shows that the average absolute quarterly distance in prices in within-chain pairs is typically lower than 7 log points, or approximately 7 percent. Figure 4b indicates that the large majority of within-chain pairs have a weekly correlation in prices above 0.7, indicating a highly correlated price setting.

How unusual is such pattern of closeness? For comparison, we display in red the pricing similarity measure for pairs of stores belonging to different chains. The measures of closeness for such between-chain pairs are completely different: the absolute distance in quarterly prices is typically above 8 log points, and the between-store correlation in prices is typically below 0.2. The differences in distribution are so large that one can almost exactly identify stores as belonging to one chain, or to different chains, just by observing the measures of distance (Online Appendix Figures 5a-b). Panel A in Table 3 summarizes the closeness in pricing both within chain and between chains[18].

A concern in this comparison is that the between-chain pairs compare stores that are more different from each other geographically than the within-chain pairs. Online Appendix Figures 3a-b and Panel B in Table 3 display very similar patterns comparing only stores within the same DMA.

A complementary concern is that the similarity of pricing for the within-chain pairs may be driven by mechanical reasons: stores within a DMA are likely to share an advertising market, and chains may be forced to charge one price within an advertising market in order to implement a mailing policy. Additionally, the cost to a chain of charging relatively uniform prices may not be

[17] The exact count differs slightly due to our selection criteria for valid store-weeks and store-quarters.

[18] In this table we continue to use a maximum of 200 stores per chain for the within-chain means.

high if the stores serve populations with very similar demographics.

To address the concern about advertising, in Figures 4c-d and in Panel C of Table 3 we restrict the pairs of stores to pairs in which the two stores belong to different DMAs[19]. Furthermore, to ensure that the stores operate in different demographic segments, we require that one of the stores in the pair be in the bottom third of nationwide income, while the other store in the pair must be in the top third of nationwide income. Applying these restrictions has only a modest impact on the distribution of the closeness measures.

In Online Appendix Figure 4 we document similar results for the closeness of pricing of stores using two auxiliary measures: the share of (near) identical prices and weekly absolute log price difference. Particularly striking is the fact that for within-chain pairs the share of identical prices is almost always higher than 20 percent and often as high as 60 to 70 percent. For the between-chain pairs, instead, the incidence of identical prices is rarely above 20 percent of pairs.

Results By Chain. The evidence presented so far documents a remarkable degree of uniformity in within-chain pricing, certainly as compared to the between-chain similarity in prices. Still, this evidence considers all chains together. We now return to an analysis chain-by-chain, as in the motivating evidence in Figures 1-3.

Figure 5a plots the two benchmark measures of within-chain pricing similarity, averaged to the chain level[20]. Each dot in the scatterplot indicates the average similarity measure for pairs of store belonging to a particular chain. For example, chain 313 has especially coordinated pricing, with an average weekly correlation in prices of over 95 percent and an average absolute quarterly distance in prices of just 1 log point (1 percent). The figure shows that near-uniformity in prices is the rule, rather than the exception: out of 73 chains, 58 chains have both an average correlation of weekly prices above 80 percent and an absolute quarterly distance in prices below 4 percent. The remaining chains have lower similarity measures, though all but three have an average correlation coefficient above 0.6 and and absolute distance below 7 log points.

Also interestingly, the two measures of pricing similarity are highly correlated: chains that are similar in one dimension are also similar in the other dimension. We reiterate that this is not mechanical, as the two measures are built to capture different dimension of pricing similarity: similarity in high-frequency sales versus similarity in low-frequency price levels. No chain appears to offer correlated timing of sales, but set at different levels of regular prices. The closest is chain 839, with average correlation 0.9 and absolute quarterly distance equal to 5 log points.

We now return to the *zone pricing* distinction seen above: some chains appear to charge largely uniform prices within a geographical region, but then differentiate prices across broad geographical regions. In particular, the geographical regions largely seems to consist of a state, or pairs of states. Thus, we decompose the measures of pricing similarity into similarity for pairs of stores within a

[19]Requiring store-pairs to be in different states as opposed to DMAs does not change the results qualitatively.
[20]For computational reasons, we use a maximum of 400 stores per chain

state, versus across state boundaries. For the subset of chains that operates in multiple states[21], we plot the within-state absolute quarterly difference in prices on the x axis, and the between-state absolute quarterly difference on the y axis. In most chains, indicated with the empty circles, there is essentially no difference between the two measures, suggesting that these chains largely charge uniform prices. Eight chains, however, stand out for having a larger difference in prices across state borders than within state. We take this as evidence indicative of *zone pricing* at the state level.[22]

In Online Appendix Figure 6, we repeat the same decomposition of within-state versus between-state pairs, but for the measure of weekly price correlation. The plot identifies broadly similar chains as having zone pricing. Two chains, however, that appear to charge zone pricing according to the absolute distance measure do not appear to do so according to the correlation measure.

Robustness. The results on pricing similarity so far are for a single high-selling item within each of 10 modules. It is possible that the patterns we find for such products may not apply to other items. Furthermore, the patterns may be related to the vertical negotiations between brands and the retailers, as there is evidence that the brands partly negotiate the pricing of the sales over time (Anderson et al., 2017).

We thus estimate the pricing closeness in food stores for (i) the 20th-highest-selling product within a category instead of the top-seller; (ii) the top-selling generic product, as opposed to the top-selling branded product, and (iii) a high-quality (high unit-price) product, compared to the lower-quality (but higher-selling) main product. The lower panels of Table 3 and in Online Appendix Figures 7a-c show that the patterns of pricing are quite similar for these alternative products, especially for the quarterly absolute price distance measure. Thus, the results are not due to unique patterns for the products we pick. The results are also similar for storable and non-storable items.

3.3 Price Response to Local Demographics

So far, we have seen that most chains have limited price variation. We now examine whether this price variation, limited as it is, is related to demographic proxies of local purchasing power, such as income. We expect stores in higher income areas to have more inelastic consumers (as we show below) and thus charge higher prices.

For each store s, we compute the average zip-code-level income y_s for consumers shopping in the store (see Section 2). We compute an average price for store s as follows: for each module j, let $\bar{p}_{jy}$ be the unweighted mean log price of products in module j and year y over all available store-weeks. Then, for each store-module-week, we calculate the demeaned weekly log price $\tilde{p}_{sjt} = p_{sjt} - \bar{p}_{jy}$, $t \in y$. We first average across weeks within each module for each store to get a store-module price

[21] To be included in this plot, we require that the chain operates at least 3 stores in each of 2 (or more) states.

[22] This method cannot identify pricing zones within a state, however, so it possibly understates the presence of pricing zones.

level $\tilde{p}_{sj} = \frac{1}{T_{s,j}} \sum_t \tilde{p}_{sjt}$[23] and then average across modules within each store to weigh each module equally to arrive at the average price level in store s $\overline{p}_s = \frac{1}{J} \sum_j \tilde{p}_{sj}$.

In Figure 6a we relate the store-level price variable $\overline{p}_s$ to the store-level income y_s within a chain. We demean both variables by the chain average and plot a bin scatterplot of the demeaned variables. The within-chain price-income relationship, while clearly statistically significant, is very flat economically: an increase in per-capita income of $10,000, equivalent to a move from the 30th to the 75th percentile, increases prices on average by only 0.72 percent. This very flat relationship is surprising in particular since higher income is likely to be associated not just with shifts in consumer elasticity, but quite plausibly also with higher costs, which should contribute to a steeper relationship as well.

A possibility is that this relationship is due to a small number of chains responding substantially to income, with no response from the other chains. In Online Appendix Figures 8a-c we display the estimated slope chain-by-chain[24]. The majority of chains have small, positive coefficients in the range between 0 and 0.01, with 27 coefficients positive and significantly different from zero. Only five chains instead have coefficients above 0.01, which itself is a fairly small effect (1 percent increase in prices for each $10,000 in income). Thus, the overall effect reflects a pattern happening within most chains, rather than a heterogeneous pattern across chains. Online Appendix Figure 9a also shows (for the food stores) that the pattern is quite similar within each of the modules and other pooled product types.

In Figure 6b we take a complementary approach to the within-chain analysis of Figure 6a, and estimate the between-chain relationship of prices and income (for food stores).[25] Namely, we relate the chain-level average price for all stores in the chain, and the chain-level average income. Interestingly, price and income are significantly related at the chain level, with a slope an order of magnitude larger than at the within-chain level: an increase of $10,000 in per-capita income at the chain level is associated with prices higher by 4.48 percent. While there are two outlier chains in terms of income, removing them does not affect the coefficient much (0.395 instead of 0.448), as Online Appendix Table 2 shows. Also, this sizable between-chain relationship holds for all modules but one, and holds for lower-selling products and for high-price (high-quality) products (Online Appendix Figure 9b).

What role does zone pricing play in these relationships? As we documented in Figures 2 and 5b, some chains have largely rigid pricing within a zone (typically a state), but then vary prices across zones. In Figure 6c, we re-estimate the within-chain relationship, but we demean the price and

[23] As we describe in the Prices subsection of Section 2, there is a bias in this price measure because missing prices are almost certainly nonsale prices.

[24] We exclude 2 chains with a very noisy estimate of the relationship (standard error above 0.02).

[25] We do not include the drugstores and mass merchandise stores since (i) we cannot compare across types of stores, given that the products are different and (ii) there are only 4 drugstore chains and only 5 mass merchandise chains, so the between-chain comparison for these groups of stores is not very informative. We nonetheless show it in Online Appendix Figure 10.

income variables at the chain-state level, thus focusing on within-zone pricing. This lowers further the slope of the price-income relationship from 0.0072 in Figure 6a to 0.0056 in Figure 6c, but it remains statistically significant. We return to this finding shortly.

In Figure 6d, we re-examine our between-chain analysis by considering between-zone pricing: we compute the average price and average income for each of the states in which the chains operate. We then demean the state-level observations at the chain level, so as to focus on the within-chain, but between-zones, pricing, and plot the chain-state observations in Figure 6d. Each point on the plot is a bin of chain-state observations: for example, chain 9 has 11 observations in the figure, corresponding to the 11 states it operates in. The between-zone analysis also provides evidence of a sizable price-income relationship: a $10,000 income increase is associated with an increase in prices of 2.16 percent, a slope about half the size as in the between-chain analysis but much larger than for the within-chain analysis. These findings are largely due to chains 9, 32, 4901, and 4904, the two food chains that operate in a number of states with wide differences in income as well as the two largest drugstore chains.

Indeed, In Figure 7a-d, we break out this zone pricing result by store type and additionally include a figure that includes only the six food chains that we label zone pricers (Figure 7b). The drugstore chains (Figure 7c) engage in more zone pricing than the other store types, but the pattern is present for the 3 types of chains we consider. We stress that this between-zone relationship is an additional, independent finding compared to the between-chain finding in Figure 6b, given that the observations in Figure 6d and Figure 7a-d are demeaned at the chain level.

In this section, we have focused on per-capita income as key demographic variable. In Online Appendix Figure 11 we show that the results are similar when estimated using instead as measure the fraction of college graduates, constructed in the same way.

3.4 Investigating the Within-Chain Response

Figures 6a and 6c raise a puzzle within the puzzle: why do chains that generally appear to charge rigid prices, at least within a pricing zone, exhibit a statistically significant, but extremely flat, response to income in their within-chain pricing? If they do intend to respond to demographics, we would expect them to do so by a larger amount, as in the between-chain relationship, or in the between-zones evidence. If they enforce rigid prices, it is unclear why there would be such clear statistical evidence of a positive relationship.

We now consider the possibility that this relationship may be due to aggregation biases in the price measure we use. As we discuss in Section 2, an averaging bias could be introduced by at least two factors: (i) differences between the weekly timing of price setting as enacted by the retailer, compared to the weekly timing with which prices are recorded in the Nielsen data; (ii) differences across stores in number of store-card users who have access to discounts or coupons.

The following example, tailored to the first case, highlights the confounds. Consider a retailer that changes price on Wednesday, but Nielsen records revenue and quantity on a Sunday-to-Saturday schedule. Suppose that the retailer introduces a discounted price p^{low} on Wednesday. Within the Nielsen-recorded week, the real price faced by consumers will be p^{high}on Sunday-Tuesday, and p^{low}on Wednesday to Saturday. In the Nielsen data, though, we only observe the weekly average, $p^{RMS} = sp^{high} + (1-s)p^{low}$, with s being the share of purchases made at high price. A first implication is that the Nielsen price will not reflect the actual price charged in either of the two weeks, but an average. That per se will not introduce bias in the analysis. But importantly, the share s of purchases made at high price is not equal to the share of time that the price is high, in this case 3/7. The share s of purchases at high price will be increasing in the income level at the store, as more inelastic consumers are less likely to chase discounts. This will introduce a bias that could produce exactly the observed within-chain relationship between average price and income. Furthermore, notice that a very similar bias arises if stores differ in the share of consumers with store cards, or coupons, and the discounted price applies only to them. The RMS observed price is a combination of the regular price and the discounted price, with the share of regular price being a function, presumably, of income.

To provide direct evidence on this biasing channel, we use the data from a major grocer used in Gopinath, Gourinchas, Hsieh, and Li (2011) and described in Section 2. This grocer does, indeed, change prices every week on Wednesday. If our model for the bias above is right, then the within-chain price-income relationship will be flatter when we use the prices from the grocer data, as opposed to the RMS data.

We match the stores in this grocer data to the Nielsen data so we can compare the two price series. Figure 8a shows a bin scatter of the within-chain relationship using the RMS price for the 133 stores in both data sets.[26] The slope is similar to the one in Figures 6a and 6c, though noisier given the small sample. Figure 8b reproduces the same exact estimate, but using the price level computed from the grocer data, which does not suffer from the day-of-week offset. The estimated slope is 0.19 percentage points flatter, from 0.27 percent to 0.08 percent, close to zero. This comparison deals with the day-of-week issue, but not with the discount card bias, which would apply to also to the grocer price. Since the grocer data also includes a non-sale price, we can directly test the impact of that. In Figure 8c, the relationship between the non-sale price and income is zero at all effects (0.02 percent). As an additional check about the importance of this bias, we use an algorithm on the Nielsen data to compute likely non-sale prices in our data, and repeat the analysis[27]. As Online Appendix Figures 12-14 show, this flattens the within-chain price-income relationship but not the between-chain relationship.

We conclude that it is very likely that, once one controls for pricing zones, pricing is completely

[26]Since this retailer is one that uses zone pricing, the relationship is demeaned by state.
[27]See Appendix Section A.1.8 for more details

rigid within a chain.

As a final use of this additional data, we take advantage of the availability of product-level cost information. Figure 8d plots a within-chain bin scatter of the wholesale cost variable for the Major Grocer against the store-level income proxy. The figure displays no evidence of a positive relationship between the two variables. This finding motivates our assumption later of constant marginal cost within a chain.

3.5 Joint Within-Between Evidence

So far, we have presented separate tests for the impact of demand determinants (like income) across stores within a chain (our 'within' evidence), across chains (our 'between' evidence), and across pricing zone (our 'between chain-state' evidence). We now present a unified test of of the three channels in Table 4. We regress for each store s the log price measure on both income for store s, the average income for all stores in the chain to which store s belongs, and the average income for all stores in a chain-state. To the extent that prices are rigid within chain, but are set at about the right level for the chain, as our previous within- and between- evidence suggest, then the chain average income will be a predictor of price setting, more so than the local income. Similarly, to the extent that there is zone pricing we expect that the state-level income will predict prices in a store, beyond the predictive power of income in a particular store. We consider separately the food stores (Panel A) from the drug and mass merchandise stores (Panels B and C), since it is only for the food stores that we can do a meaningful between-chain comparison.

We start from a naive specification that just regresses for food stores the price level in stores s to the income in store s (Column 1). This specification yields a signification relationship of 0.0175 (s.e. 0.0047). A similar relationship is sometimes estimated in studies that examine the impact of determinants of prices, like income. Column 2 shows that this association is almost entirely due to the chain-level income measure (coefficient of 0.0404), with the coefficient on own-store income reduced to just 0.0044. This confirms the earlier finding that most of the price variation is driven by between-chain, as opposed to within-chain, associations.

We then add income at the chain-state level to consider the impact of zone pricing; we do so both with, and without chain fixed effects (Columns 4 and 3, respectively). The results show that zone pricing at the state level is an important determinant of prices as well, further reducing the impact of the own-store income variable.

For the drug and mass merchandise stores (Panels B and C) we cannot reliably test the between-chain hypothesis, given the small number of chains, but we can consider the impact of zone pricing. The results in Column 4 show that, like for food stores, the income at the state level is a stronger determinant of the pricing in a store than the income for that particular store.[28]

[28]For mass merchandise stores, there is a negative relationship between prices and income when not including chain

Overall, these results confirm the previous findings: the pricing level appears to be set to reflect the average determinant of purchasing at the chain level, or at the chain-state level, in a form of zone pricing. There is only much more limited evidence of responsiveness to local income at the store level.

4 Demand estimation and optimal prices

4.1 Model

In the previous section we showed that firms appear not to respond to local income in their store-specific prices, but they do respond to the overall income level of the areas where they operate in setting the average chain-level price. How large should the response of prices to income be? We provide a simple benchmark model to present a counterfactual of the optimal pricing for a chain setting prices across stores. We stress that we view the assumptions needed for this counterfactual exercise as unlikely to be exactly satisfied in reality. At a minimum, though, this provides a check on the order of magnitude of the deviation from predictions, and as assessment of the profits losses from pricing uniformity under this benchmark model.

Consider the monopolistic pricing decision of a multi-store chain that aims to maximize the sum of the profits across the different stores s and products j. We assume that the demand function is of the constant elasticity type in each store, $q_{sj} = k_{sj} p_{sj}^{\eta_s}$, with a price elasticity η_s which depends on store s. We return to the assumption of constant-elasticity demand below. We interpret the assumption of monopolistic competition as in the trade literature: firms face competition, which is reflected in the demand elasticity. We also assume constant marginal cost of product j across the different stores, with a possible fixed cost: $C(q) = c_j q_{sj} + C_s$. Thus, the chain maximizes

$$max_{q_{s,j}} \sum_{s,j} p_{sj}(q_{sj}) q_{sj} - c_{sj} q_{sj} - C_s.$$

As well known, the first order conditions yield

$$p_{sj}^* = \frac{\eta_s}{1+\eta_s} c_{sj}$$

or in log terms

$$log\left(p_{sj}^*\right) = log\left(\eta_s / (1+\eta_s)\right) + log\left(c_j\right). \tag{1}$$

Notice that under these assumptions we can infer the optimal pricing of a chain provided we know the store-level elasticity η_s, assuming an average mark-up.

fixed effects because, among the largest two mass merchandise chains, the one operating in, on average, higher income areas has lower prices (see Online Appendix Figure 10b).

4.2 Elasticity estimates

The model above requires estimates of the price elasticity of demand at the store level. As our benchmark measure of elasticity, we estimate the response of log quantity at the weekly level to the weekly log price product-by-product, for each store s. More precisely, we estimate

$$log\left(q_{sjt}\right)=\alpha_i+\eta_s log\left(p_{sjt}\right)+\gamma_i X_{sjt}+\epsilon_{sjt}. \tag{2}$$

That is, for each store s, we regress log quantity on log price using all weeks t and all products j with non-missing observations. The coefficient on the log price is the estimated price elasticity, $\hat{\eta}_s$. We use price variation for all 9 years and all 10 products in order to maximize precision. As controls X_{sjt}, we include year*product fixed effects (to capture the fact that some products vary across years within a module) and 52 week-of-year*product fixed effects to capture product-specific seasonality effects. We cluster the standard errors by a bi-monthly period, thus allowing for correlation across products, as well as over time within a 2-month period.

These elasticity estimates miss two important margins: inter-temporal substitution and cross-product substitution. That is, we assume that, controlling for price in week t, the quantity sold in week t does not depend on the price set in previous weeks; yet, stockpiling behavior (as an example) would violate this condition. Second, we also assume that a sale in period t for a top-selling orange juice does not affect the sale of other juice products; to the extent that there is product substitution, the profit calculations above are incorrect. We revisit these assumptions below.

Setting the concerns about substitution aside for now, we document that the elasticity estimates from (2) are well-behaved: the elasticities range from -4 to -1, a well-behaved distribution. The standard errors for the elasticity estimates, displayed in Figure 9b, range mostly between 0.05 and 0.2 for food stores and between 0.2 and 0.4 for drugstores and mass-merchandise stores for which we use fewer products, implying that the elasticity estimates are precise, with t statistics typically in the double digit range.

Still, there is a degree of noise left in the elasticity estimates, which we take into account with a simple empirical shrinkage procedure. To estimate the amount of shrinkage, we re-estimate the elasticity separately using just the first 26 weeks of year year and again using the next 26 weeks of each year; label these elasticity estimates $\hat{\eta}_{1,s}$ and $\hat{\eta}_{2,s}$. We then ask what is the optimal shrinkage of $\hat{\eta}_1$as a predictor of $\hat{\eta}_2$. We compute the mean squared error $\left[\left(1-\rho\right)\hat{\eta}_{1,i}+\rho\overline{\eta}_1-\hat{\eta}_{2,i}\right]$, where $\overline{\eta}_1$ is the overall average towards which $\hat{\eta}_{1,i}$ is shrunk. Online Appendix Figures 15a-c display the mean squared error as function of the shrinkage for each store type. There is a slight improvement in the prediction accuracy with some shrinkage, but the estimated optimal shrinkage is just $\hat{\rho}=.104$ for food stores, though it is slightly larger at $\hat{\rho}=.305$ for drugstores and $\hat{\rho}=.408$ for mass-merchandise stores. We apply this shrinkage correction to our overall measure of elasticity. Figure 9a displays the distribution of the shrunk elasticity estimates across the 22,680 stores in our sample.

Validation. This provides evidence that the elasticities are precisely estimated, but an equally important concern is about mis-specification: to what extent does the logQ-logP relationship assumed in (2) reflect the demand curve? Is the assumption of constant elasticity approximately correct? Figure 9c shows that this is indeed the case, to a perhaps surprising degree. The figure presents a bin scatter of $log\,(q)$ on $log\,(p)$; to mirror the specification in (2) we use the residuals of such variables from regressions on the controls X. The relationship between the two log variables is remarkably linear. The relationship in Figure 9c aggregates across all products and tens of thousands of stores of all types. Visual inspection of this relationship by product and store-by-store generally yields similarly well-behaved lines (if with different slopes); some additional examples are in Online Appendix Figures 16a-b.

In Online Appendix Figures 16c-d we provide two additional pieces of evidence validating the elasticity estimates for food stores. First, we document that the log price variable explains about half of the remaining variation (in terms of R^2) after controlling for the Xs. Second, we run a regression that augments specification (2) by including also the prices charged in weeks $t-2$ and $t-4$, as well as in week $t+4$. The coefficients on these variables, while statistically significant and in line with the stockpiling predictions, are an order of magnitude or more smaller than the coefficients on price in week t. Furthermore, they are not larger for storable products, like toilet paper and canned soup, than for non-storables, like yogurt.

Determinants of Elasticity. Next, we examine if the estimated elasticity correlates with expected determinants of consumer willingness to pay, such as income. A bin scatterplot (Figure 9d) shows that the estimated elasticity η_s is a remarkably monotonic (and in fact linear) function of the income for each store s within each chain.

Table 5 provides more systematic evidence on the determinants of the estimated elasticity. Consistent with Figure 9d, an increase of \$10,000 is associated with an increase of the elasticity of 0.140 (s.e. 0.014), a point estimate that remains very similar with the addition of chain fixed effects (Column 2). In columns 3 and 4, we add as determinants the share of college graduates, the median home price, and controls for the percent urban share. We also add a simple measure of competition with other stores: indicators for the number of other food stores within 5 kilometers of the store. The coefficients generally have the expected sign, with income as the strongest determinant, and a weak, but correct-signed, effect of the competition proxies.[29]

4.3 Comparing observed and optimal prices

In this section, we bring to the data the specification (1) which predicts the optimal pricing as a function of the store-level elasticity. This allows us to benchmark the observed price variation to

[29]Column 5 in Online Appendix Table 3 shows that it is important to control for the percent urban variables, as without those the competition variables have the opposite sign (though their effect is not significant).

the model-predicted one. In particular, we estimate

$$log\left(p_s\right)=\alpha+\beta log\left(\hat{\eta}_s/\left(1+\hat{\eta}_s\right)\right)+\epsilon_s. \tag{3}$$

Specification (3) derives under the model specification (1) under the assumption that the marginal cost is constant across all stores s, and after pooling across products j. The model in particular makes the prediction that under optimal pricing $\hat{\beta}=1$, that is, the coefficient on the $log\left(\hat{\eta}_s/\left(1+\hat{\eta}_s\right)\right)$ term should be 1. If the chains under-respond to the elasticity variation, instead, we will observe $\hat{\beta}<1$. For our benchmark specification, we instrument the elasticity term, $log\left(\hat{\eta}_s/\left(1+\hat{\eta}_s\right)\right)$, henceforth "log elasticity," with the store-level income to more fully address the measurement error in the elasticity term[30]. The standard errors are clustered by chain in food stores and by chain*state in drugstores and mass-merchandise stores to allow for any within-chain correlation in errors.[31]

First Stage. Figures 10a-d display graphical evidence of the first stage for our specification, relating the log elasticity term, $log\left(\hat{\eta}_s/\left(1+\hat{\eta}_s\right)\right)$, to income y_s. We display the evidence with the same decomposition used for the reduced-form results above (Figure 6a-d): we display the within-chain specification (Figure 10a), the between-chain specification for the food stores (Figure 10b), the within-chain-state specification (Figure 10c), and the zone pricing, between-chain-state relationship (Figure 10d). Remarkably, the relationship is similar across all these dimensions, with a first stage coefficient varying between 0.03 and 0.055.

The first stage that we use in the regression is displayed in Table 5, Columns 6-8. We relate the log elasticity to income with chain fixed effects, estimating the relationship separately for the food stores, the drugstores, and the mass merchandise stores. The first stage would be quite similar if we did not include chain fixed effects (Column 5), but the within-chain is the cleanest variation, given possible compostitional effects in the between-chain comparison.[32]

IV Estimates. Table 6 presents in Column 1 the estimates of specification (3) for the within-chain price variation, instrumenting the log elasticity with income using the first stage documented above. In Column 1 we focus on the within-chain pricing, including chain fixed effects, and further including chain-state fixed effects to control for zone pricing in Column 2. In Column 3 we focus on the zone pricing running the regression at the chain-state level, including chain fixed effects. Finally, in Column 4 we estimate the between-chain relationship.[33]

[30]For this specification we winsorize the store elasticity $\hat{\eta_s}$ at -1.2. This happens very rarely in the case of our benchmark weekly elasticity estimates but more frequently for weekly index and quarterly top-product estimates.

[31]More precisely, for food stores we cluster at the 'parent_code' level, so as to allow for correlation in pricing between two chains (as identified by a separate 'retailer_code') which fall under the same 'parent_code'. For drugstores and mass-merchandise stores, we cluster at the 'parent_code'*chain level.

[32]While the products used for the elasticity computation remain constant across chains, different chains have different sales patterns across different products. To the extent that different products have different average elasticities, this can induce compositional differences in the estimated store-level elasticities across chains. This is much less likely to occur within a chain because of the similarity of pricing within chain.

[33]We use the same first stage for all the specification, treating the regression as a two-sample IV and bootstrapping by 'parent_id' (the chain variable) for food stores and by chain-state for the drug and mass merchandise stores. If instead we had used the respective first stage for the between-chain specification, the estimate would be similar, but

Considering first the food stores (Panel A), the estimated coefficient on the log elasticity term, $\hat{\beta} = 0.092$ (*s.e.* 0.034) indicates a response of prices to elasticity that, while statistically significant, is an order of magnitude smaller than the model prediction of $\beta = 1$. This estimate is even lower after controlling for the chain-state fixed effects (Column 2). These results thus confirm the qualitative conclusion in Figures 6a and 6c, that the within-chain price response is much smaller than predicted by the model. Furthermore, we saw that even this moderate price response to elasticity is likely due to an aggregation bias in the price measure.

In Column 3, we focus on the zone pricing across states, as in Figure 6d. The results imply a substantial response of income to the elasticity term, though smaller than predicted by the model, $\hat{\beta} = 0.351$ (*s.e.* 0.193).

Then in Column 4, we focus on the between-chain relationship, by estimating specification (3) at the chain-level. The estimated coefficient on the log elasticity term in this between-chain regression, $\hat{\beta} = 0.944$ (*s.e.* 0.220), indicates a response that is now consistent with the model: we cannot reject a slope $\beta = 1$. This provides a magnitude for the substantial price-income relationship in the between-chain graph in Figure 6b.

The results for the drug stores (Panel B) indicate a larger within-chain response, but still significantly smaller than predicted by the model: $\hat{\beta} = 0.287$ (*s.e.* 0.040). The between-chain-state relationship (zone pricing) is consistent with the model predictions: $\hat{\beta} = 0.858$(*s.e.* 0.267). The results for the mass merchandise stores (Panel C) are intermediate between the ones for the food stores and those for the drug stores.

Overall, the within-chain relationship is between 4 and 10 times flatter than the model would predict, the between-chain relationship (for food stores) is in line with the model, and the between-state relationship varies from about 3 times smaller than the model implies (for food stores) to in line with the model (drug stores).

OLS Estimates. While our main focus is on the specification which instruments the elasticity with income, in Online Appendix Figure 4 we present the parallel results to the IV ones for the OLS specification in Online Appendix Table 4 and Online Appendix Figures 17-19. We find qualitatively similar results, but the point estimates for the price-log elasticity relationship are about 3 times or more smaller. Thus, this reinforces the conclusion that the within-chain price-elasticity relationship is more than an order of magnitude too flat to be consistent with the model. However, now the between-chain relationship and the zone pricing relationship are also clearly smaller than implied by the model. We favor the IV results since they take care of additional forms of measurement error that our simple shrinkage estimation may not capture.

Robustness. We now consider a series of robustness checks, focusing on the IV results in food

much more noisy, given that there are only 64 chains (and this approach would not be possible for the drug and mass merchandise stores). Importantly, as we show in Figures 10a-d, the point estimates in the first stage are similar across the different specifications.

stores, in Table 7. First, one may be concerned that the results are sensitive on using income as an instrument for log elasticity. In Panel A of Table 7, we show that the results are very similar using a broader range of demographic and competition variables, as in Table 5. This is not surprising, as the income is the strongest determinant of log elasticity.

Further, we examine whether the specific choice of top-selling name-brand items is driving the results. Thus, we replicate the results using different goods to form the price series in the dependent variable. In Panel D of Table 7, we use the 20th most-available good in each of the 10 product categories we use. There is no evidence that this different product makes a difference. Next, in Panel E, we consider the role of branding by a top-selling generic that is common to many chains within a subset of the modules considered. The within-chain relationship of pricing to income or elasticity remains very similar to the one for the top-selling branded good. In Online Appendix Table 5, we show that these patterns persist when using generic top sellers within chains (Panel C), and some high-quality (high-price) products (Panel D). Thus, the exact specification, or product chosen, does not change the results.

4.4 Elasticity and Substitution

While the patterns above are similar for different goods, the results may also reflect two important biasing factors in the elasticity estimates: intertemporal substitution and cross-product substitution. It is possible that stores within a chain differ in their single-product short-run elasticity, but that these differences reflect also stronger patterns of substitution. That is, the stores for which we estimate a more elastic own-price elasticity may also have a high cross-price elasticity. In this scenario, in the high elasticity stores a price sale on product j generates higher sales for product j, but only at the cost of reduced sales for a range of substitute products. In this case, the grocer does not benefit by setting lower prices in the more elastic store (unlike what our simple model suggests). Similar considerations hold for the scenario in which the differences in estimated elasticity across stores within a change reflect also differential patterns of intertemporal substitution.

Quarterly Elasticity. To address the concern of intertemporal substitution, we re-estimate the elasticities at the quarterly level. That is, we average the weekly log price and log units sold across all weeks in a quarter, and then re-estimate our main equation (2). The controls X in this case still include the year*product fixed effects, and include 3 quarter-of-year*product fixed effects.

The estimated store-level quarterly elasticities are smaller (in absolute value) than the benchmark ones (Online Appendix Figure 21a), as expected, but the two measures are highly correlated (Online Appendix Figure 21b). Importantly, the quarterly elasticity measure passes the same validation exercises as our benchmark measure, as Online Appendix Figures 21c-e document: (i) the log-log specification is approximately linear; (ii) the standard errors of the estimated elasticity are still relatively small (at .15 to .4), even if clearly larger than in the benchmark elasticities,[34] and (iii) the measure is still highly correlated with local income.

Given these results, in Panel B of Table 7 we reestimate the results using this different elasticity measure. The within-chain relationship is again estimated to be much too flat compared to the model prediction. The between-chain relationship is an order of magnitude larger, but now significantly smaller than the model predicts ($\hat{\beta} = 0.340$). This is not surprising: given that the estimated elasticities are now closer to -1, the implied price response should be even larger. Overall, though, the qualitative findings are robust to using a medium-run elasticity which addresses the intertemporal substitution.

Price Index. To address the complementary concern of cross-product substitution, we construct a price index for each module, as detailed in Section 2. We then re-estimate our main equation (2). As in the case of the quarterly elasticities, the estimated index price elasticities are smaller (in absolute value) than the benchmark ones (Online Appendix Figure 22a), as expected, but highly correlated with the benchmark elasticity (Online Appendix Figure 22b). The index elasticity also passes the same validation tests (Online Appendix Figures 22c-e).

Given these results, in Panel C of Table 7 we reestimate the results using the index-price-based elasticity measure, while still using as price variable the price for the products in Table 2 (that is, we are not using the index price as dependent variable). The within-chain relationship is again estimated to be much too flat compared to the model prediction. Notice that we cannot re-estimate the between-chain specifications given that the price indices are not comparable across chains. In Online Appendix Table 5, Panel B we show that the within-chain results are similar using the price index as dependent variable.

Overall, the findings are robust to using alternative measures of elasticity that capture the two most important margins of substitution.

4.5 Yearly Average Price versus Weekly Average Price

Our analysis so far has focused on the weekly average price. That is, for each store s and product j, we have taken the equal weighted average across the different weeks t to compute the average price. A different number of interest is the yearly average price for store s and product j, which is the ratio of the annual revenue to the annual units sold for a store-product. The two prices are different in an important way, since consumers are more likely to purchase a product when it is on sale. Thus, the yearly average price will tend to be lower than the weekly average price. More importantly, the yearly average price will mechanically respond to the price elasticity, as stores with more price-elastic consumers will have consumers shop proportionately more when prices are lower.

[34] We cluster the standard errors for the quarterly elasticities at the quarterly level, allowing for correlation across the 10 products.

Thus, stores with more elastic price elasticity will have lower yearly average prices, for given weekly average prices.

We want to clarify that the two averages are relevant for different purposes. Since our primary interest in in the price setting of firms, we have focused the analysis thus far on the weekly average price, which is the closest we can get to the price posted by the stores.[35] The yearly average price is of interest to consider the implications of the price posted for the prices that consumers ultimately pay, taking into account the substitution margin across weeks. This why we turn to it now.

Figures 11a-d reproduce the key findings in Figure 6a-d, comparing the yearly average price to the weekly average price. As expected, the yearly average price is more responsive to within-chain differences in income than the weekly average price (Figures 11a and 11c), with a slope that is about twice as steep. Similarly, the between-state zone-pricing relationship is also stronger using the yearly average price (Figure 11d). The between-chain relationship for food stores, instead, is not much affected (Figure 11b).

In Table 8 we present the regression results for our benchmark IV strategy, comparing the results for yearly average prices (Panel B) to our benchmark results on weekly average prices (Panel A). In particular, in Panel B, instead of using the log of the weekly average price, we use as the dependent variable the log of the yearly average price. The within-chain coefficient ($\hat{\beta} = 0.223$), while more than twice as large as the benchmark estimate in Panel A, is still 5 times smaller than the model prediction under our benchmark of optimal pricing ($\beta = 1$). Thus, even taking into account this margin of adjustment does not bring the level of prices up to what is expected in light of the model. Still, it is interesting to note how the presence of sales works as a partial "automatic stabilizer", guaranteeing that consumers in more-elastic stores pay lower prices over the year, even in presence of uniform pricing.

4.6 Lost Profits

An important implication of the model is that it allows us to compute the lost profits relative to a benchmark in which firms do the optimal pricing, as given by (1). To do this, we assume that empirical marginal costs are equal for each chain-product but are free to differ across products and across chains. We assume that each store within a chain is of equal size and that the pooled elasticity is the relevant elasticity for all products.

We estimate the markup M_c for each chain c using the mean elasticity for stores in each chain (the mean markup is 39%). We then average prices within each chain-module, using both average price posted and average price paid, and define the marginal cost for each chain as the ratio of average price to markup. Note that since the price posted and price paid are not identical, there

[35] As we discussed, even the weekly price may differ from the price posted. For example, if there is a sale for a product and the sale price only applies to consumers with the store card, the weekly price will record a combination of the regular price (for consumers without the card) and the sale price (for the consumers with card).

are two different possible marginal costs that we use.

We assume a demand function with constant elasticity: $q = kp^{\eta}$, but we set the scaling variable to be equal to $k = 1$ for all stores. Profits are thus estimated as $\Pi = p^{\eta}(p - MC)$ for each store-module, for various values of p: the model price using mean chain elasticity, the model price using mean chain-state elasticity, the model pricing using store-level elasticity, and the empirically observed average price. We sum across all store-modules within each chain and then express loss profits as a percent of actual profits: $\frac{\Pi_{theo} - \Pi_{actual}}{\Pi_{actual}}$.

5 Interpretations

In the previous section we documented a set of findings about firm pricing in retail stores, and most importantly: (i) the large majority of chains charge largely uniform prices across all their stores, and thus do not respond to local income, or local demand elasticity; (ii) the chains do appear to instead respond to local income in setting the *overall* level of prices in their stores, with magnitudes approximately consistent with what one would expect given a simple monopolistic competition model; (iii) for a small number of chains that do zone pricing, the pricing across the zones does respond to local income; (iv) the magnitude of the losses from price uniformity is sizable, on the order of 8 percent of profits at the chain level.

We now consider which explanations may make sense of these facts. Some traditional explanations do not appear to apply to this setting, among them **menu costs** (Mankiw, 1985). Grocery stores change prices regularly to implement sales. Thus, it is implausible that a menu cost limits the ability to set different prices at the store level, especially since store-level heterogeneity in income is persistent, and thus local prices would have to be updated only rarely.

A behavioral explanation that is also implausible in this setting is that firm managers have **limited attention** (e.g., Gabaix and Laibson 2006) with respect to the determinants of optimal pricing at the store level. It is hard to imagine that managers are literally not aware, or even optimally inattentive, with regards to the local incomes, or price elasticities, given their access to data, and to consulting firms in this regard, and especially the fact that we examine the role of local income, averaged over several years. This is an obvious variable to observe.

Another possible explanation is that committing to uniform or zone pricing benefits chains by allowing them to soften price competition. Dobson and Waterson (2008) present a model of this **tacit collusion** explanation, and Adams and Williams (2017) find mixed support for it using data from the hardware industry. To test for it in our context, we compare the within-chain response of prices to income for stores with no competitors nearby, and for stores with 5+ competitors nearby (Online Appendix Figures 24a-c). If tacit collusion binds individual stores, we would expect more price response to income in the absence of local competitors. It is possible, though, that the pricing decisions are made at the chain level and thus we compare the extent of non-uniform pricing as a function of the stores in a chain that are isolated (that is, with no competitors nearby) (Online Appendix Figures 24d-e). Either way, we do not find much evidence supporting this model in our setting.

Another possibility is that the price uniformity may be due to a constraint posed by the **advertising** of coupons. Most likely, the advertising markets are the Nielsen DMAs. Thus, advertising constraints would tend to force price uniformity within a DMA, the relevant advertising zone, but not between DMAs. In Online Appendix Figures 25a-b and 26a-d we compare the zone pricing at the state level to the zone pricing at the DMA level (after taking residuals for state-chain fixed effects). For both food stores and drug stores, we find less evidence of zone pricing at the DMA level than at the state level, and about the same for mass merchandise stores. Thus, it does not appear that firms are designing their pricing around advertising constraints.

We discuss more in detail two remaining explanations. The first is one of managerial decision-making costs, or **managerial inertia**. Managers may perceive a cost in deviating from the traditional pricing in the industry, which has indeed been, it turns out, uniform pricing. The managers may not be well incentivised to take the change, while fearing the cost in case a price change backfires.

A different explanation is that managers would like, per se, to price to the local demand elasticity, but they refrain from doing so because of **fairness concerns** among consumers. If consumers respond negatively to price differentiation across the stores, perhaps by boycotting a chain, tailoring prices to a store may not be worthwhile. There is certainly anecdotal evidence that fairness constraints may matter. In a report on the UK grocery pricing, the UK Competition Commission writes "*Asda said that it would be commercial suicide for it to move away from its highly publicized national EDLP pricing strategy and a breach of its relationship of trust with its customers, and it would cause damage to its brand image, which was closely associated with a pricing policy that assured the lowest prices always*" and "*Morrisons stated that adopting a policy of local prices would be contrary to its long-standing marketing and pricing policy, it would damage its brand and reputation built up over many years and would adversely affect customer goodwill, as well as being costly to implement and manage.*" (Competition Commission, 2003)

The two models—managerial inertia and consumer fairness—share some common components. In both cases, the model is consistent with the between-chain results, as firms can still set the right overall level of prices, even as they are concerned, or inertial, about store-specific pricing. In addition, we can model both explanations in terms of the firm facing a fixed cost in deciding whether to price flexibly (as opposed to uniformly). The fixed cost captures either the managerial cost or the expected fairness cost of pricing to market. More precisely, assume that for each store s the firm chooses to price to elasticity if

$$max_{p_{s,j}} \sum_j p_{sj} q_{sj}\left(p_{sj}\right) - c_{sj} q_{sj}(p_{sj}) - C_s - K \geq max_{\bar{p},j} \sum_j \bar{p}_j q_{sj}\left(\bar{p}_j\right) - c_{sj} q_{sj}\left(\bar{p}_j\right) - C_s. \quad (4)$$

That is, the firm decides whether to price to store s, incurring a fixed cost K, or instead set an overall uniform price $\bar{p}_j$, which maximizes profits subject to uniform prices. To be more precise, the fixed costs could apply to two levels. First, the firm could decide store-by-store whether to price to elasticity in store s; in this interpretation the fixed cost K is per store that is not priced uniformly. Or the firm could decide at the chain level whether to price uniformly, or price to elasticity in every store; in this case, the fixed cost K is firm-wide. Under either interpretation, the fixed cost captures the managerial costs or anticipated risk of negative publicity from consumers.

In the first version—that the fixed cost applies at the store level—we expect a threshold policy, in which firms will be more likely to price to store for stores with elasticities more substantially different from the average elasticity, since for these stores the average price $\bar{p}$ is more distant from the optimal prices, and thus the losses larger. Thus, if we rank stores within a chain by elasticity, or an elasticity determinant such as income, we should be more likely to see store-specific pricing for stores at the tails of the distribution. With this in mind, we revisit Figure 6a, which displays a bin scatter of within-chain prices as function of within-chain variation in income. The graph shows no evidence that the more extreme bins (for stores with about $15,000 higher, or lower, income than average for the chain) behave differently from the other stores. Rather, they are on the regression line. Thus, the evidence does not seem to support for this version of the fixed cost model, unless one assumes extremely high costs.

In the second version, the decision is considered at the chain level: the chain computes if the gain from targeted pricing in (4) is larger than the fixed cost K. For each chain, assuming constant marginal costs, we compute the optimal profit under flexible pricing versus with uniform pricing, as outlined in Section 4.6. Chains with a wider distribution of elasticities across their stores will have a larger estimated loss from uniform pricing. The x axis on Figure 12 shows the distribution across chains of this measure which varies from 2-3 percent of profits to over 20 percent of profits. On the y axis, this scatterplot displays for a chain the average quarterly absolute price difference, our measure of price dissimilarity across stores. The scatterplot shows a weak, though, positive association between loss of profits and dissimilarity of pricing.

Overall, this evidence suggests that the the implicit costs of flexible pricing are large, in the range of 10-20 percent of profits. This suggests that firms believe in very large costs from consumers perceiving unfair pricing, or that managerial inertia has very sizable costs.

A final explanation that we consider is **firm learning**: firms may be learning to price to elasticity, especially as access to data increases. While learning is not an explanation of the average finding, it is interesting to ask whether firms are moving over time to flexible pricing from the first years in our sample (2006-08) to the most recent years (2012-14). Online Appendix Figures 27a-b provide no evidence that this is the case: there is no chain that appears that have switched over time to more flexible pricing, and the overall within-chain price-income relationship has remained about the same in the early versus later years.

6 Implications

In this Section, we consider the implications of our findings of price uniformity for a variety of economics contexts.

6.1 Inequality

Jaravel (2016) among others brings attention to the role of store pricing for the rise of inequality in the past decades, and shows that the introduction of novel products catering to higher-income consumers lowered the price for such goods, itself contributing to rising income inequality.

As we document now, price rigidity by retail stores has implications for inequality as well. In particular, we compare the observed average level of prices in areas with different income, with the counterfactual level of prices that one would expect if firms priced flexibly as in our benchmark model. We compute the observed level of prices at a particular income level by simply taking the average price charged by stores with local income in that range. For the counterfactual, we compute the optimal price under flexible pricing, taking the observed income of stores as a given. We do this for food stores (Figure 13a), drug stores (Figure 13b) and mass merchandise stores (Figure 13c).

As the blue circles in Figure 13a show, areas with higher income have higher average prices: an extra $10,000 of local income increases prices in average by about 2 percent. This relationship is consistent with our between-chain relationship (e.g., Table 6 Column 4): chains operating in higher average income areas charge higher prices. Yet, this price-income slope is much flatter than expected if firms were pricing flexibly to the elasticity. Under flexible pricing (green points), the price increase associated with $10,000 higher local income would be about a 5 percent increase, more than twice as large. The difference occurs because of the lack of within-firm pricing variation, which thus flattens the response. The pattern is similar for drug stores (Figure 13b). For mass merchandise stores, the observed price-income relationship is in fact negatively sloped, due to the fact that of the two major chains, the one operating in higher income areas has lower prices. Even there, the counterfactual price (green dots) has a more positive slope with respect to income.

These patterns have quantitatively important implications for inequality: by this calculation, low-income areas (average income of about $20,000) pay about 3 percent higher prices than they would pay under flexible pricing, and high-income areas (average income of about $60,000) pay about 8.5 percent lower prices than under flexible pricing. Thus, price rigidity contributes in a quantitatively important way to inequality. It is not obvious, though, that it would contribute to increases in inequality, as opposed to a level effect that is constant over time. Importantly, consolidation between retailers could increase this pattern.

An interesting aspect of this finding is that it runs counter to the fairness explanation outlined above. Uniform pricing actually triggers “unfair” pricing in the sense of increasing inequality.

6.2 Response to local shocks

A second implication of our findings relates to the response of local prices to macroeconomics shocks (Beraja, Hurst, and Ospina, 2016; Stroebel and Vavra, 2014). Our results imply that the response of prices to demand shocks will be smaller for local shocks than for economy-wide shocks. This occurs because chains that charge uniform pricing will respond to economy-wide shocks in their overall price level, but they will respond to local shocks only to the extent that the local shocks affect a sizable fraction of their stores. Given that most chains span several states, a localized shock will induce a fairly small local response.

We provide a calibration of the size of the effects in Online Appendix Table 8 for food stores. We predict the response to a 1% income shock using the estimated store response to own and chain average income as in Table 4, Column 2. We do this separately for simulated shocks occurring at the county, DMA, state, and national level and for two different price variables: our benchmark price variable (the average weekly price), which is closest to the price posted by the stores, and the average yearly price, which more closely tracks the price paid by consumers taking into account the intertemporal substitution. The percentages shown are the response to a 1% income shock in the indicated locality as a percent of the response of a 1% nationwide shock.

The table indicates that the response to a state-level shock would be only 50 percent (Column 1) or 57 percent (Column 2) as large as the response to a nation-wide shock. This is because many chains are spread across state lines, and would thus respond imperfectly to a local shock. The response drops to 18 percent (Column 1) or 30 percent (Column 2) for shock that occurs at the county level. This exercise, thus, suggests that price uniformity can have first-order implications for the response of prices to macro shocks.

6.3 Incidence of trade costs and taxation

A third implication of uniform pricing relates to the estimation and incidence of trade costs. A large literature estimates trade costs by examining differences in the prices of specific products at geographically separated retail stores. Prior studies are surveyed by Fackler and Goodwin (2001) and Anderson and van Wincoop (2004). As a recent example, Atkin and Donaldson (2015) use prices in the Nielsen RMS data to estimate trade costs, accounting explicitly for the source locations of the products and the possibility of spatially varying markups.

Setting aside for a moment the adjustment for markups, this strategy will estimate trade costs to be larger the more prices vary across space. Uniform pricing would thus lead trade costs to be underestimated. At an extreme, if all stores were owned by a single chain that practiced uniform pricing, the estimated trade costs would be zero. In the observed data, the extent to which they are underestimated will depend on the size and geographic distribution of chains.

How would uniform pricing affect adjustments for markups? Atkin and Donaldson (2015) propose an innovative strategy that infers the extent of market power from the observed passthrough of price shocks in origin locations to prices in stores further away. While they would ideally use the origin wholesale price, this is not available in the data so they use the origin retail price as a proxy. Uniform pricing will tend to increase the estimated passthrough, as it increases the correlation between changes in retail prices in the origin with prices in other markets. It will therefore tend to reduce the level of estimated markups, while (correctly) implying less variation in markups across space. The extent of these effects again depends on the size and distribution of chains.

Both of these points relate to the estimation of trade costs. Uniform pricing also affects the true incidence of these costs. Just as we noted above that uniform pricing tends to raise prices in high-income areas and lower them in low-income areas, so too here it will tend to raise prices in locations close to where products are produced and lower them in remote locations. It thus shifts the incidence of trade costs away from those who actually purchase transported goods and toward those whose goods travel shorter distances.

7 Conclusion

In this paper, we show that most large US grocery and drug-store chains in fact set uniform or nearly-uniform prices across their stores. We show that limiting price discrimination in this way costs firms significant short-term profits. We find managerial costs to be the most plausible explanation for this pattern, possibly along with consumer fairness concerns. We show that the result of nearly-uniform pricing is a significant dampening of price adjustment, and that this has important implications for the pass-through of local shocks, the incidence of trade costs, and the extent of inequality.

References

Adams, Brian and Kevin R. Williams. 2017. "Zone Pricing in Regional Oligopoly." link

Anderson, Eric, Benjamin A. Malin, Emi Nakamura, Duncan Simester, and Jón Steinsson. 2017. "Informational Rigidities and the Stickiness of Temporary Sales." *Journal of Monetary Economics*, in press. link

Anderson, James E. and Eric van Wincoop. 2004. "Trade Costs." *Journal of Economic Literature*, 42(3): 691–751.link

Atkin, David and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." link

Beraja, Martin, Erik Hurst, and Juan Ospina. 2016. "The Aggregate Implications of Regional Business Cycles." *NBER Working Paper No. 21956.* link

Bils, Mark and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy,* 112(5): 947-985. link

Bloom, Nicholas and John Van Reenen. 2007. "Measuring and Explaining Management Practices Across Firms and Countries." *Quarterly Journal of Economics*, 122(4): 1351-1408. link

Broda, Christian and David E. Weinstein. 2008. "Understanding International Price Differences Using Barcode Data." *NBER Working Paper No. 14017.* link

Cavallo, Alberto, Brent Nieman, and Roberto Rigobon. 2014. "Currency Unions, Product Introductions, and the Real Exchange Rate." *The Quarterly Journal of Economics,* 129(2): 529-595. link

Chintagunta, Pradeep K., Jean-Pierre Dubé, and Vishal Singh. 2003. "Balancing Profitability and Customer Welfare in a Supermarket Chain." *Quantitative Marketing and Economics*, 1(1):111-147. link

Cho, Sungjin and John Rust. 2010. "The Flat Rental Puzzle." *The Review of Economic Studies*, 77(2): 560-594. link

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong. 2015. "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *American Economic Review*, 105(3): 993-1029.

Competition Comission. 2003. Safeway plc and Asda Group Limited (owned by Wal-Mart Stores Inc); Wm Morrison Supermarkets PLC; J Sainsbury plc; and Tesco plc: A Report on the Mergers in Contemplation. Cm 5950. TSO, London, UK.

Competition Commission. 2005. Somerfield plc and Wm Morrison Supermarkets plc: A report on the acquisition by Somerfield plc of 115 stores from Wm Morrison Supermarkets plc.

DellaVigna, Stefano and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence" *Quarterly Journal of Economics*, 119(2), 1 353-402.

Dobson, Paul W. and Michael Waterson. 2005. "Chain-Store Pricing Across Local Markets." *Journal of Economics & Management Strategy*, 14(1): 93-119. link

Dobson, Paul W. and Michael Waterson. 2008. "Chain-Store Competition: Customized vs. Uniform Pricing." *Warwick Economic Research Papers No. 840*. link

Eizenberg, Alon, Saul Lach, and Merav Yiftach. 2017. "Retail Prices in a City." link

Ellickson, Paul B. and Misra, Sanjog. 2008. "Supermarket Pricing Strategies." *Marketing Strategy*, 27(5): 811-821. link

Ellison, Sara, Christopher M. Snyder, and Hongkai Zhang. 2017. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets." *CESifo Working Paper Series No. 6285*. link

Fackler, Paul L. and Barry K. Goodwin. 2001. "Spatial Price Analysis." *Handbook of Agricultural Economics, vol. 1*: 971–1024. link

Fitzgerald, Terry and Juan Pablo Nicolini. 2014. "Is there a Stable Relationship Between Unemployment and Future Inflation? Evidence from U.S. Cities." *Federal Reserve Bank of Minneapolis Working Paper 713.*

Gabaix, Xavier and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets" *Quarterly Journal of Economics*, 121(2): 505-540.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li. 2011. "International Prices, Costs, and Markup Differences." *American Economic Review*, 101(6): 2450-2486.

Goldfarb, Avi and Mo Xiao. 2011. "Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets." *American Economic Review*, 101(7): 3130-3161.link

Handbury, Jessie and David E. Weinstein. 2015. "Goods Prices and Availability in Cities." *The Reveiw of Economic Studies*, 82(1): 258-296. link

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. 2014. "Learning Through Noticing: Theory and Evidence from a Field Experiment." *Quarterly Journal of Economics*, 129(3): 1311-1353. link

Heidhues, Paul and Botond Koszegi. 2018. "Behavioral Industrial Oganization", in Handbook of Behavioral Economics, Elsevier.

Hoch, Stephen J., Byung-Do Kim, Alan L. Montgomery, and Peter E. Rossi. 1995. "Determinants of Store-Level Price Elasticity." *Journal of Marketing Research,* 32(1): 17-29.link

Hortaçsu, Ali, Fernando Luco, Steven L. Puller, and Dongni Zhu. 2017. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets."link

Hortaçsu, Ali and Steven L. Puller. 2008. "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market." *RAND Journal of Economics*, 39(1): 86-114.link

Jaravel, Xavier. 2016. "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector." link

Kaplan, Greg and Guido Menzio. 2015. "The Morphology of Price Dispersion." *International Economic Review*, 56(4): 1165-1206. link

Mankiw, N. Gregory. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Quarterly Journal of Economics,* 100(2): 529-537.

Massey, Cade and Richard H. Thaler. 2013. "The Loser's Curse: Decision Making and Market Efficiency in the National Football League Draft." *Management Science*, 59(7): 1479-1495. link

McMillan, Robert Stanton. 2007. "Different Flavor, Same Price: The Puzzle of Uniform Pricing for Differentiated Products." link

Miravete, Eugenio J., Katja Seim, and Jeff Thurk. 2013. "Complexity, Efficiency, and Fairness of Multi-Product Monopoly Pricing." *CEPR Discussion Paper No. DP9641*. link

MMC. 1997a. "Domestic Electrical Goods I: A Report on the Supply in the UK of Television, Video Cassette Recorders, Hi-FI Systems and Camcorders." Monopolies and Mergers Commission, Cm 3675, London: TSO.

MMC. 1997b. "Domestic Electrical Goods II: A Report on the Supply in the UK of Washing Machines, Tumble Dryers, Dishwashers and Cold Food Storage Equipment." Monopolies and Mergers Commission, Cm 3676, London: TSO.

Montgomery, Alan L. 1997. "Creating Micro-Marketing Pricing Strategies Using Supermarket Scanner Data." *Marketing Science*, 16(4): 315-337. link

Nakamura, Emi and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *The Quarterly Journal of Economics*, 123(4): 1415–1464,link

Nakamura, Emi and Jón Steinsson. 2013. "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications." *Annual Review of Economics*, 5: 133-163. link

Orbach, Barak Y. and Liran Einav. 2007. "Uniform prices for differentiated goods: The case of the movie-theater industry." *International Review of Law and Economics*, 27(2): 129-153. link

Romer, David. 2006. "Do Firms Maximize? Evidence from Professional Football." *Journal of Political Economy*, 114(2): 340-365.

Seim, K., and M. Sinkinson. 2016. "Mixed pricing in online marketplaces." *Quantitative Marketing and Economics*, 14(2), 129–155.

Shiller, Ben and Joel Waldfogel. 2011. "Music for a Song: An Empirical Look at Uniform Pricing and Its Alternatives." *Journal of Industrial Economics*, 59(4): 630-660. link

Stroebel, Johannes and Joseph Vavra. 2014. "House Prices, Local Demand, and Retail Prices." *NBER Working Paper No. 20710.* link

Thommasen, Øyvind, Howard Smith, Stephan Seiler, and Pasquale Shiraldi. 2017. "Multi-Category Competition and Market Power: A Model of Supermarket Pricing." *American Economic Review*, 107(8): 2308-2351. link

Zhu, Jian-Da. 2014. "Effect of Resale on Optimal Ticket Pricing: Evidence from Major League Baseball Tickets." link

A Appendix

A.1 Data

A.1.1 Store Selection.

In the RMS data, Nielsen provides a basic categorization of stores into five “Channel Codes”: Convenience, Food, Drug, mass-merchandiser, and Liquor. In the HMS data, there are more detailed “Retailer Channel Codes” and each store is assigned to one of 66 mutually exclusive categories such as Department Store, Grocery, Fruit Stand, Sporting Goods, and Warehouse Club. Our starting sample of food stores includes all stores that are categorized as “Food” stores in the RMS data. All the food stores selected in the final sample fall into the “Grocery” store category in the HMS channel code categorization[36].

Store open and close. Our elasticity estimates are potentially biased by stores entering and leaving the Nielsen dataset (which could be due to things like especially low “closeout” prices or low quantities due to stockouts). We do the same pooled linearity plots (residuals of logP and logQ after removing seasonality and module FE) and look only at the residuals from the weeks one month after entering and prior to leaving the sample. These points are not concentrated in any particular region and still appear near the line of best fit for all store-weeks. We also plot price and quantity sold over time for individual entering and leaving stores. Although some stores have lower quantity sold prior to exiting the sample, overall there are no uniform patterns across stores.

A.1.2 Product selection

We select 10 modules (product categories) based on commonly available and highly-sold products. These products include five that belong to product groups used in Hoch, Kim, Montgomery, and Rossi (1995) (soup, cookies, OJ, soda, and toilet paper), as well as products used in Montgomery (1997) (OJ).

Within a module (e.g., soda), we select a high-selling product (e.g., 12-pack cans of Coke). The product choice aims to ensure that (i) the product is available across as many chains and stores as possible (to ensure comparability across stores and across chains), and that (ii) within a store, it is sold in as many weeks of the year as possible (since otherwise the price is not recorded). Formally, we select the top-availability UPC as the product within a module-year with the highest number of week-store observations with positive sales. We do this for each module, repeatedly year by year. For three modules, this determines the selected product, which thus varies across the years. For the remaining 7 modules, we modify this procedure to be able to keep a constant product across all years.[37] Namely, we consider all products that are present in all nine years, and whose coverage is at most 10 percentage points below that of the top product in a given module and year. Among these, we select the product with the highest availability across years as defined above.

[36]The starting sample of 11,501 Food stores also contains some Discount Stores and Warehouse Clubs, as well as some (likely mislabeled) drugstores.

[37]For the 3 other modules, it was not possible to find a constant product across the years without sacrificing too much the availability objective.

For Drugstores, we replace the top-availability soda UPC with the fifth-availability soda UPC as the top four products go on temporary price reductions extremely rarely and thus bias our elasticity estimates towards zero.

The "generic top-seller within chain" selection is done at the chain level, considering only generic products. We use the Nielsen identifier "CTL BR" to identify (masked) store-brand products. The products that we select in each module across chains may not be comparable.

We choose a different set of generic products to make between-chain comparisons of generic product pricing possible. The procedure is identical to our product selection procedure for top products except that we consider only generic products instead of excluding them. This is possible because Nielsen assigns the same (masked) UPC to products it deems similar. We make a further refinement to ensure that they are products of similar quality: we require that the average price for each store-product is within 20% of each other for stores in the same DMA[38]. However, many of our top branded products actually fail this test so we are erring on the side of being too strict with this requirement. Still, for many products we still have low availability. We consider only the four modules with the highest availability–all above 80%–across all stores (soup, cookies, soda, and yogurt) and construct a pooled price level including only these products.

A.1.3 Prices.

Week offset. To be more precise, prices and units are aggregated over the period Sunday to Saturday for **most** but not all retailers. According to Nielsen: "For scanning data, not all retailers provide weekly data using a Sunday to Saturday definition. Some retailers provide data based on their promotion week, which varies by retailer. Nielsen maps non-Saturday ending weeks received from retailers to the best fit Saturday."

Suspiciously low prices. We noticed that there are 1,118 observations of price = .01 and units sold < 10 in the top products we select. Since most of the products have average prices above .50 (See Table 1 Panel D), and because there is no associated spike in units sold, we believe that these observations are invalid. There are similar issues of lesser frequency with prices between .02 and .10. We decide to drop all prices <= .10 as our log-log elasticity estimation is very sensitive to these outliers.

A.1.4 Pairs Dataset for the Analysis of Store Pricing Similarity

We have a more stringent store selection criteria for the pairs data. Since the measures are pairwise at the weekly (or quarterly) level, we want to ensure a sufficient number of overlapping weeks in each pair. To do this, we define a valid module as a module with non missing observations for at least 60% of all possible weeks[39] (quarters with at least 6 weeks of non missing data within each quarter) over the nine years of data. For a store from our sample of 9,415 stores to be eligible for the

[38]Our understanding is that Nielsen only guarantees identically-sized products when assigning products the same UPC

[39]Note that these are *all* possible weeks out of nine years, i.e. 468 weeks. This is different than our availability measure, where the denominator is the number of weeks where the store has nonzero sales in all products in the ten modules we select.

pairs data, we define stores with at least 7 out of the 10 modules valid by this definition to be "good stores." For each chain, we first sample from the good stores. The remaining stores are sampled if necessary. We emphasize again that the average quarterly log price that we use is the unweighted average log weekly price.

In the within-chain pairs data, we limit the number of stores in each chain to 400 for computational reasons (since the number of pairs scales with the number of stores squared), and we further limit the number of stores per chain to 200 in the distributional histograms for weighting reasons. Out of the 64 chains we select, only five chains have more than 400 stores and only ten chains have more than 200 stores.

For the between-chain pairs, we begin with the set of stores that we sampled for the within-chain pairs. First, we sample one store per chain-DMA if there are multiple chains in the DMA. If only one chain operates in the DMA, we sample two stores. We then drop stores from oversampled chains to reduce the sample to a total of two stores per DMA, with one caveat: we do not drop any chain completely, so some DMAs do have more than two stores in the final between-chain sample.

A.1.5 Demographics

All demographics are zip-code level data from the 2008-2012 5-year ACS. These years represent the middle five years of the Nielsen sample (which covers 2006-2014). We explain how we aggregate this zip-code level demographics into store-level demographics in Section 2.

There is one store in our sample that has missing median home price data. We impute this value by regressing median home price on the other demographics (income, fraction with a bachelor's degree, race, and percent urban) on our sample of 9,415 stores. There are three drugstores that are only visited by one household each which reports a PO Box zip code as its zip code. We use county-level demographics for these three stores.

A.1.6 Competition Measures

We use the HMS panel data to help us construct a measure of competition based on geodesic distance. First, we assume that each HMS household lives at the center of its zip code. For each of the stores in the HMS dataset[40], we use a trip-weighted average of the coordinates of each household in order to arrive at an imputed location for the store. We then count the number of stores within various distances of each store by geodesic distance[41].

A.1.7 Matching Stores from Nielsen RMS to Major Grocer Data

We choose the retailer that has the most stores in the Major Grocer Data. The two datasets have about 1.5 years of overlap covering all of 2006 and part of 2007. While the data from the Major Grocer has 5-digit zip codes, Nielsen data only has 3-digit zip codes. We thus take the sum of units

[40]Since we are not concerned about ownership in this measurement, we use all food stores in the dataset, not just the 9,415 in our sample (that we are confident we can assign to a chain). On the other hand, we do not track store opening and closing so we are implicitly assuming that if the store is open at all within the nine years, it is open for the entirety of 2006-2014.

[41]i.e. distance as the crow flies

sold in 2006 for our top products in 10 modules in both the Major Grocer data and for all stores in the Nielsen sample. We then use the Stata user-defined function (.ado) reclink and perform a fuzzy match, requiring a match of 3-digit zip code and allowing the sum of units sold (in string form) to be slightly mismatched. This results in 134 matches belonging to a single Nielsen retailer_code and 1 match belonging to a different retailer_code. We then limit the set of stores in the Nielsen data to those from the majority retailer and perform the fuzzy match again, but we continue to fail to match the single store. To check the matches, we use a different "manual" matching method where we round each 2006 yearly units sold in both the Nielsen and Major Grocer datasets to the nearest 2% or 5, whichever is larger. We then examine the best matches belonging to each store-module. The modal store always matches the result we obtain using reclink.

A.1.8 Imputing Nonsale Prices from Nielsen RMS Data

We attempt to extract nonsale prices from the average prices in the Nielsen RMS Dataset. First, we only want to consider prices that are "high enough." For each store-year-module, define $p^{80}_{s,j,y}$ to be the 80% percentile price in store s, year y, and module j. Then, we want to ensure that there is only one unique price charged during the week by keeping week t if and only if the same price is recorded for three weeks in a row (or two weeks in a row for the first and last weeks of each year): $p_{s,j,t} = p_{s,j,t-1} = p_{s,j,t+1}$ for $t \in [2, 51]$, week 1 if $p_{s,j,1} = p_{s,j,2}$, and week 52 if $p_{j,52} = p_{j,51}$. From this set of "unique" price-weeks, we keep only those where $p_{s,j,t} \geq p^{80}_{s,j,y}$. We then calculate the price level as detailed in Section 3.3, with the one difference that we omit any years with missing prices from the average. The value that we demean each store-module-year's price by remains the average of all prices, as opposed to the average of nonsale prices[42].

We compare the results we get using this algorithm with the data from the Major Grocer. We ignore the week offset and match the first week in the Nielsen data to the first week of the Major Grocer data, the second week to the second, and so on[43]. Almost all store-product-weeks match the Major Grocer true price data exactly, and in fact the discrepancies seem to be an issue with the MG data (for example, yogurt price of .7945936). There are also a few cases where the MG nonsale price is .80, the MG true price is .75, and this is a "sale" that lasts longer than three weeks so we categorize .75 as a nonsale price as it is above the 80th percentile of yearly prices.

A.2 Midweek Price Changes

In this section, we solve a simple case of offset weeks using constant-elasticity demand and show that this timing bias could explain the observed slope.

[42]This is done both because not all store-years have valid nonsale prices and to facilitate comparisons between this price measure and our benchmark measure

[43]This should not matter because the weeks that we keep in the Nielsen nonsale price are in the middle of a period where the price does not change anyways

Figure 1. Examples of *Uniform Pricing*

Figure 1a. Pricing for Chain 128, Orange Juice

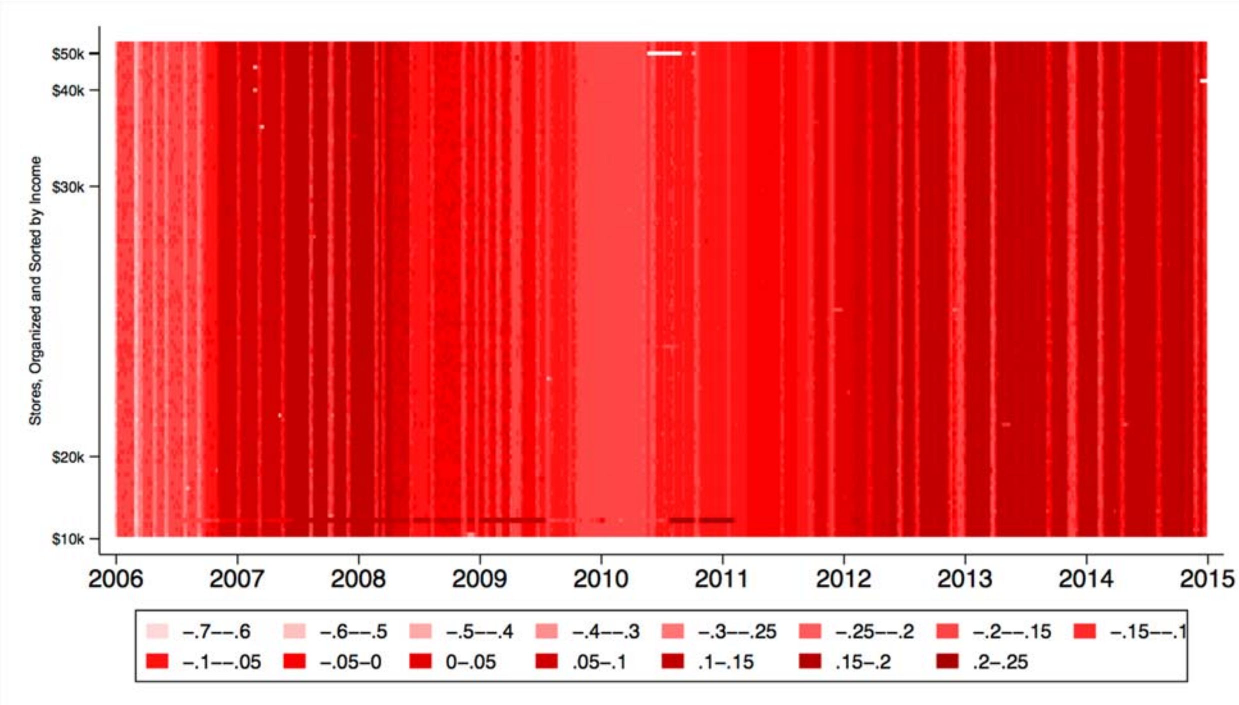

Figure 1b. Pricing for Chain 128, 5 Different Products

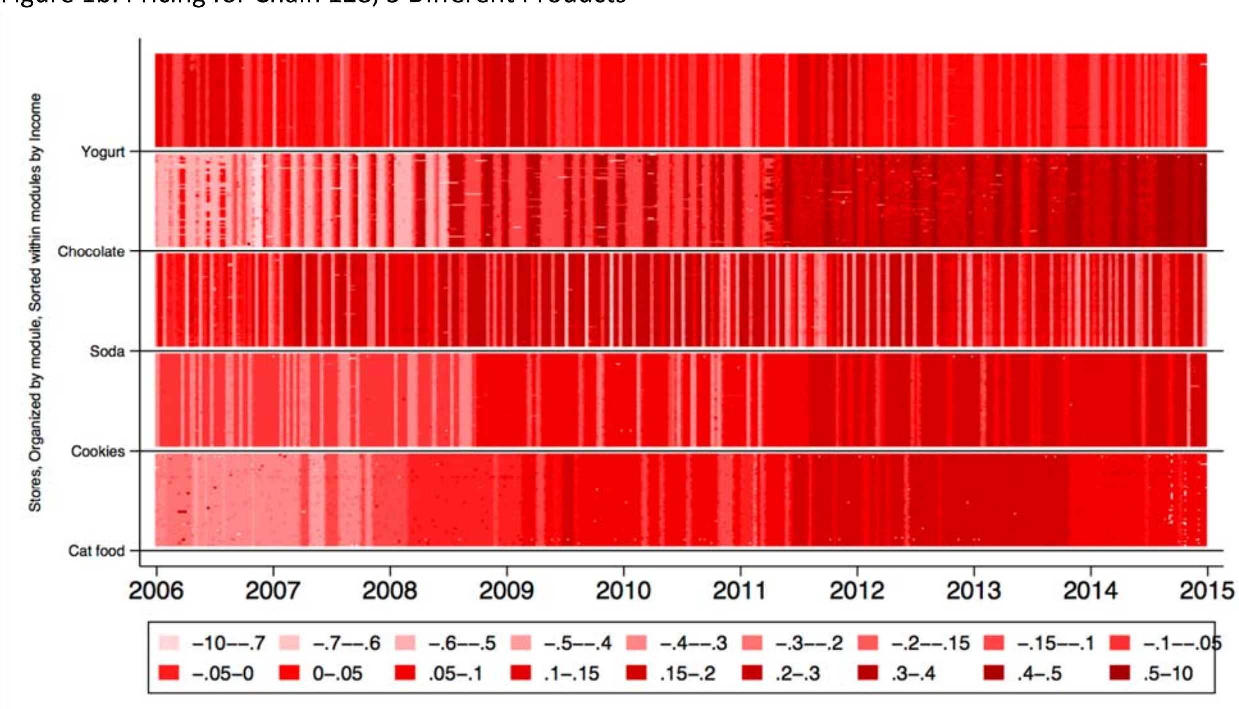

Notes: Plots depict demeaned log prices. Darker colors indicate higher price and are blank if price is missing. Each column is a week t. Each row is a store, and stores are sorted by measure of store-level income per capita. In Figure 1a, dividers are $10,000s. In Figure 1b, the same 50 stores appear for each product.

Figure 2. Example of *Zone Pricing*: Chain 130, Orange Juice

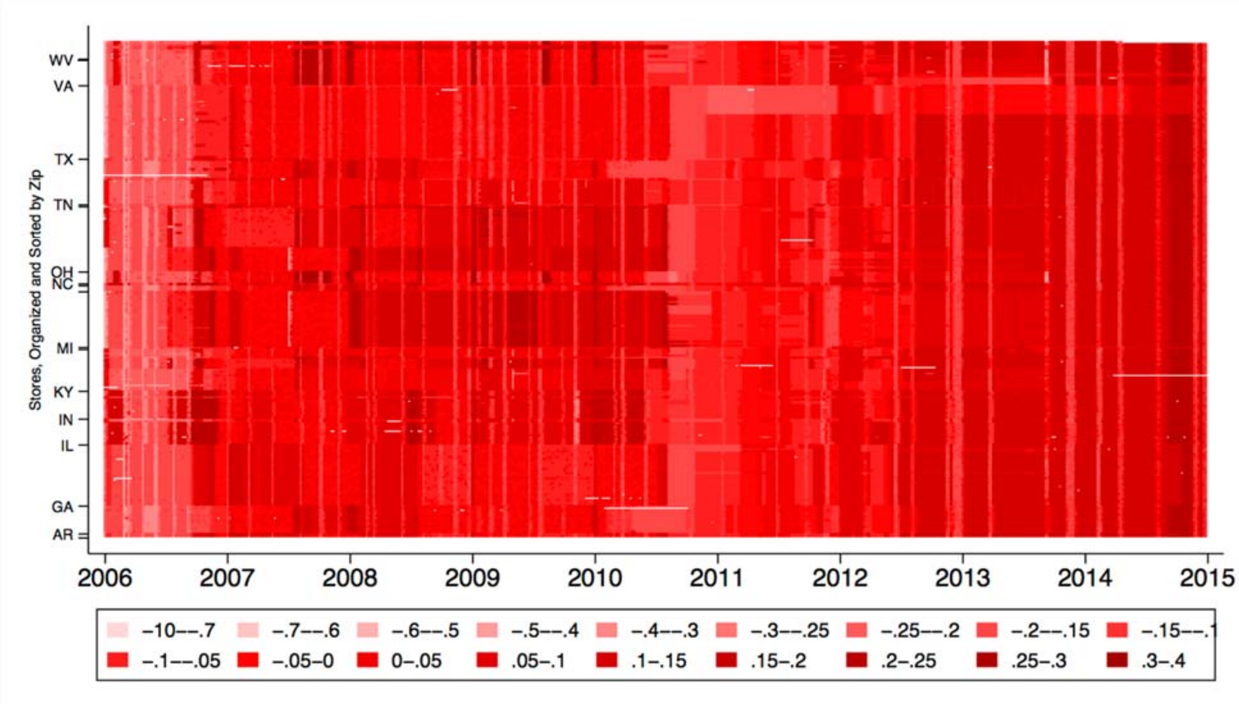

Notes: Plots depict demeaned log prices. Darker colors indicate higher price and are blank if price is missing. Each column is a week t. Each row is a store, and stores are sorted by three-digit zip code within each state divider.

Figure 3. Example of Individualized Pricing: Chain 868, Orange Juice

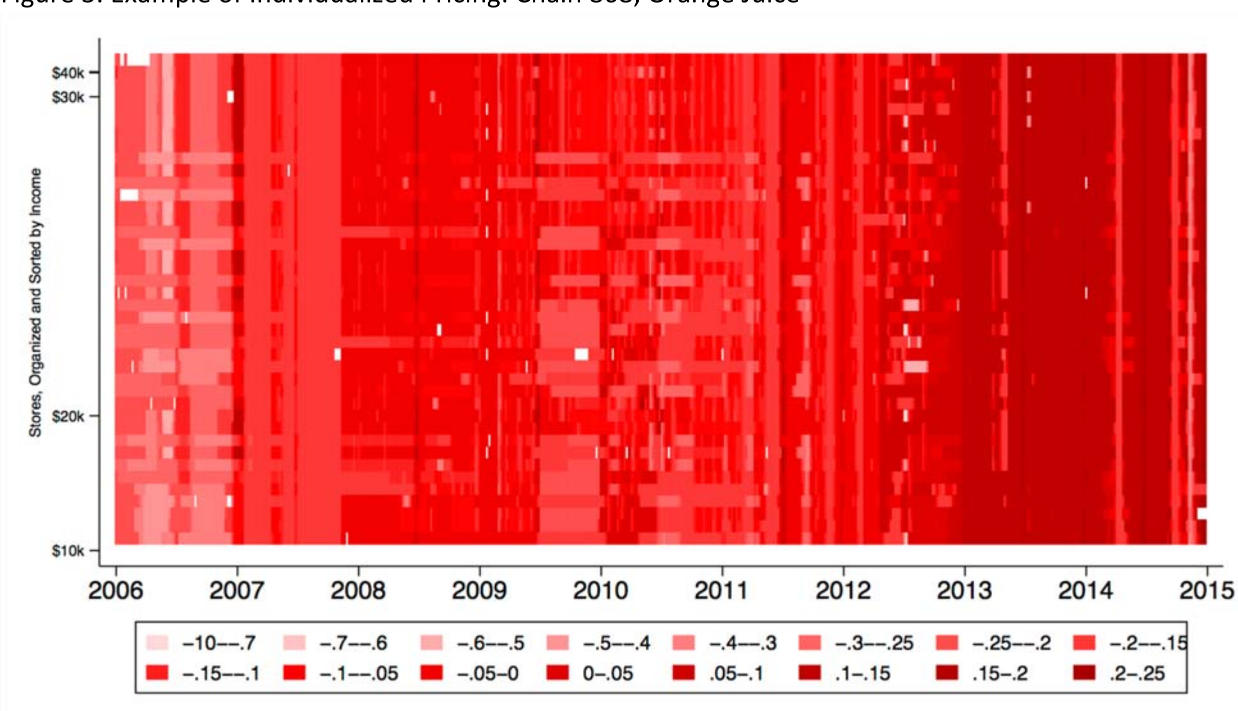

Notes: Plots depict demeaned log prices. Darker colors indicate higher price and are blank if price is missing. Each column is a week t. Each row is a store, and stores are sorted by measure of store-level income per capita.

Figure 4. Similarity in Pricing Across Stores: Same-Chain comparisons versus Different-Chain Comparisons.
Figures 4a-b. All pairs. Quarterly absolute difference in log prices and weekly correlation of log prices

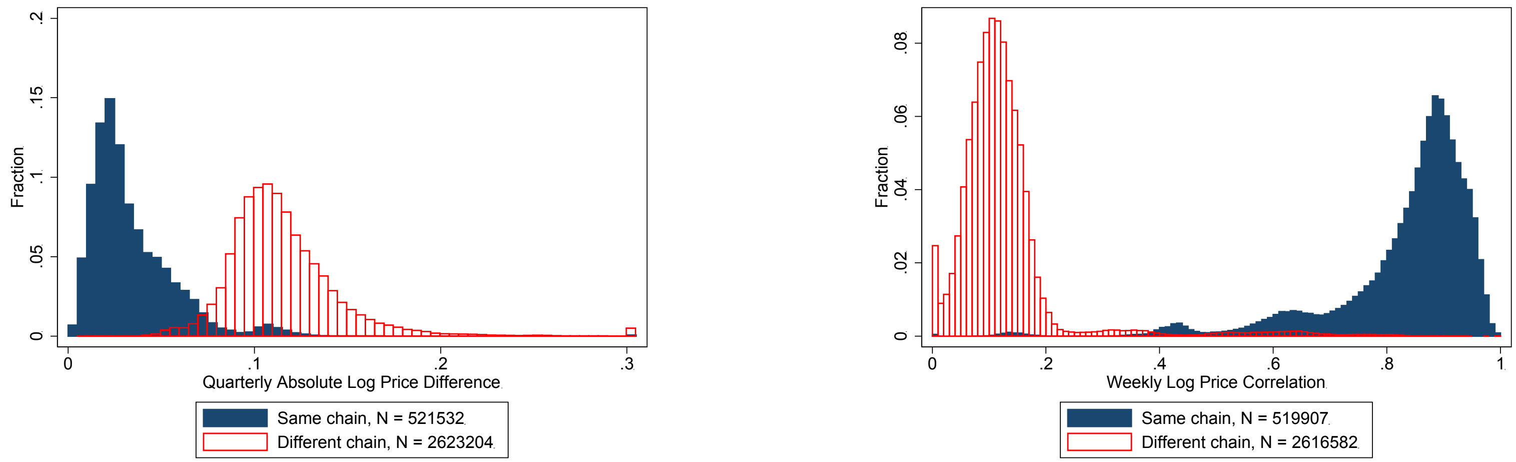

Figure 4c-d. Comparisons Across DMA and top third vs. bottom third of income only.

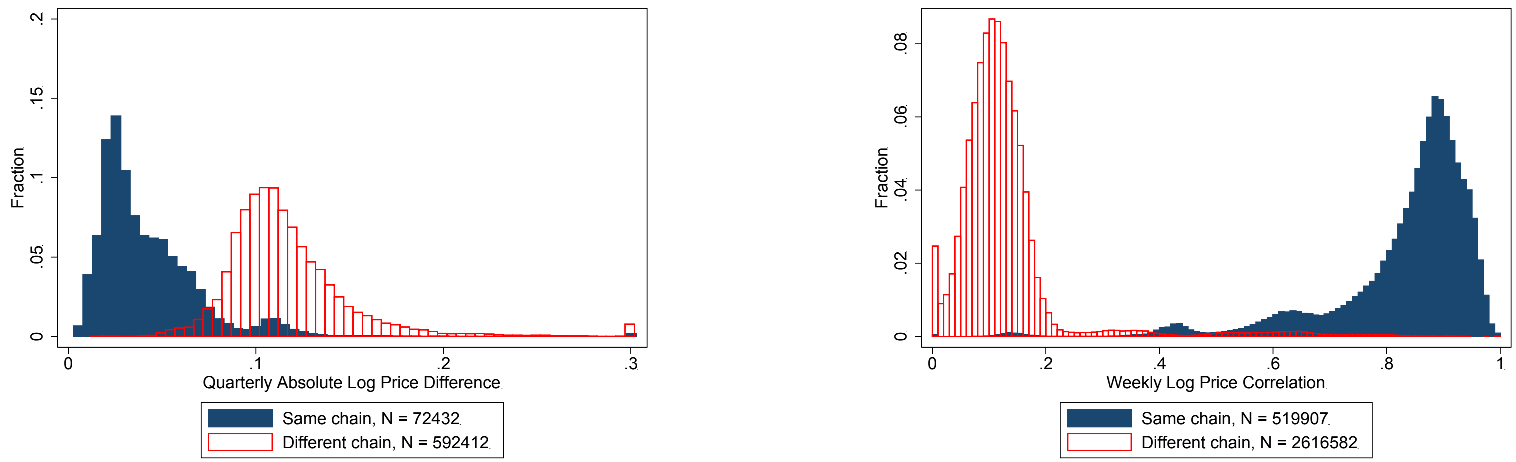

Notes: Each observation is a store-pair. "Same chain" mean same retailer_code. "Different chain" means both different retailer_code and different parent_code. Store pairs within a chain display markedly different pricing patterns compared to pairs in different chains. This relationship holds even when restricting the sample to pairs that should be the most differentiated, such as store pairs in different DMAs and in very different income areas (Panel c and d): even within chains, there should not be any advertising constraints and fairness should not be too large a concern. Quarterly Absolute Log Price Differences are Winsorized at .3 and Weekly Correlation are Winsorized at 0. A maximum of 200 stores per chain are used in the same chain distributions (red outlines) to avoid overweighting the 10 largest chains.

Figure 5. Similarity in pricing, Chain-Level Measure

Figure 5a. Quarterly Similarity in Pricing versus Weekly Correlation of Prices, by Chain

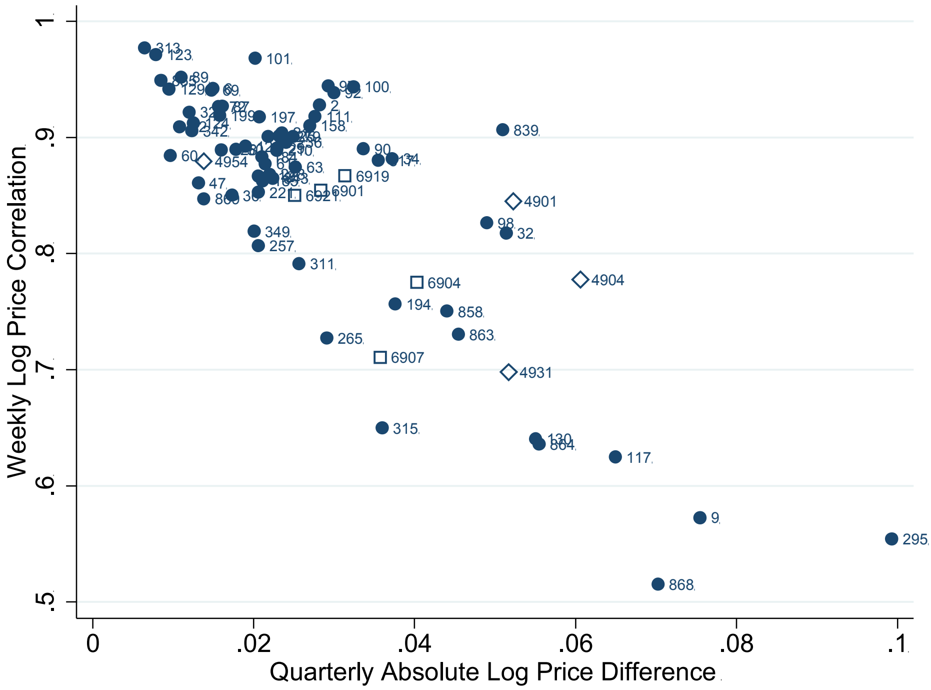

Figure 5b. Within-State Price vs Between-State Price Quarterly Absolute Log Price Difference by Chain

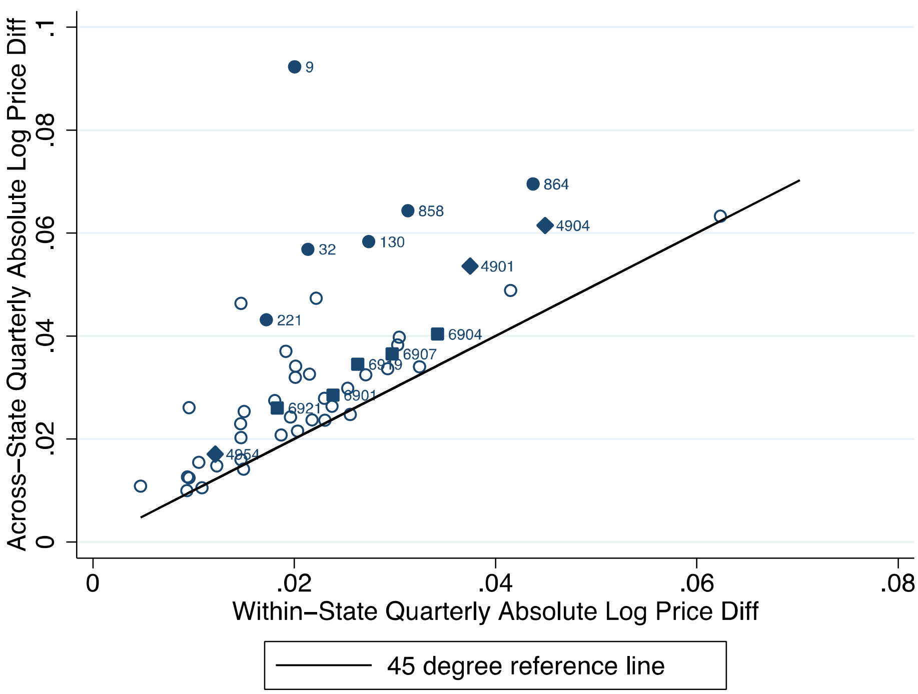

Notes: Circles represent food stores, diamonds represent drugstores, and squares represent mass-merchandise stores. In Figure 5b, each observation is a chain that operates at least three stores in multiple states. Chains that differentiate pricing geographically are labeled. For computational reasons, a maximum of 400 stores per chain are used, which affects only the largest nine chains.

Figure 6. Price versus Store-Level Income

Figure 6a. Price versus Income: Within-Chain

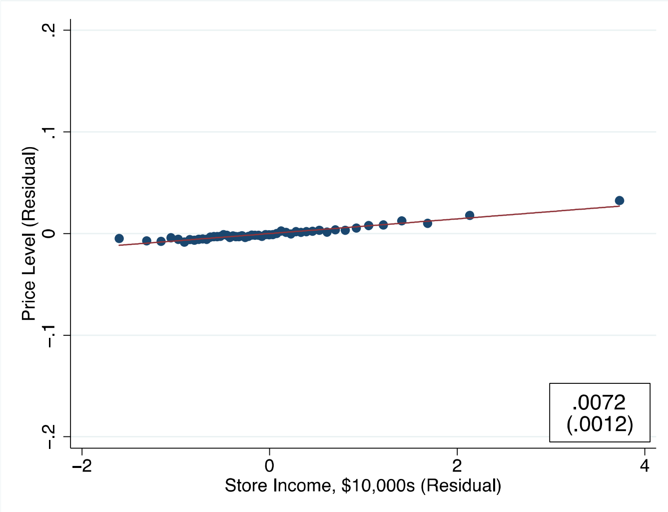

Figure 6b. Price versus Income: Between Chains (Food Stores Only)

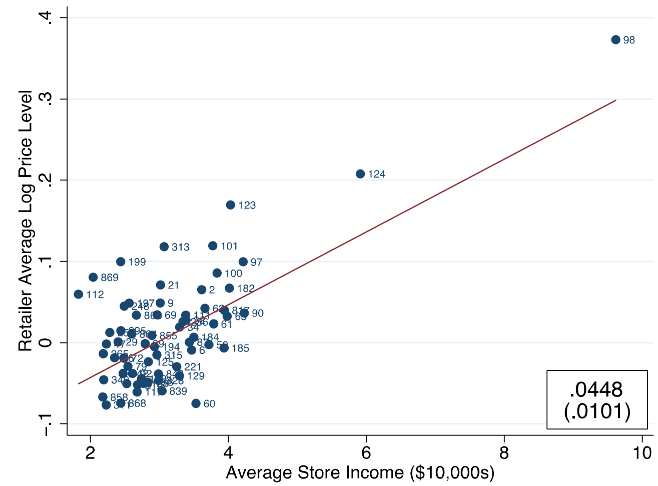

Figure 6c. Price versus Income: Within-Chain-State

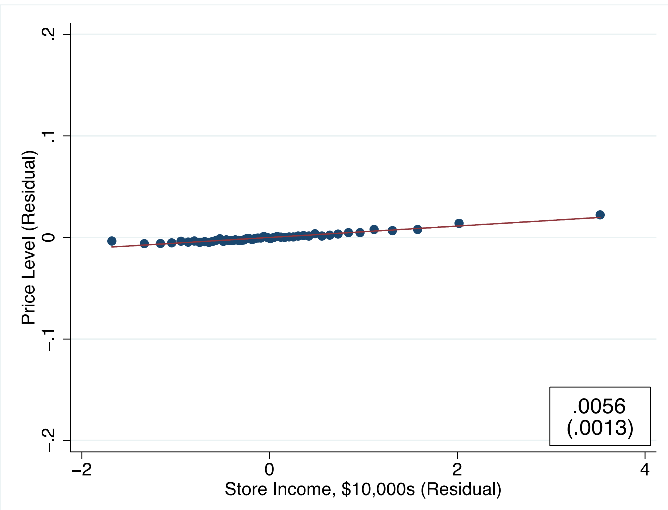

Figure 6d. Price versus. Income: Between Chain-State

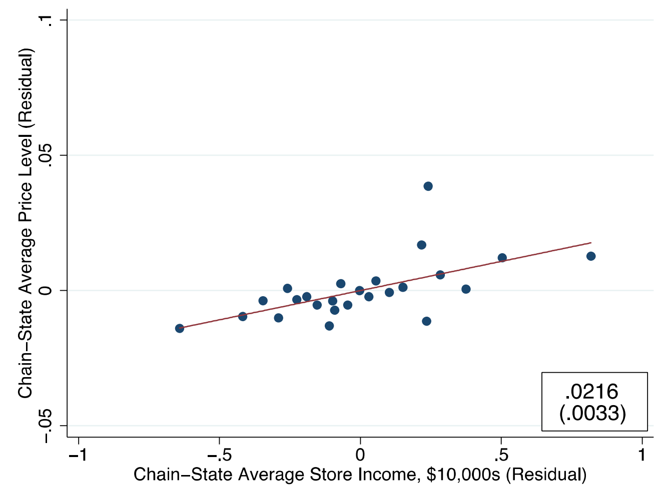

Notes: Standard errors clustered by parent_code. Axes ranges have been chosen to make the slopes visually comparable. Analytic weights equal to the number of stores in each aggregation unit are used in Figures 6b and 6d. In Figure 6a, residuals are after removing Chain FE. In Figure 6c, residuals are after removing ChainXState FE. In Figure 6b, labels indicate Chain. If we exclude outlier Chains 98 and 124, the regression results become .0395 (.0119). In Figure 6d., each observation is one of 25 bins of chain-state averages.

Figure 7. Zone Pricing

Figure 7a. Food Stores (All Chains), State Zones

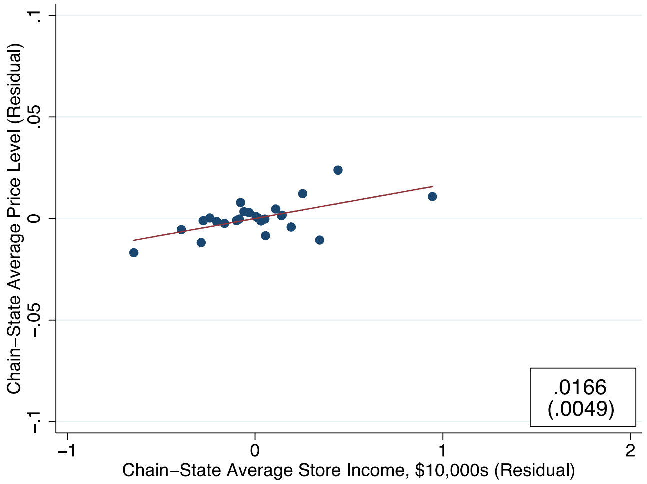

Figure 7b. Food Stores (Zone-Pricing Chains Only), State Zones

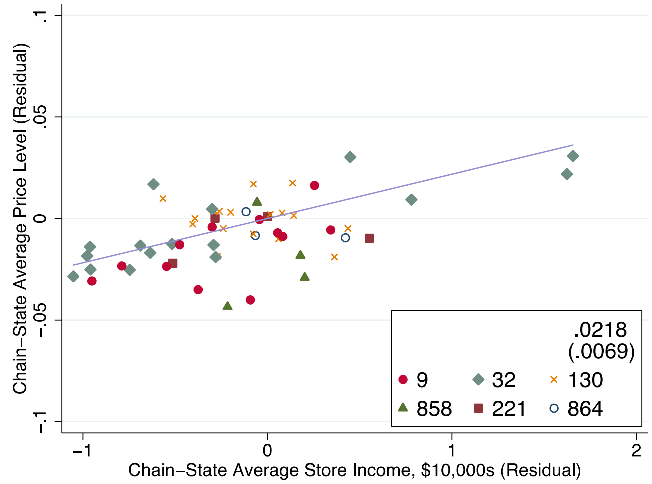

Figure 7c. Drug Stores, State Zones

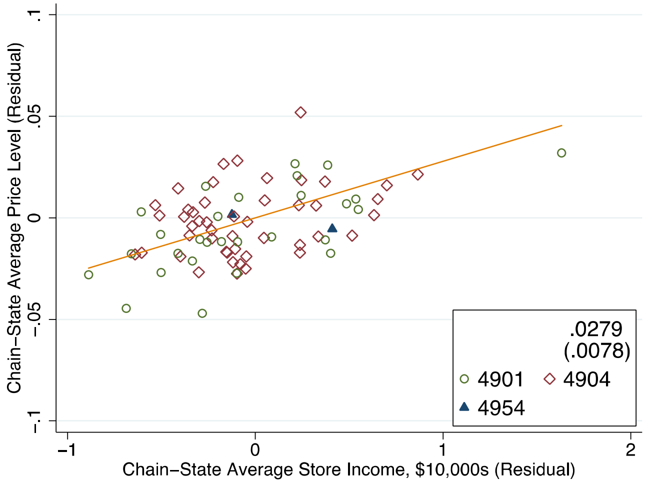

Figure 7d. Mass Merchandise Stores, State Zones

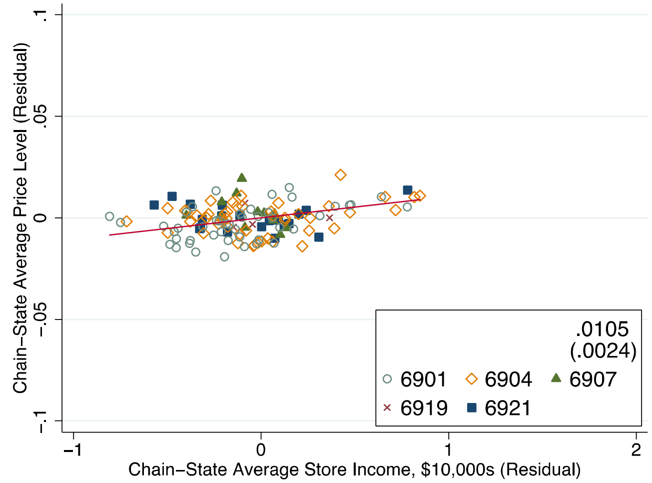

Notes: Standard Errors are clustered by parent_code* state. In Figure 7a., each observation is one of 25 bins of chain-states. In Figures 7b., 7c., and 7d., each observation is an individual chain-state.

Figure 8. Price Response to Income: Investigation using Major Grocer Data

Figure 8a. Nielsen Data: Average Weekly Price

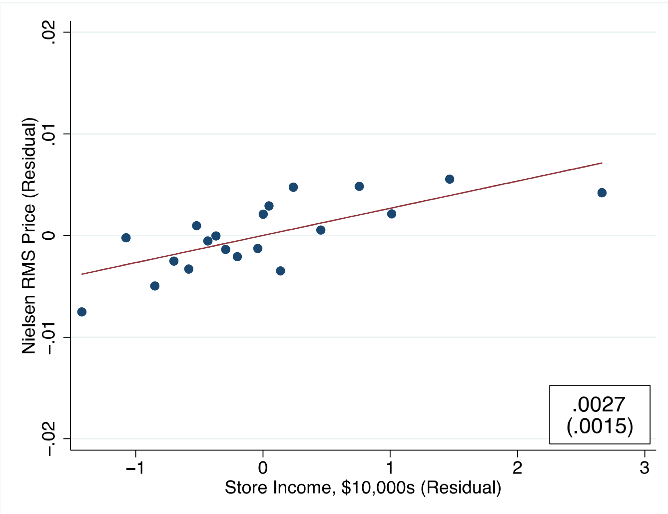

Figure 8b. Data from Major Grocer: Average weekly price

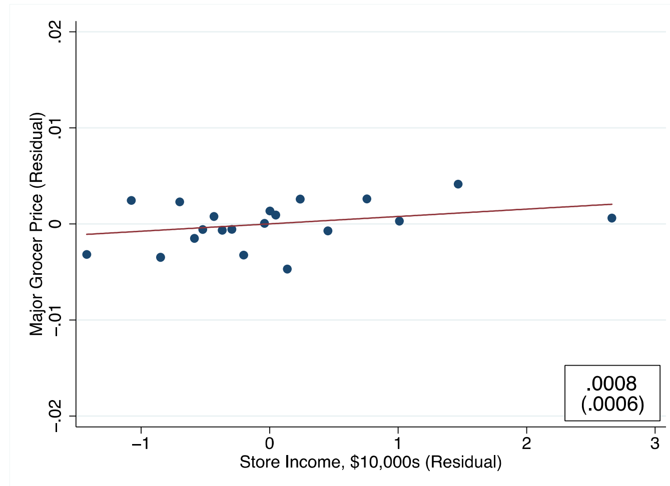

Figure 8c. Data from Major Grocer: Nonsale Price

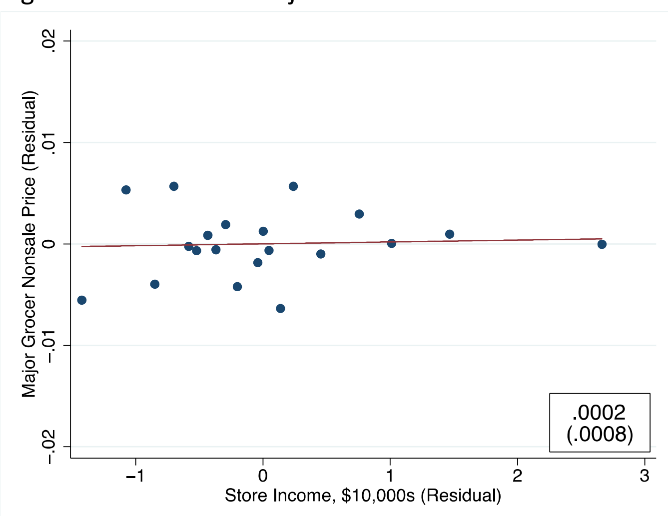

Figure 8d. Data from Major Grocer: Wholesale Cost

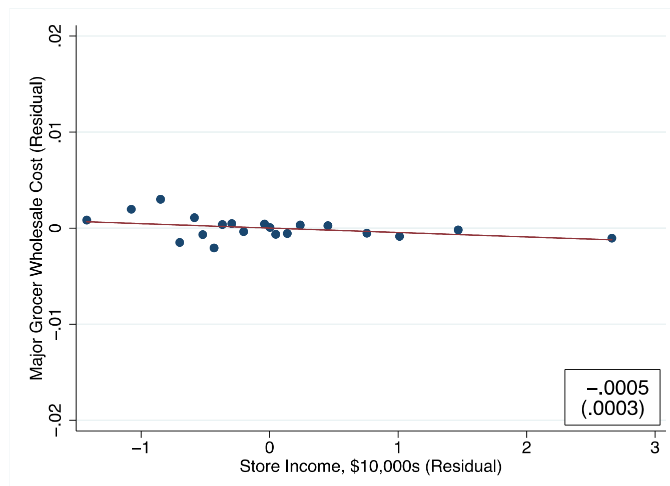

Notes: Stores were matched using 3-digit zip code and total 2006 yearly expenditure for the 10 products we selected. Price Level is calculated using the same 10 top products we select. In each figure, there are 20 quantiles representing 133 stores from the Major Grocer. Values plotted are the residuals after removing state fixed effects, and robust standard errors are used. Price levels are based on are 2006 prices only and are thus not identical to our benchmark top-product price level. Wholesale Cost (Figure 7d) does not include transport or storage costs and is before supplier discounts.

Figure 9. Elasticity Estimates and Validation

Figure 9a. Elasticity Estimates

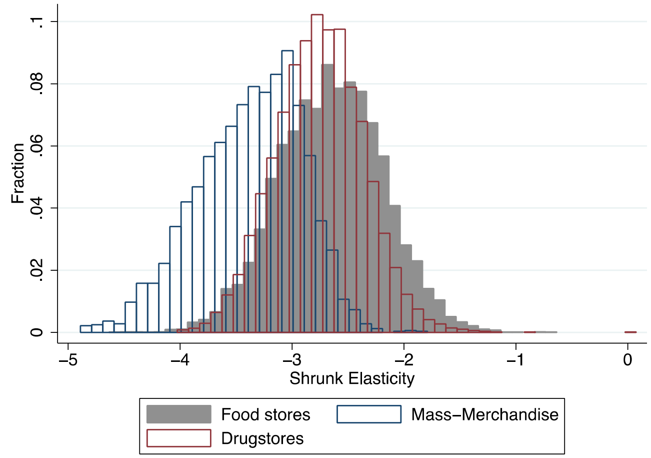

Figure 9b. Elasticity Estimates: Distribution of Standard Errors

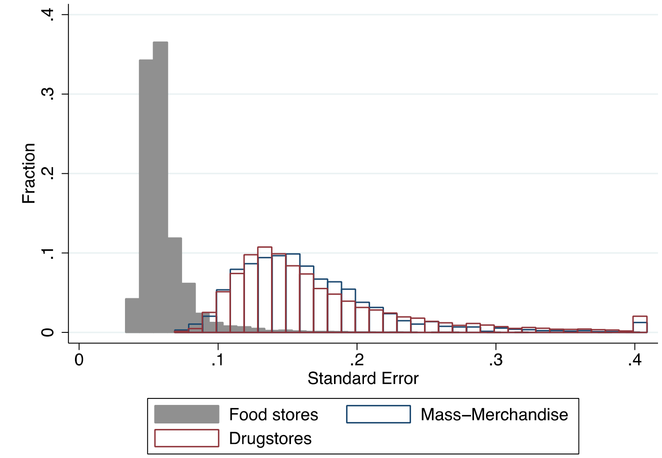

Figure 9c. Validation I. Linearity of Log Q and log P

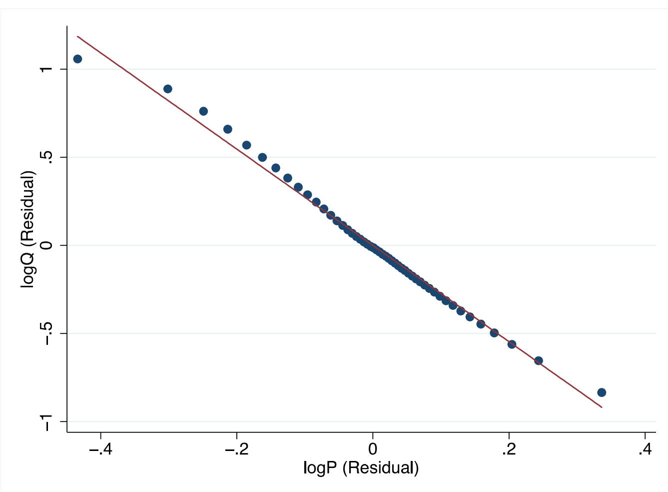

Figure 9d. Validation II. Relationship with store-level income

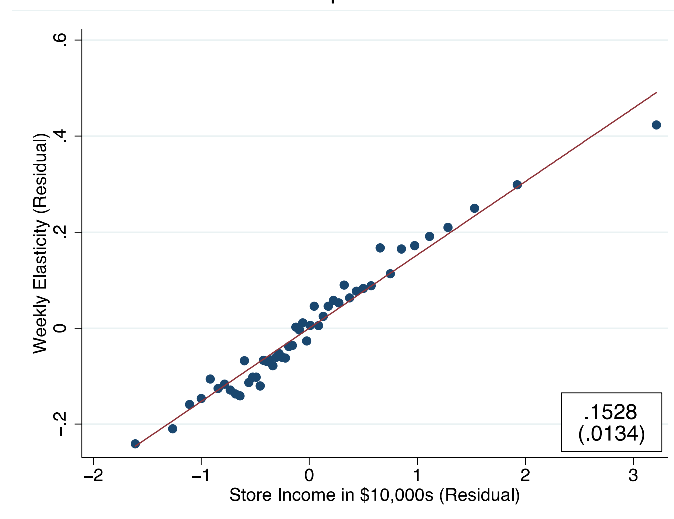

Notes: Figure 9c. 50 quantiles representing 60,552,601 store-module-weeks. Residuals are after taking out module*week of year and module*year FE. Figure 9d. 50 quantiles representing 22,680 stores. Residuals are after removing Chain FE. Standard errors are clustered by parent_code.

Figure 10. Elasticity versus Price, Instrumenting with Income, First Stage

Figure 10a. First Stage, Income and Elasticity within chains

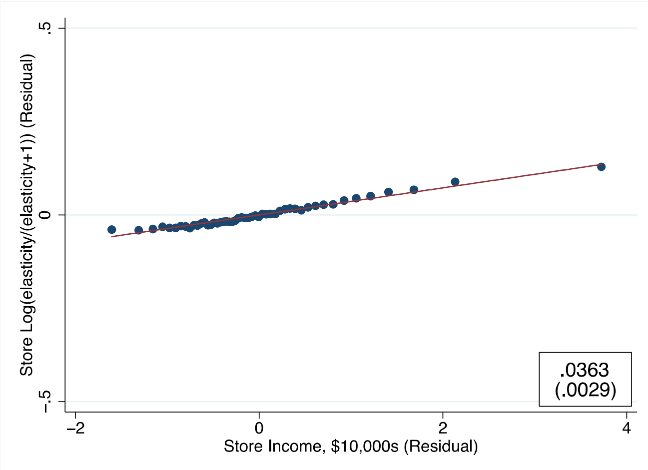

Fig. 10b. First Stage, Between Retailer (Food stores only)

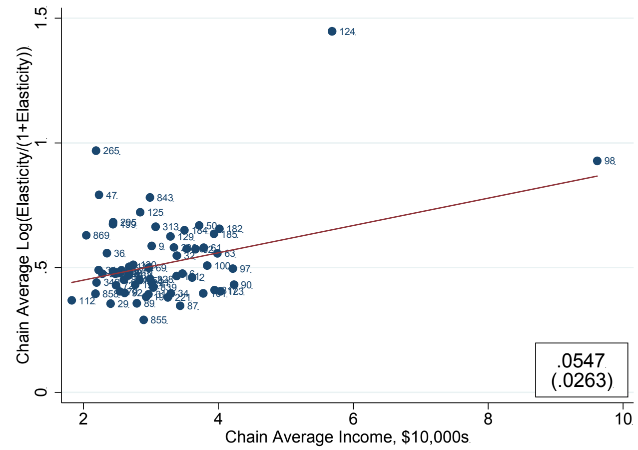

Figure 10c. First Stage, within chain-state

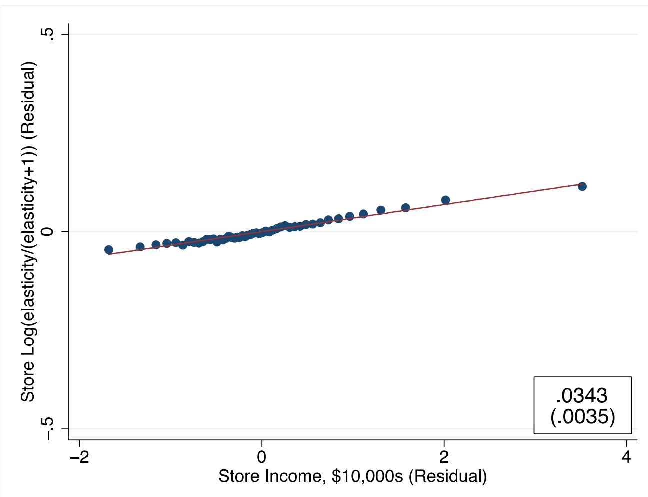

Figure 10d. First Stage, Income and Elasticity, Between Chain-State Averages

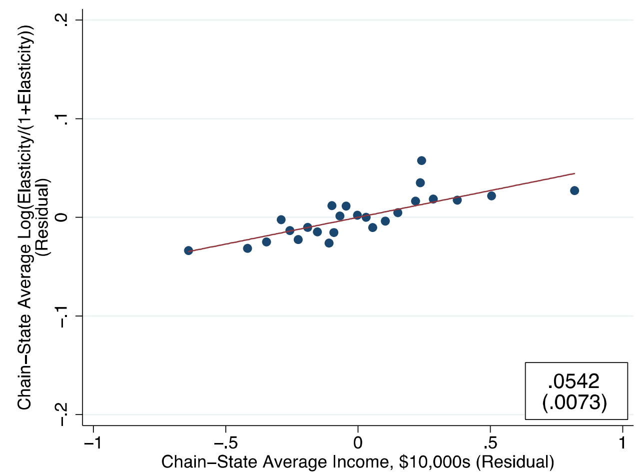

Notes: Axes ranges chosen to make slopes visually comparable. Standard errors are clustered by parent_code in Figures 10a, 10b, and 10c and are clustered by parent_code*state in Figure 10d. Figure 10a: 50 quantiles representing 22,680 stores. Residuals are after removing Chain FE. Figure 10b: The Chain-level average log(e/(e+1)) was calculated by Winsorizing elasticity first and then taking the average log(e/(e+1)). Figure 10c: 50 quantiles representing 22,680 stores. Residuals are after removing ChainXState FE. Figure 10d: 25 quantiles representing 396 chain-state means. Residuals are after removing Chain FE.

Figure 11. Weekly Average Price versus Yearly Average Price

Figure 11a. Within-Chain Price vs Income

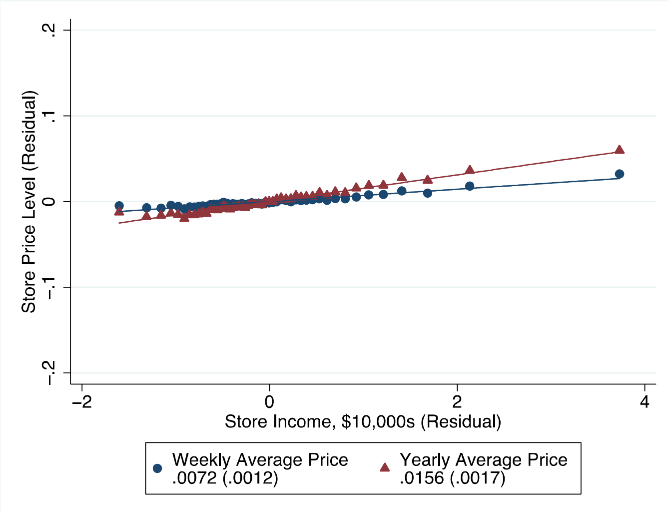

Figure 11b. Between-Chain Price vs Income (Food Stores only)

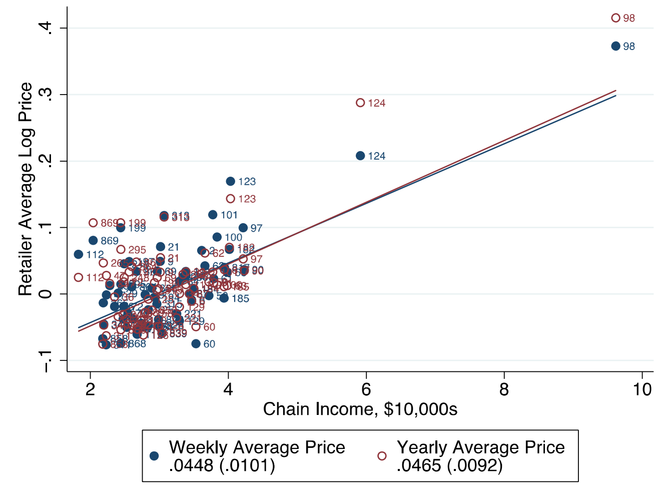

Figure 11c. Within-Chain-State Price vs Income

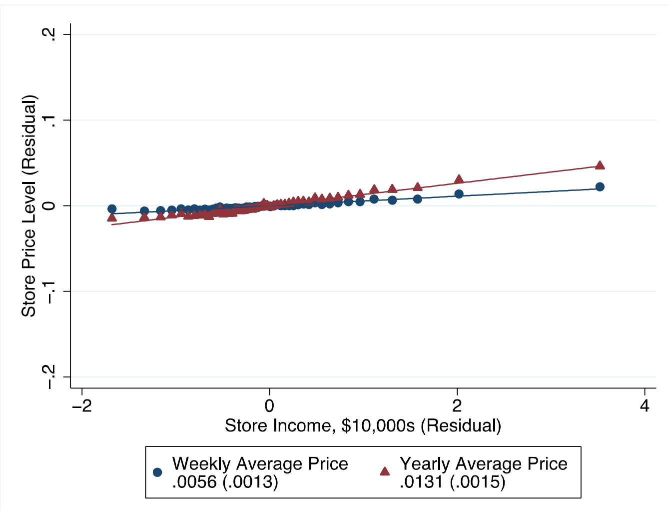

Figure 11d. Between Chain-State Price vs Income

Notes: Residuals are after removing Chain FE. Standard errors are clustered by parent_code. Price Paid is average yearly price paid and is normalized such that the average store has Price Paid = 0. Figures 11a. and 11c. have 50 bins representing 22,679 stores. Figure 11d. has 25 bins representing 396 chain-states.

Figure 12. Profit Loss from Uniform Pricing Versus Price Uniformity at Chain Level.

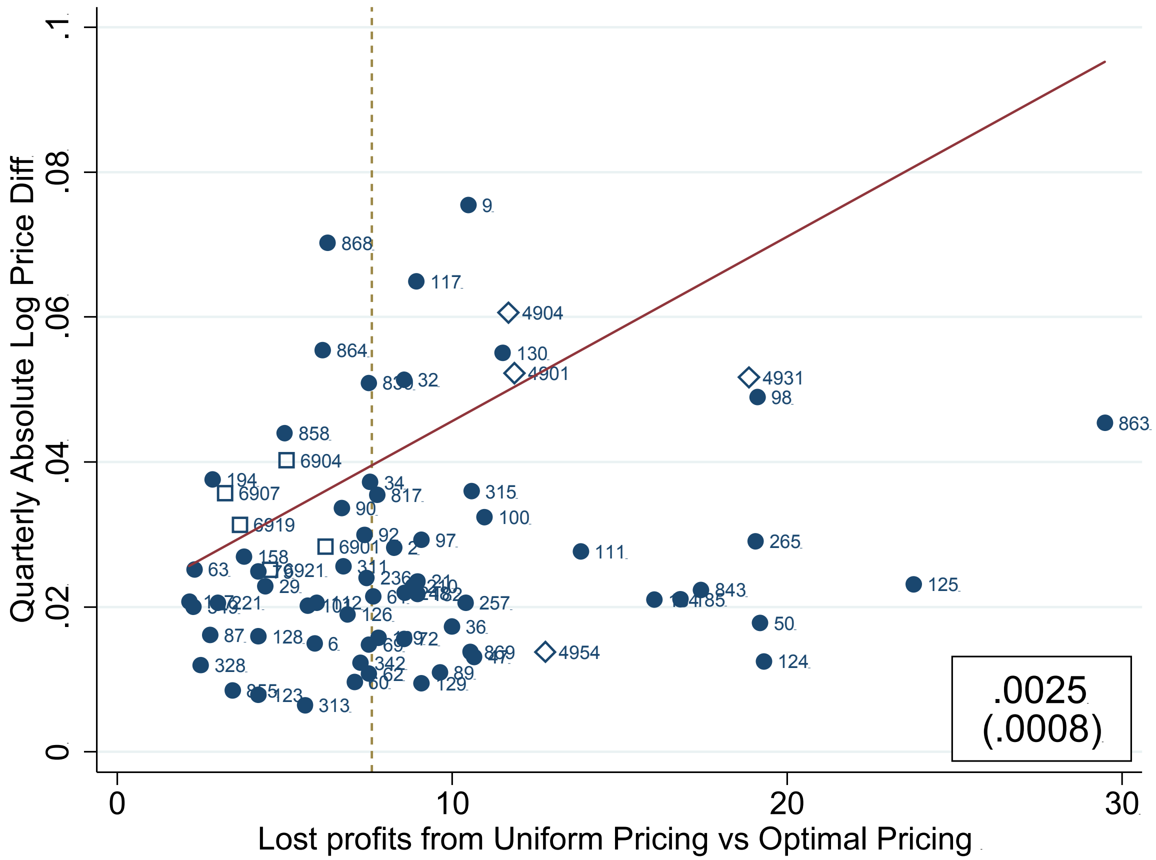

Notes: We approximate Operating Margin as 8.3*(variable costs) based on estimates from Montgomery 1997. Elasticities are Winsorized at -1.2 prior to calculating theoretical lost profits. Dashed vertical line indicates median value of 7.65%. Chain 295 (0.099, 65.25%) has been omitted from both the scatterplot and the regression line. The coefficient including Chain 295 is .0024 (.0008).

Figure 13. Price Rigidity and Inequality: Prices in Areas with Different Income.

Figure 13a. Food Stores

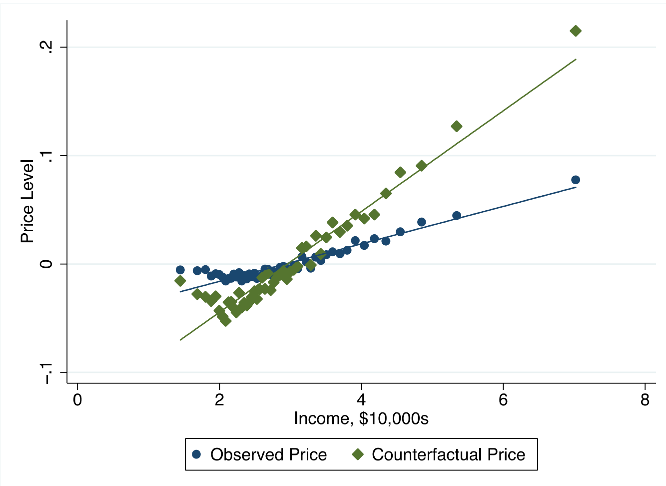

Figure 13b. Drugstores

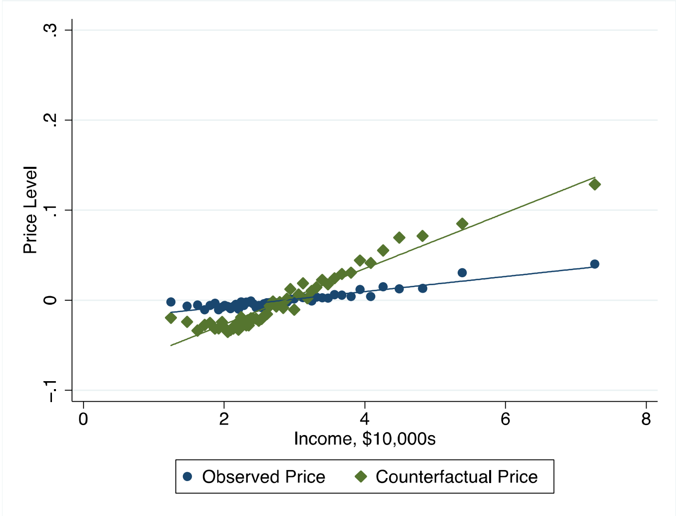

Figure 13c. Mass Merchandise stores

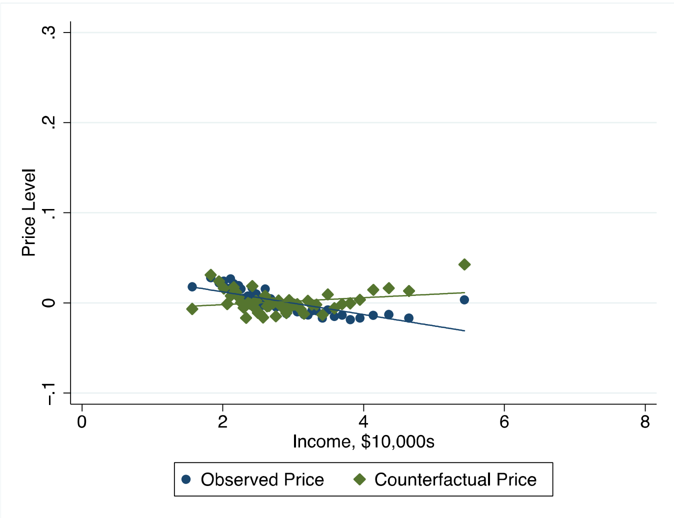

Notes: Counterfactual Price uses flexible pricing, applying the model of monopolistic competition to each chain. We allow marginal cost to vary by chain by keeping the average chain price level relative to other chains unchanged from observed relationships. See Online Appendix Figure 23 for versions that allows marginal cost to vary by chain and for predicted elasticities for food stores.

Table 1. Sample Formation and Summary Statistics: Stores and Chains

	No. of Stores	No. of Chains	No. of States	Total Yearly Revenue	
	(1)	(2)	(3)	(4)	
Panel A. Sample Formation					
Initial Sample of Stores	38,539	326	48+DC	$224 billion	
Store Restriction 1. Stores do not Switch Chain, >= 104 weeks	24,489	119	48+DC	$193 billion	
Store Restriction 2. Store in HMS dataset	22,985	113	48+DC	$192 billion	
Chain Restriction 1. Chain Present for >= 8 years	22,771	83	48+DC	$191 billion	
Chain Restriction 2. Valid Chain	22,680	73	48+DC	$191 billion	
Final Sample, Food Stores	9,415	64	48+DC	$136 billion	
Final Sample, Drug Stores	9,977	4	48+DC	$21 billion	
Final Sample, Merchandise Stores	3,288	5	48+DC	$34 billion	
Panel B. Store Characteristics	**Mean**	**25th**	**Median**	**75th**	
Average per-capita Income	$29,000	$22,450	$26,900	$33,450	
Percent with at least bachelor degree	21.0%	9.3%	17.8%	29.0%	
Number of HMS Households	28.3	11	21	37	
Number of Trips of HMS Households	862	196	502	1162	
Number of Competitors within 5 km	2.3	0	1	3	
Number of Competitors within 10 km	8.0	1	3	10	
Panel C. Chain Characteristics, Food Stores	**Mean**	**25th**	**Median**	**75th**	
Number of Stores	147	30	66	156	
Number of DMAs	7.4	2	4	8	
Number of States	3.4	1	2.5	4	
Panel D. Chain Characteristics, Drug Stores	**Chain 4901**	**Chain 4904**	**Chain 4931**	**Chain 4954**	
Number of Stores	3000	6853	55	69	
Number of DMAs	118	201	9	6	
Number of States	32	48+DC	1	2	
Panel E. Chain Characteristics, Merchandise Stores	**Chain 6901**	**Chain 6904**	**Chain 6907**	**Chain 6919**	**Chain 6921**
Number of Stores	1565	1311	138	30	244
Number of DMAs	190	189	36	13	48
Number of States	47+DC	48	13	11	22

Notes: Valid chains are those in which at least 80% of stores with that retailer_code have the same parent_code and in which at least 40% of stores never switch parent_code or retailer_code. Total Product Revenue is total revenue for our selected products over the nine-year sample. Availability is number of store-weeks with nonzero sales divided by number of store-weeks in which stores in our sample have positive sales in all products belonging to the 10 modules

Table 2. Summary Statistics: Products

		Constant Product	Yearly Product Revenue by Store (in $)	Average Price	Weekly Availability
		(1)	(2)	(3)	(4)
Panel A. Product Characteristics, Food Stores					
Canned Soup	(Campbell's Cream of Mushroom 10.75 oz)	Y	$3,400	$1.18	99.7%
Cat Food	(Purina Friskies 5.5 oz)	Y	$450	$0.49	93.9%
Chocolate	(Hershey's Milk Chocolate Bar 1.55 oz)	Y	$1,650	$0.72	99.7%
Coffee		N	$6,400	$8.45	96.1%
Cookies	(Little Debbie Nutty Bars 12 oz)	Y	$2,100	$1.51	97.3%
Soda	(Coca-Cola 12pk cans)	Y	$34,100	$3.99	99.9%
Orange Juice	(Simply Orange 59 oz)	Y	$5,400	$3.54	99.1%
Yogurt	(Yoplait Low Fat Strawberry 6 oz)	Y	$1,900	$0.64	99.3%
Bleach		N	$1,950	$2.04	96.9%
Toilet Paper		N	$7,000	$8.60	94.9%
Panel B. Product Characteristics, Drug Stores		(1)	(2)	(3)	(4)
Soda	(Coca-Cola 12pk cans)	Y	$3,600	$4.30	93.9%
Chocolate	(Hershey's Milk Chocolate Bar 1.5 oz)	Y	$625	$0.72	95.7%
Panel C. Product Characteristics, Merchandise Stores		(1)	(2)	(3)	(4)
Soda	(Coca-Cola 12pk cans)	Y	$13,300	$4.12	98.5%
Chocolate	(Hershey's Milk Chocolate Bar 1.55 oz)	Y	$725	$0.70	97.9%
Cookies		N	$2,150	$2.57	92.9%
Bleach		N	$2,700	$2.23	94.6%
Toilet Paper		N	$7,600	$8.70	93.2%

Notes: Valid chains are those in which at least 80% of stores with that retailer_code have the same parent_code and in which at least 40% of stores never switch parent_code or retailer_code. Total Product Revenue is total revenue for our selected products over the nine-year sample. Availability is number of store-weeks with nonzero sales divided by number of store-weeks in which stores in our sample have positive sales in all products belonging to the 10 modules

Table 3. Similarity in Pricing Across Grocery Stores, Within-Chain vs. Between-Chain

Measure of Similarity:	**Absolute Difference in Log Quarterly Prices**		**Correlation in (De-Meaned) Weekly Prices**		**Share of Identical Prices (Up to 1 Percent)**	
Wtihin vs. Between:	**Same Chain**	**Different Chain**	**Same Chain**	**Different Chain**	**Same Chain**	**Different Chain**
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Benchmark UPCs, All Store Pairs						
Mean	0.034	0.117	0.836	0.128	0.529	0.107
Standard Deviation	(0.023)	(0.034)	(0.127)	(0.106)	(0.189)	(0.048)
Number of Pairs	491,941	2,620,810	490,077	2,616,142	489,901	2,614,537
Panel B. Benchmark UPCs, Store Pairs Within a DMA						
Mean	0.022	0.115	0.902	0.135	0.619	0.117
Standard Deviation	(0.014)	(0.039)	(0.057)	(0.152)	(0.152)	(0.091)
Number of Pairs	140,989	10,361	140,648	10,369	140,644	10,360
Panel C. Benchmark UPCs, Store Pairs Across DMA, Top 33% income vs Bottom 33% Income Only						
Mean	0.042	0.118	0.808	0.124	0.457	0.106
Standard Deviation	(0.027)	(0.037)	(0.140)	(0.100)	(0.193)	(0.047)
Number of Pairs	60,673	589,645	59,529	588,625	59,496	588,170
Panel D. Generic Product UPCs, All Store Pairs						
Mean	0.032	NA	0.647	NA	0.611	NA
Standard Deviation	(0.026)	NA	(0.193)	NA	(0.201)	NA
Number of Pairs	377,225	NA	373,008	NA	373,008	NA
Panel E. Non-Top Selling UPCs, All Store Pairs						
Mean	0.034	0.117	0.805	0.095	0.578	0.101
Standard Deviation	(0.020)	(0.024)	(0.130)	(0.116)	(0.182)	(0.050)
Number of Pairs	332,195	1,930,054	309,550	1,783,377	309,550	1,783,377
Panel F. Higher Unit Price Items, 8 products in 3 modules only, All Store Pairs						
Mean	0.028	0.152	0.788	0.118	0.642	0.132
Standard Deviation	(0.016)	(0.051)	(0.135)	(0.120)	(0.178)	(0.066)
Number of Pairs	327,457	1,938,276	274,555	1,551,106	274,555	1,551,106

Notes: See Appendix for details on the store sample. The pool that stores are selected from consists of stores that meet our other selection criteria and also have for at least 7 modules nonmissing data for at least 60% of all quarters with minimum six weeks of nonmissing data (columns (1) and (2)) or 60% of all weeks (Columns (3) - (6)). A maximum of 200 stores per chain are used to avoid overweighting the five largest chains. Generic Between-Chain Store Pair Comparisons are not currently possible because we have selected different products for each chain-module.

Table 4. Determinants of Pricing

Dependent Variable:	**Log Prices in Store s**			
	(1)	(2)	(3)	(4)
Panel A. Food Stores				
Own Store Income	0.0175***	0.0044***	0.0029***	0.0029***
	(0.0047)	(0.0013)	(0.0003)	(0.0003)
Chain Average Income		0.0404***	0.0284**	
		(0.0101)	(0.0129)	
Chain-State Average Income			0.0136*	0.0136*
			(0.0069)	(0.0069)
Fixed Effects				Chain
Observations	9,415	9,415	9,415	9,415
R-squared	0.134	0.290	0.296	0.925
Panel B. Drug Stores				
Own Store Income	0.0084***		0.0075***	0.0075***
	(0.0012)		(0.0008)	(0.0008)
Chain-State Average Income			0.0103	0.0203***
			(0.0107)	(0.0074)
Fixed Effects				Chain
Observations	9,968		9,968	9,968
R-squared	0.056		0.063	0.470
Panel C. Mass Merchandise Stores				
Own Store Income	-0.0126***		0.0029***	0.0029***
	(0.0031)		(0.0010)	(0.0010)
Chain-State Average Income			-0.0699***	0.0076***
			(0.0099)	(0.0019)
Fixed Effects				Chain
Observations	3,288		3,288	3,288
R-squared	0.043		0.272	0.916

Notes: In Panel A, standard errors are clustered by parent_code. In Panels B and C, standard errors are clustered by parent_code*state.

Table 5. Determinants of Store-Level Price Elasticity

Dependent Variable:	Store s Shrunk Estimated Price Elasticity				Store s Log((elasticity/(1+elasticity))			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Demographic Controls								
Income Per Capita	0.140***	0.143***	0.0722***	0.0601***	0.0386***	0.0474***	0.0321***	0.0220***
(in $10,000)	(0.0137)	(0.0087)	(0.0201)	(0.0212)	(0.0057)	(0.0046)	(0.0020)	(0.0013)
Fraction with College			0.458***	0.485***				
Degree (or higher)			(0.1136)	(0.1309)				
Median Home Price			0.0037*	0.0049***				
(in $100,000)			(0.0020)	(0.0018)				
Controls for Urban Share			X	X				
Controls for Number of Competitors w/in 5km								
1-4 Other Grocery Stores			-0.0119	0.0040				
			(0.0174)	(0.0142)				
5-9 Other Grocery Stores			-0.0167	-0.0022				
			(0.0226)	(0.0215)				
10+ Other Grocery Stores			-0.0690*	-0.0544				
			(0.0393)	(0.0433)				
Fixed Effect for Chain		X	X			X	X	X
Fixed Effect for Chain*State				X				
Sample:	All Stores				All Stores	Food Stores	Drug Stores	Merch. Stores
R Squared	0.083	0.652	0.669	0.750	0.100	0.697	0.353	0.565
Number of Observations	22,660	22,660	22,660	22,660	22,660	9,415	9,957	3,288

Notes: Standard errors are clustered by parent_code for all columns except for columns (7) and (8), where they are clustered by parent_code*state. All independent variables are our estimate of store-level demographics at the zip-code level based on Nielsen Homescan (HMS) panelists' residences. Data from 2012 ACS 5-year estimates. Percent with College Degree (or higher) is the percent of adults 25 and older with at least a bachelor's degree. Controls for Urban Share are a set of dummy variables for Percent Urban for values in [.8, .9), [.9, .95), [.95, .975), [.975, .99), [.99, .999), and [.999, 1].

Table 6. Responsiveness of Log Prices to Store-Level Log Elasticity

Dependent Variable:	Log Prices in Store *s*		Average Price for Chain-State	Avg. Log Prices for Chain *c*
Specification:	Within-Chain, IV		Between-Chain-State, IV	Between-Chain, IV
	(1)	(2)	(3)	(4)
Panel A. Food Stores				
Log (Elast. / (Elast.+1)) in Store *s*	0.0919***	0.0605***		
	(0.0339)	(0.0100)		
Mean Log (Elast. / (Elast.+1)) in State-Chain Combination			0.351**	
			(0.193)	
Mean Log (Elast. / (Elast.+1)) in Chain c				0.944***
				(0.220)
Fixed Effect for Chain	X		X	
Fixed Effect for Chain-State		X		
Number of Observations	9,415	9,415	171	64
Panel B. Drug Stores				
Log (Elast. / (Elast.+1)) in Store *s*	0.287***	0.231***		
	(0.0400)	(0.0293)		
Mean Log (Elast. / (Elast.+1)) in State-Chain Combination			0.858***	NA
			(0.267)	
Fixed Effect for Chain	X		X	
Fixed Effect for Chain-State		X		
Number of Observations	9,972	9,972	83	
Panel C. Mass Merchandise Stores				
Log (Elast. / (Elast.+1)) in Store *s*	0.187***	0.134***		
	(0.0492)	(0.0436)		
Mean Log (Elast. / (Elast.+1)) in State-Chain Combination			0.478***	NA
			(0.112)	
Fixed Effect for Chain	X		X	
Fixed Effect for Chain-State		X		
Number of Observations	3,288	3,288	142	

Notes: In Panel A, bootstrap clusters are parent_codes. In Panels B and C, bootstrap clusters are parent_code*state. Elasticities are Winsorized to -1.2. Means are average Log Model (not Log Model of average elasticity).

Table 7. Log Prices and Store-Level Log Elasticity, Robustness (Food Stores)

Dependent Variable:	**Log Prices in Store *s***			**Average Price for Chain-State**	**Average Log Prices for Chain *c***	
Specification:	**Within-Chain, IV w/ income**	**Within-Chain, IV, All Variables**		**Between-Chain-State, IV w/ Income**	**Between-Chain, IV w/ income**	**Between-Chain, IV w/ All Vars.**
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Benchmark Product						
Log (Elast. / (Elast.+1)) in Store *s*	0.0919***	0.101***	0.0643***			
	(0.0339)	(0.0349)	(0.0105)			
Mean Log (Elast. / (Elast.+1))				0.351*	0.936***	0.915***
in Chain c				(0.202)	(0.196)	(0.187)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171	64	64
Panel B. Elasticity Computed at Quarterly Horizon						
Log (Elast. / (Elast.+1)) in Store *s*	0.0396**	0.0389***	0.0253***			
	(0.0161)	(0.0097)	(0.0032)			
Mean Log (Elast. / (Elast.+1))				0.151**	0.409***	0.410***
in Chain c				(0.0591)	(0.0933)	(0.102)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,403	9,403	9,403	171	64	64
Panel C. Elasticity Computed with Indices						
Log (Elast. / (Elast.+1)) in Store *s*	0.0388***	0.0367***	0.0253***			
	(0.0141)	(0.0116)	(0.0038)			
Mean Log (Elast. / (Elast.+1))				0.149**	NA	NA
in Chain c				(0.0600)	NA	NA
Fixed Effect for Chain	X	X				
Fixed Effect for Chain*State			X			
Number of Observations	9,258	9,258	9,258	171		
Panel D. 20th-top selling product						
Log (Elast. / (Elast.+1)) in Store *s*	0.0960***	0.105***	0.0749***			
	(0.0251)	(0.0266)	(0.0101)			
Mean Log (Elast. / (Elast.+1))				0.299**	0.936***	0.915***
in Chain c				(0.146)	(0.196)	(0.187)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171	64	64
Panel E. Generic, comparable across chains						
Log (Elast. / (Elast.+1)) in Store *s*	0.0835**	0.0975**	0.0532***			
	(0.0408)	(0.0466)	(0.0204)			
Mean Log (Elast. / (Elast.+1))				0.345	1.486***	1.481***
in Chain c				(0.454)	(0.383)	(0.350)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,296	9,296	9,296	171	61	61

Notes: Standard errors are bootstrapped. Bootstrap samples are clustered at the parent_code level. 100 replications are used. Elasticities are Winsorized at -1.2. Panels B and C do not have the full sample of 9,415 stores because we excluded elasticity estimates with large standard errors. Generic Products in Panel E meet criteria where we think they are likely to be similar products. However, only four modules have sufficient availability and even those products are not sold in all stores.

Table 8. Log Prices and Store-Level Log Elasticity, Price Posted vs. Price Paid (Food Stores)

Dependent Variable:	**Log Prices in Store *s***			**Average Price for Chain-State**	**Average Log Prices for Chain *c***	
Specification:	**Within-Chain, IV w/ income**	**Within-Chain, IV, All Variables**		**Between-Chain-State, IV w/ Income**	**Between-Chain, IV w/ income**	**Between-Chain, IV w/ All Vars.**
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Price Variable is Average Log Price Posted (Benchmark)						
Log (Elast. / (Elast.+1)) in Store *s*	0.0919***	0.101***	0.0643***			
	(0.0339)	(0.0349)	(0.0105)			
Mean Log (Elast. / (Elast.+1)) in Chain c				0.351*	0.936***	0.915***
				(0.202)	(0.196)	(0.187)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171	64	64
Panel B. Price Variable is Log of Average Yearly Price, instead of Average Weekly Price						
Log (Elast. / (Elast.+1)) in Store *s*	0.223***	0.231***	0.197***			
	(0.0316)	(0.0326)	(0.0149)			
Mean Log (Elast. / (Elast.+1)) in Chain c				0.479***	0.979***	0.936***
				(0.154)	(0.239)	(0.220)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171	64	64

Notes: Standard errors are bootstrapped. Bootstrap samples are clustered at the parent_code level. 100 replications are used. Elasticities are Winsorized at -1.2.

Table 9. Estimated Loss of Profits at Chain Level

	Mean	25th	Median	75th	
Panel A. Food Stores					
Loss of Profits Comparing Optimal Pricing to Actual Pricing	8.78%	5.21%	7.13%	9.35%	
Loss of Profits Comparing Optimal Pricing to Uniform Pricing	9.71%	5.70%	7.82%	10.31%	
Loss of Profits Comparing Optimal Pricing to State-Zone Optimal Pricing	7.66%	4.19%	7.16%	8.80%	
Panel B. Drugstores	**Chain 4901**	**Chain 4904**	**Chain 4931**	**Chain 4954**	
Loss of Profits Comparing Optimal Pricing to Actual Pricing	8.97%	8.55%	14.17%	8.97%	
Loss of Profits Comparing Optimal Pricing to Uniform Pricing	12.03%	11.51%	18.72%	12.19%	
Loss of Profits Comparing Optimal Pricing to State-Zone Optimal Pricing	9.17%	8.15%	18.72%	11.73%	
Panel C. Mass Merchandise Stores	**Chain 6901**	**Chain 6904**	**Chain 6907**	**Chain 6919**	**Chain 6921**
Loss of Profits Comparing Optimal Pricing to Actual Pricing	6.10%	5.00%	3.10%	3.49%	4.29%
Loss of Profits Comparing Optimal Pricing to Uniform Pricing	6.39%	5.20%	3.23%	3.63%	4.49%
Loss of Profits Comparing Optimal Pricing to State-Zone Optimal Pricing	4.23%	4.06%	2.18%	0.73%	1.74%

Notes: Elasticities are Winsorized at -1.2. "Actual Pricing" is within-chain price-level predicted with elasticity by store type. Uniform Pricing and State-Zone Optimal Pricing assume that the chain prices to the mean elasticity in the chain or in the chain-state.

Appendix Figure 1. Store Locations

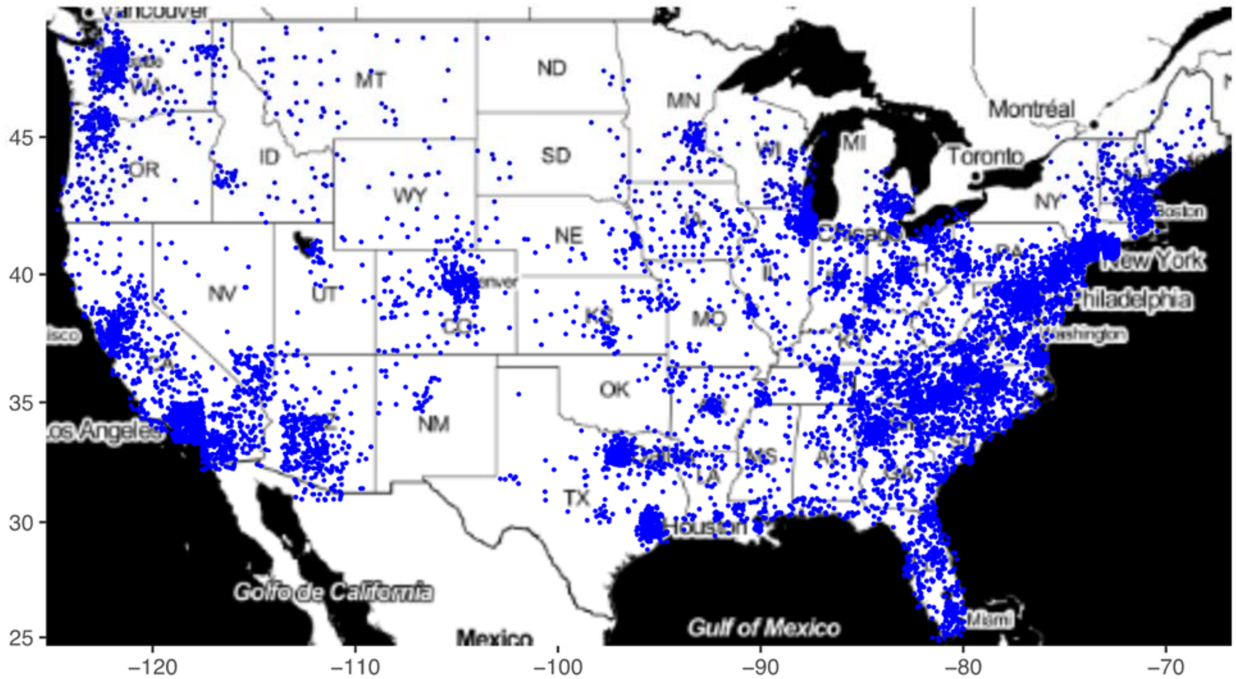

Note: Stores are placed at the midpoint of the county given in the RMS dataset, but locations are jittered so that stores do not overlap. In some cases, this may cause stores near state borders to be placed in the wrong state. This was a tradeoff we made to show more accurately the number of stores located in very large counties in Arizona and Southern California.

Online Appendix Figure 1a. Additional Examples of Chains with Uniform Pricing: Chain 2

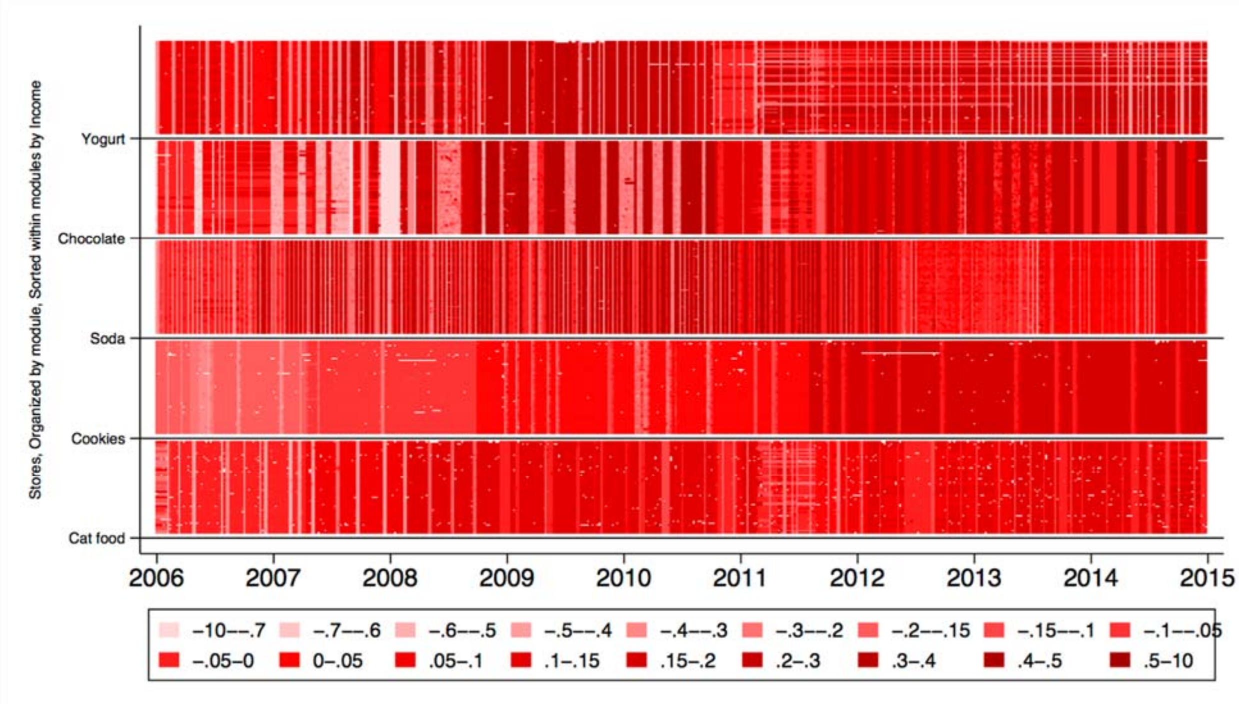

Online Appendix Figure 1b. Additional Examples of Chains with Uniform Pricing: Chain 79

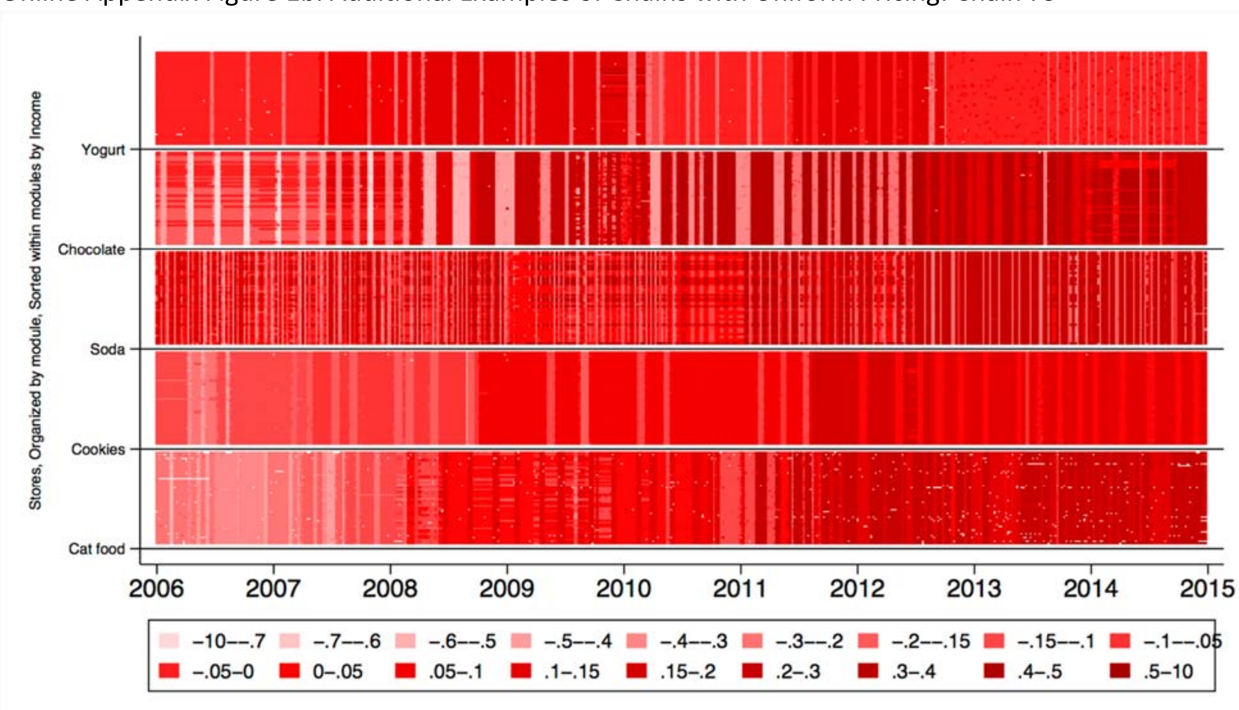

Notes: Plots depict demeaned log prices. Darker colors indicate higher price and are blank if price is missing. Each column is a week t. Each row is a store, and stores are sorted within products by measure of store-level income per capita. The same 50 stores appear for each product.

Online Appendix Figure 2. Additional Examples of Chains with Geographic Pricing Blocks

Online Appendix Figure 2a. Example of Chain with Geographic Pricing Blocks: Chain 9, Orange Juice

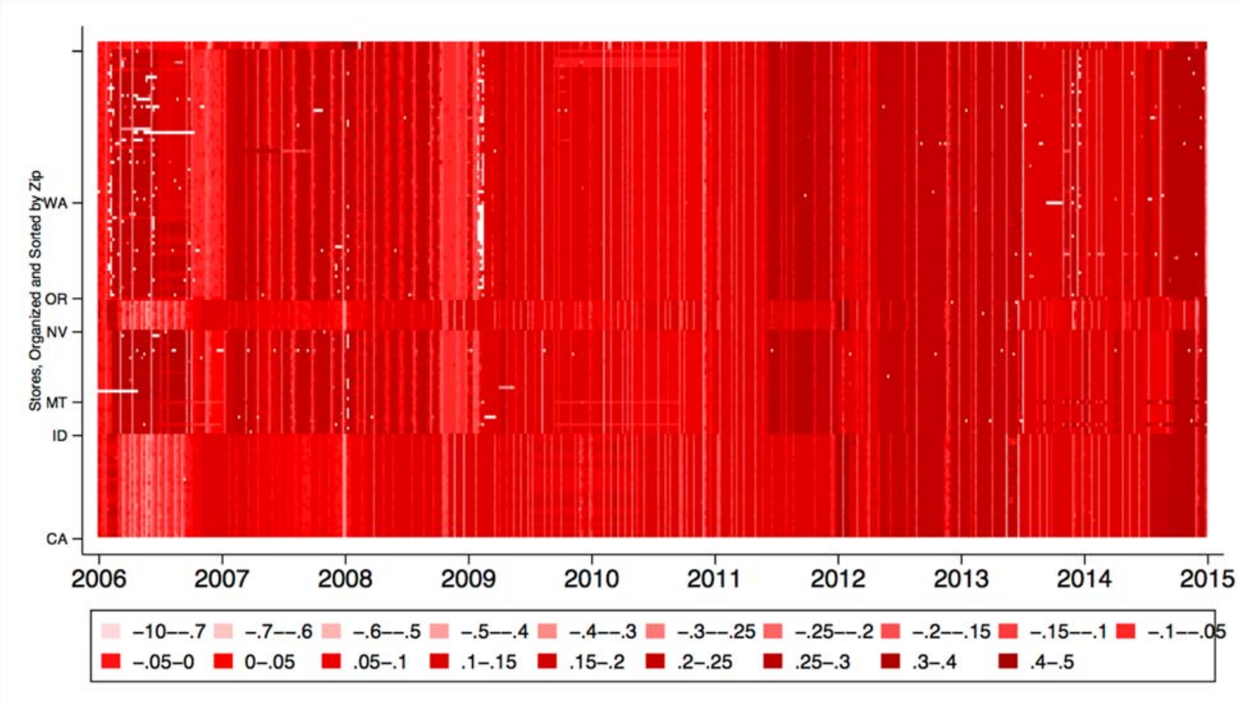

Online Appendix Figure 2b. Example of Chain with Geographic Pricing Blocks: Chain 32, Orange Juice

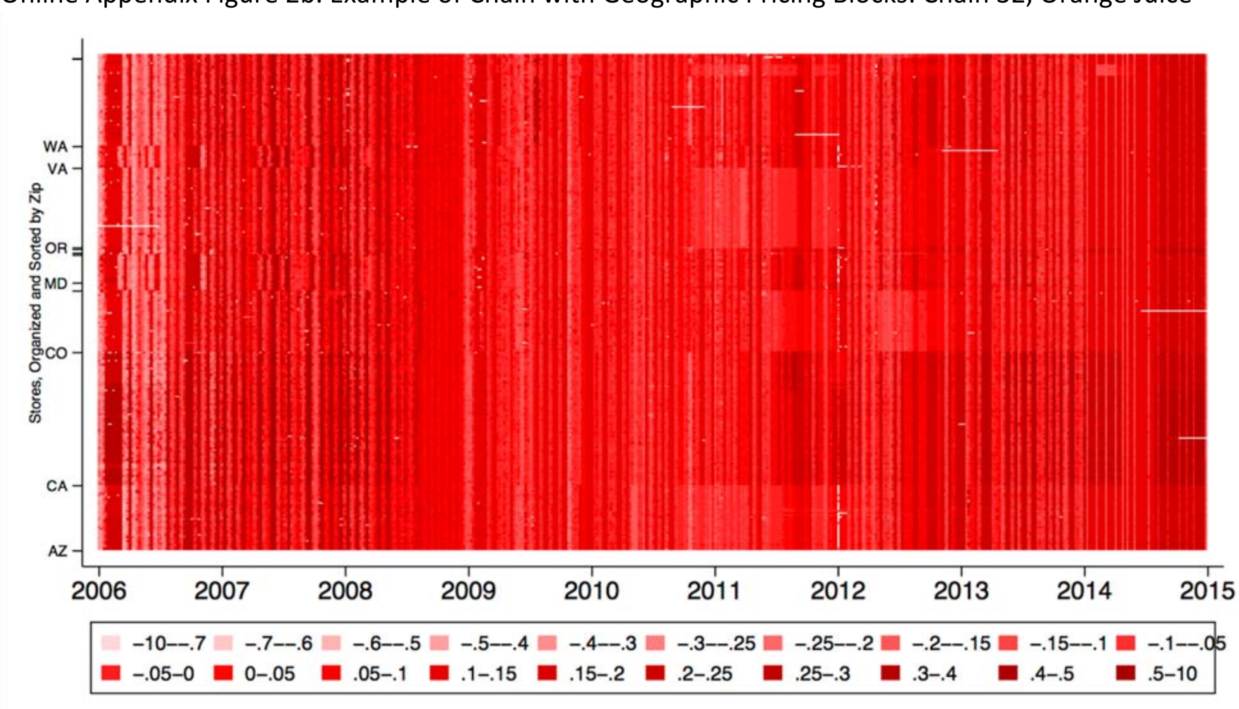

Notes: Plots depict demeaned log prices. Darker colors indicate higher price and are blank if price is missing. Each column is a week t. Each row is a store, and stores are sorted by three-digit zip code within each state divider.

Online Appendix Figure 3. Same as Figure 4 but within DMA

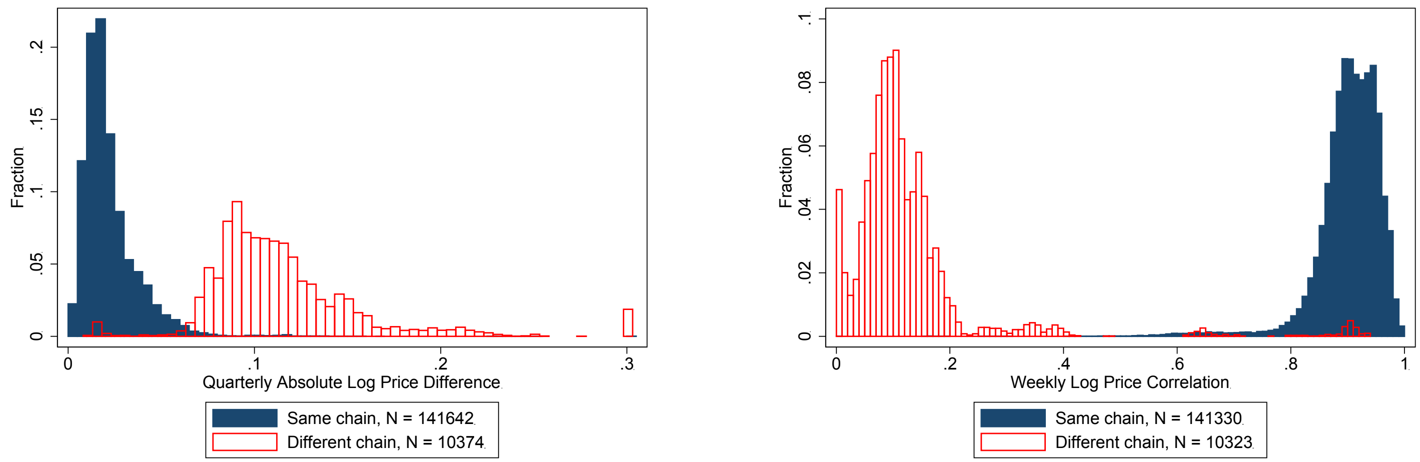

Notes: Each observation is a store-pair. "Same chain" mean same retailer_code. "Different chain" means both different retailer_code and different parent_code. Store pairs within a chain display markedly different pricing similarity compared to pairs in different chains. This relationship holds even when restricting the sample to pairs that should be the most similar, such as store pairs in the same DMA. Quarterly Absolute Log Price Differences are Winsorized at .3 and Weekly Correlation is Winsorized at 0. A maximum of 200 stores per chain are used in the same chain distributions to avoid overweighting the largest chains.

Online Appendix Figure 4: Pairs, Alternative Measure of Price Similarity: Weekly Absolute log price difference and Share Identical

Online Appendix Figure 4a-b. All pairs

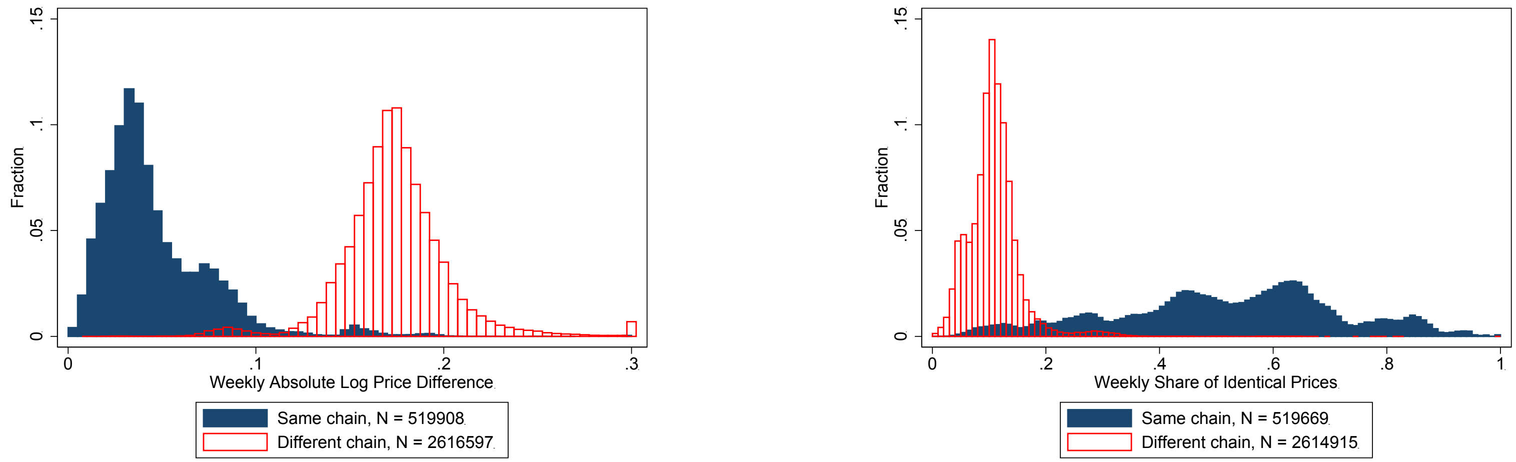

Online Appendix Figure 4c-d. Comparisons Across DMA and top third vs. bottom third of income only.

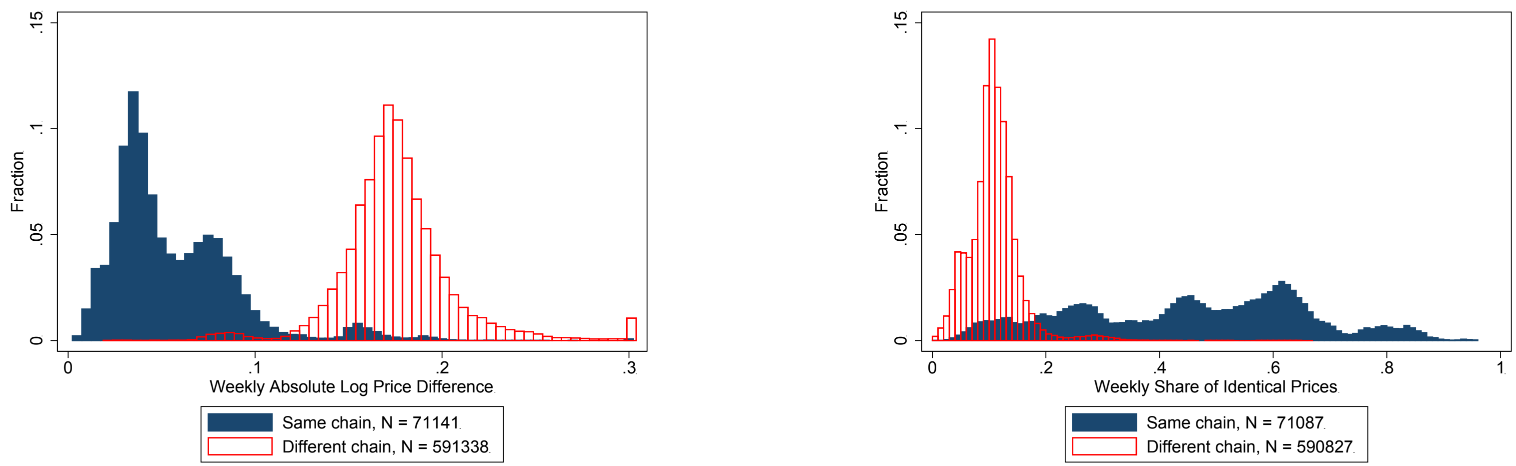

Notes: Each observation is a store-pair. Store pairs within a chain display markedly different pricing similarity compared to pairs in different chains. This relationship holds even when restricting the sample to pairs that should be the most differentiated, such as store pairs in different DMAs and in very different income areas (Panel b): even within chains, there should not be any advertising constraints and fairness should not be too large a concern. Quarterly Absolute Log Price Differences Winsorized at .3. A maximum of 200 stores per chain are used in the same chain histograms to avoid overweighting the largest chains.

Online Appendix Figure 5. Probability of Being in Same Chain as Function of Pricing Distance

Online Appendix Figure 5a. By Quarterly Absolute Log Price Difference

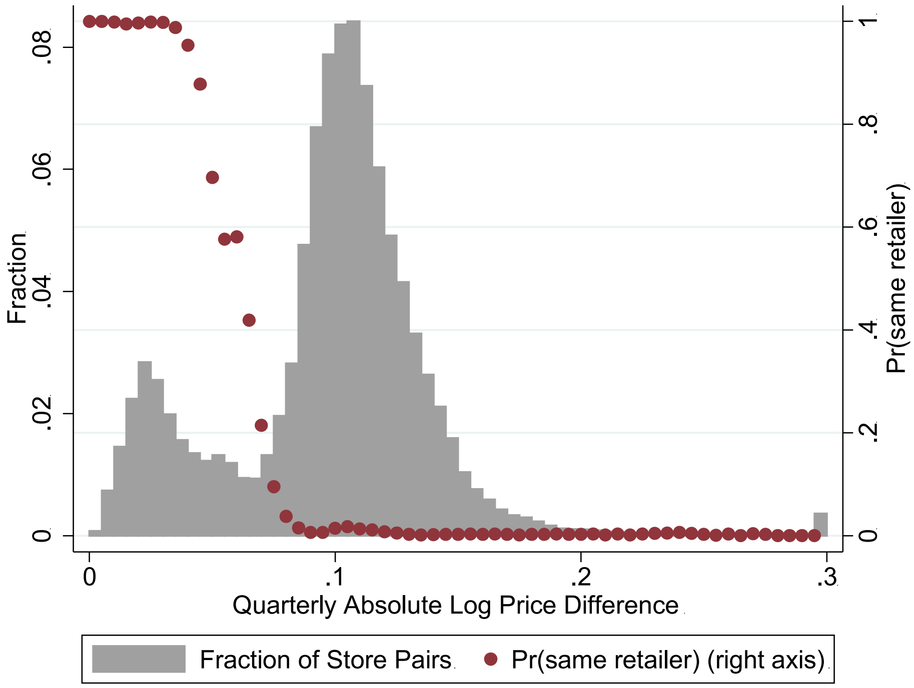

Online Appendix Figure 5b. By Correlation of Weekly Log Prices

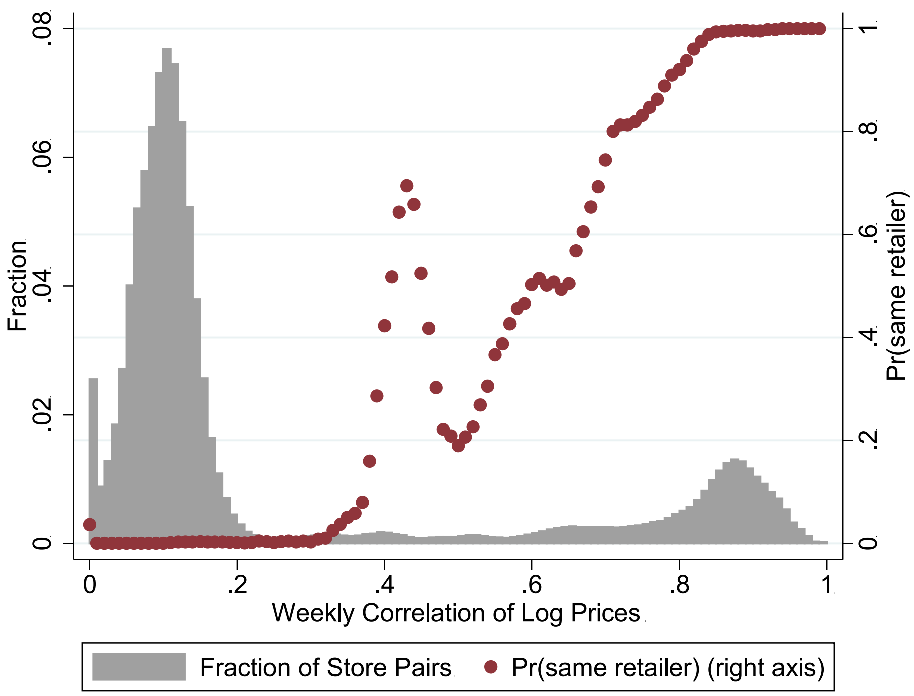

Notes: Empirically observed probability of being in the same retailer is plotted with the distribution of all pairs in the background (grey histogram). A maximum of 200 stores are used for within-chain pairs. The base rate of two stores being in the same chain is .1628 for quarterly absolute log price difference and .1628 for weekly correlation.

Online Appendix Figure 6. Between and Within-State Weekly Correlation of Log Prices

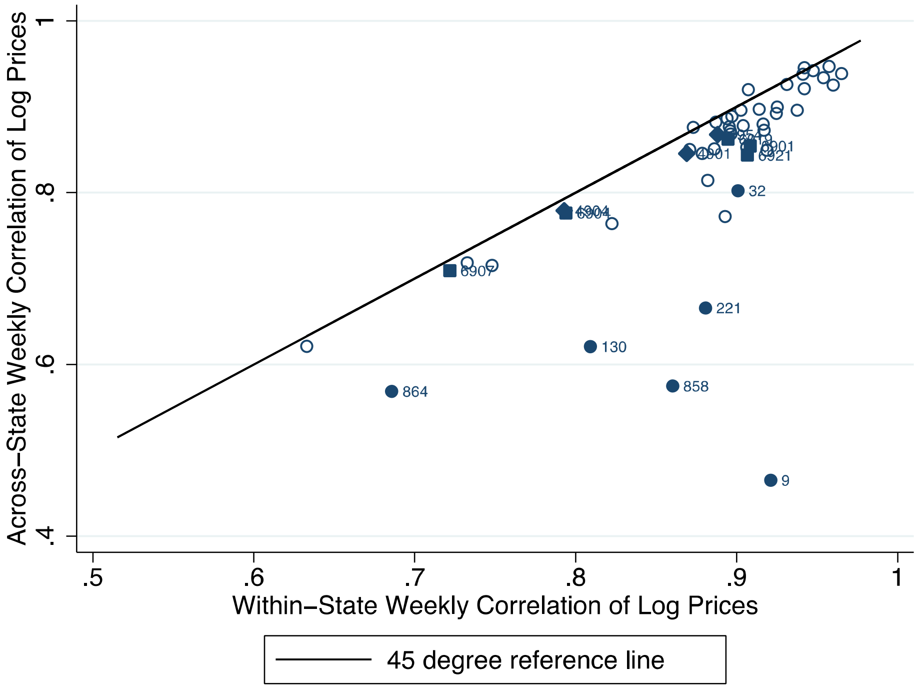

Notes: Each observation is a chain that operates at least three stores in multiple states. Chains of note are labeled.

Online Appendix Figure 7. Robustness of Key Fact, by Retailer. Quarterly Absolute Log Price Difference (Food stores only)

Onl. App. Fig. 7a. Vs. 20^{th} Availability Items

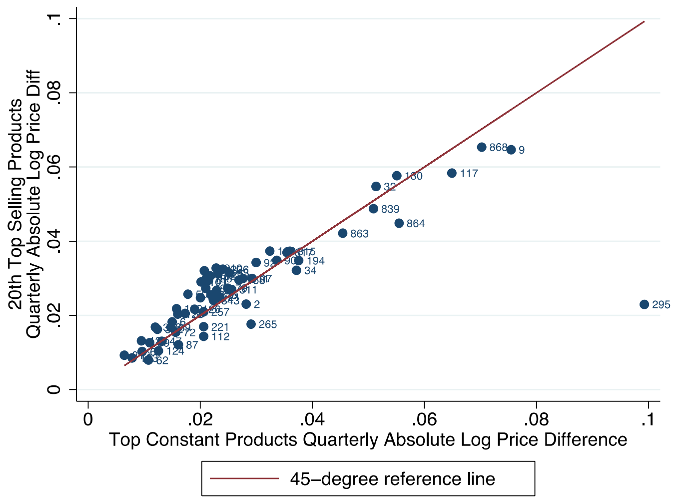

Onl. App. Fig. 7b. Vs. Top Selling Generic

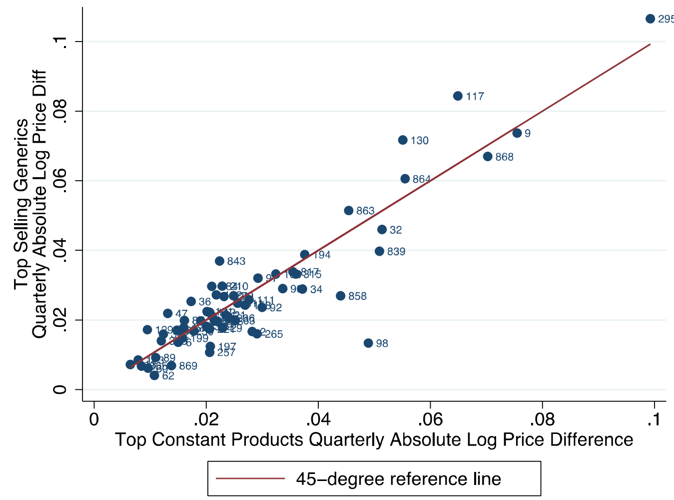

Onl. App. Fig. 7c. Vs. High Quality Unit-Price Items

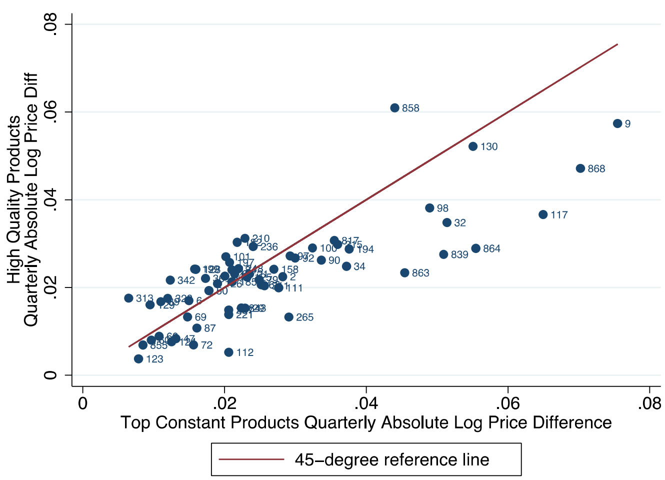

Notes: 20^{th} Availability and Top Generic products contain ten products, one from each module. High Quality products contain eight products, three each from Coffee and Cookies and two from Chocolate. A maximum of 400 stores are used per retailer.

Online Appendix Figure 8. Within-Chain Response of Prices to Income, By Chain

Online Appendix Figure 8a. Food Stores

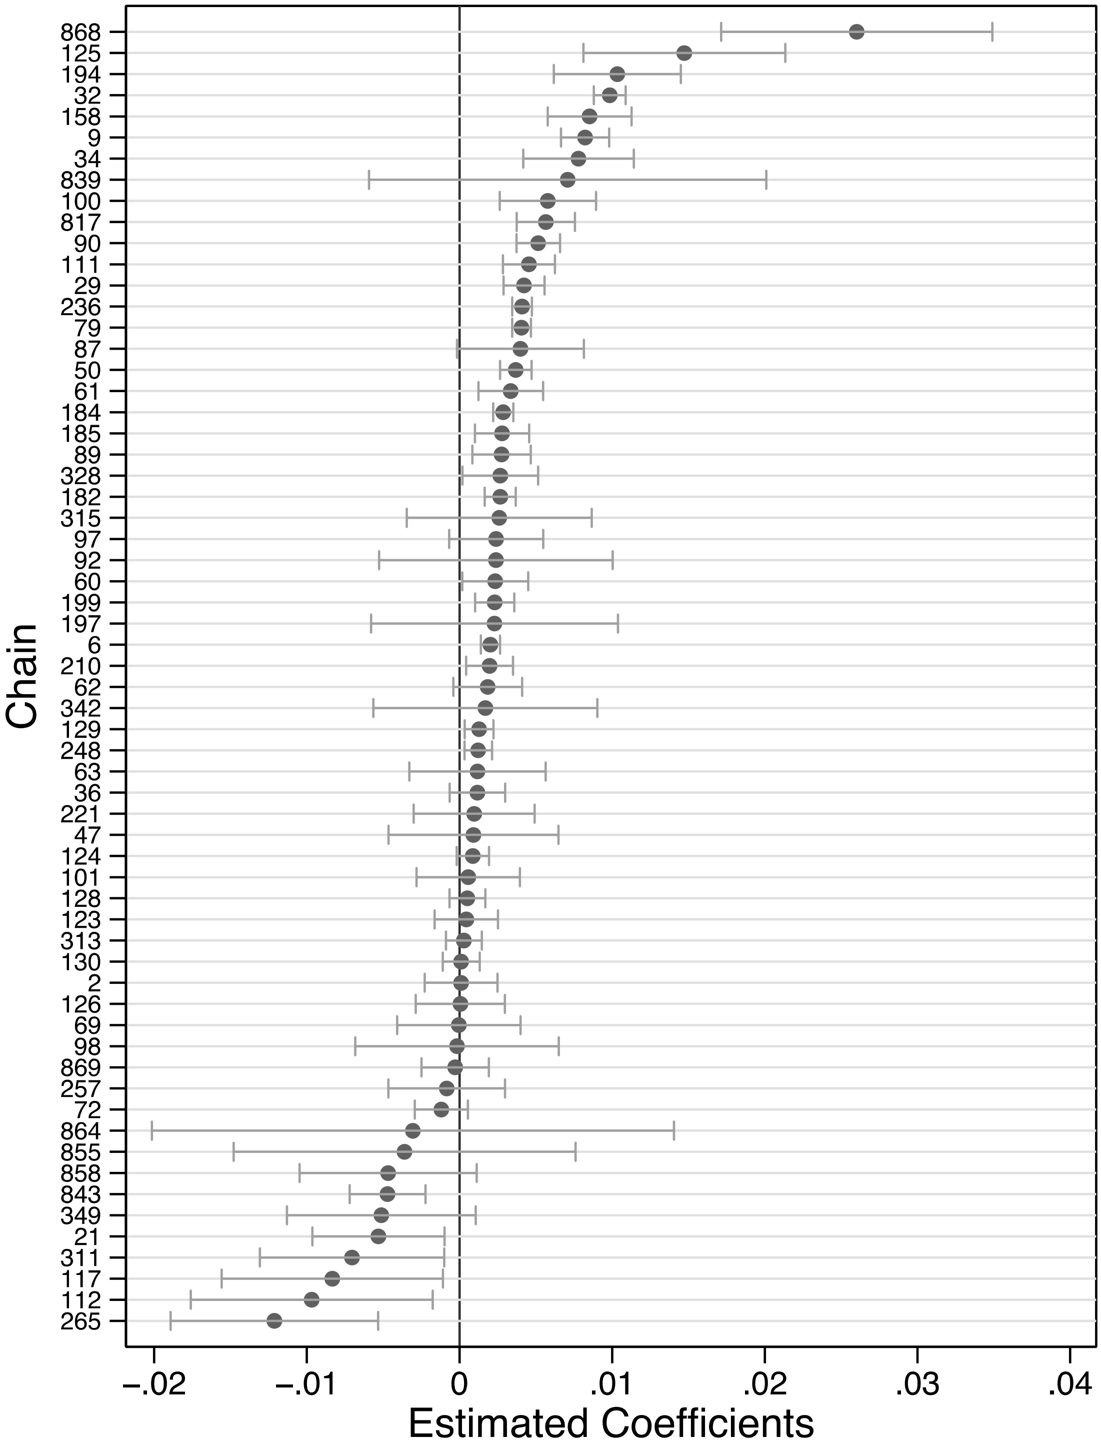

Notes: Plotted are the coefficients for independent store-level regressions of price on income (in $10,000s) for each chain and 95% confidence intervals based on robust standard errors. A coefficient of 0.01 means that within chain *c*, prices are set 1 log point (1%) higher for an increase in income of $10,000. Two chains with SE > .02 are omitted (one, Retailer 295, has coefficient (Robust SE) 0.015 (.0254)).

Online Appendix Figure 8b. Drugstores

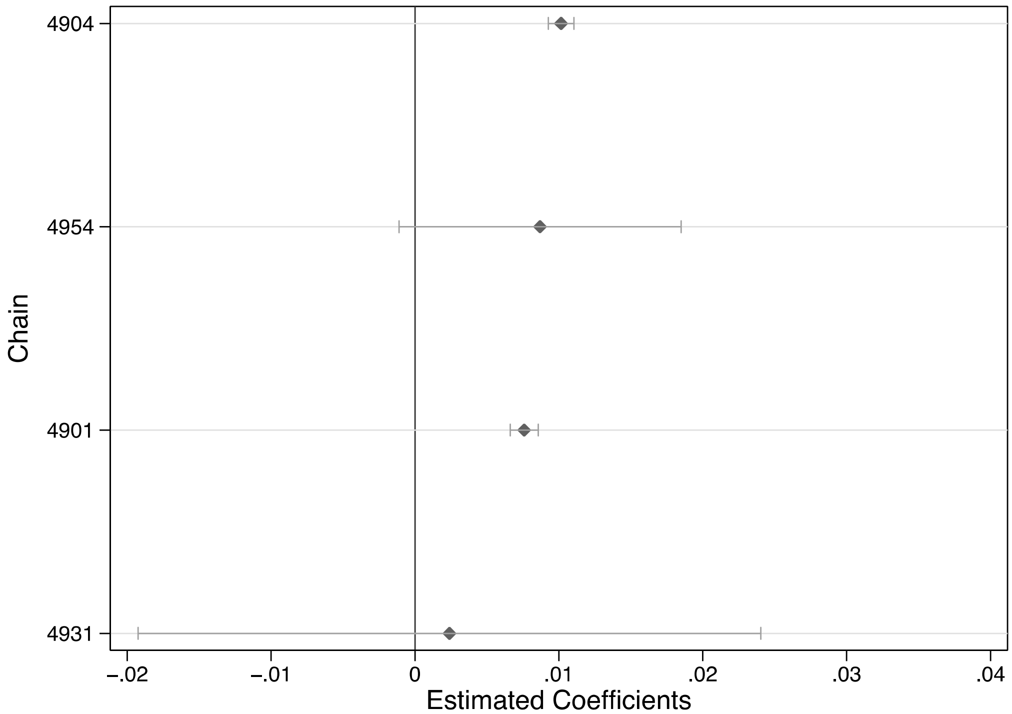

Online Appendix Figure 8c. Mass-Merchandise Stores

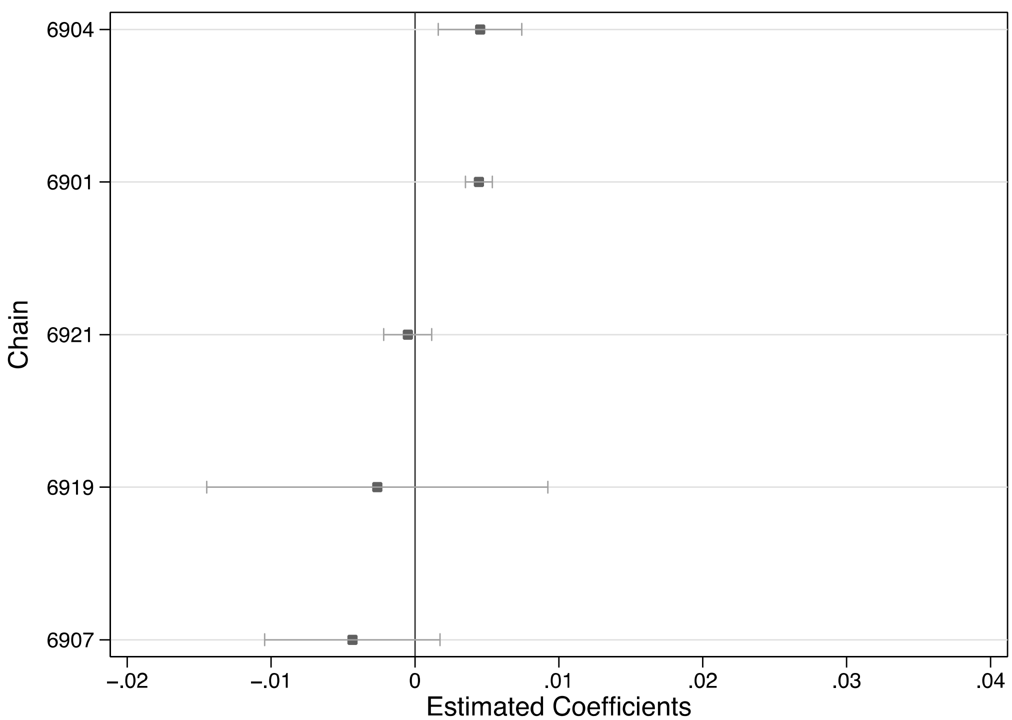

Notes: Plotted are the coefficients for independent store-level regressions of price on income (in $10,000s) for each chain and 95% confidence intervals based on robust standard errors. A coefficient of 0.01 means that within chain *c*, prices are set 1 log point (1%) higher for an increase in income of $10,000.

Online Appendix Figure 9. Price vs. Income by module/product group (Food stores only)

Online Appendix Figure 9a. Within-chain Price vs. Income Regression

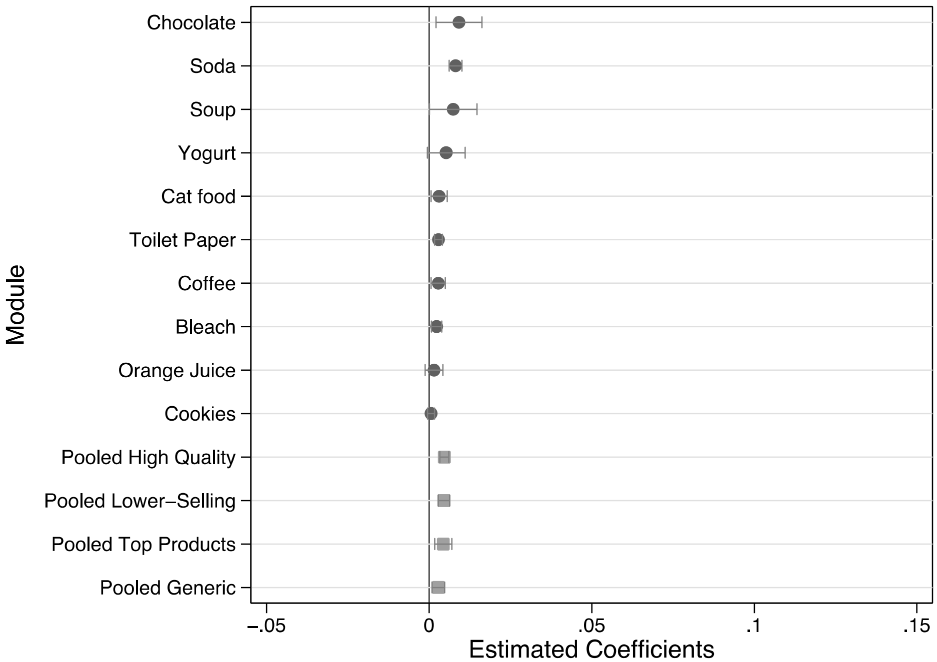

Online Appendix Figure 9b. Between-Chain Price vs. Income Regression (Food Stores Only)

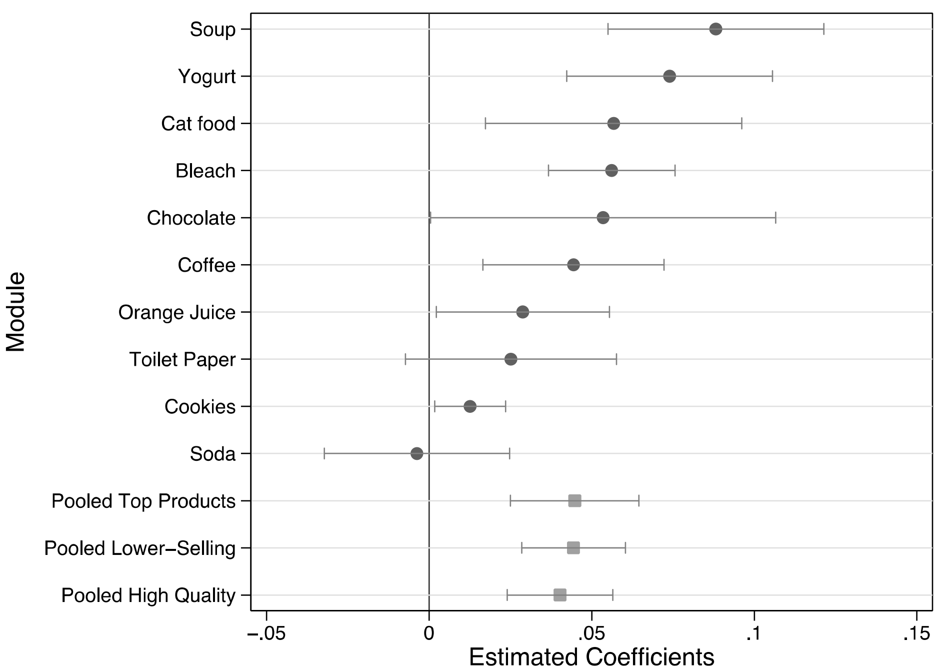

Notes: Online Appendix Figure 9a plots the coefficients for independent store-level regressions of price on income (in $10,000s) with chain fixed effects for each module and 95% confidence intervals based on standard errors clustered by parent_code. A coefficient of 0.01 for module *m* means that for stores within chain *c*, prices for module *m* are set 1 log point (1%) higher given an increase in income of $10,000. Figure 9b plots the same relationships but for chain averages using analytic weights equal to number of stores with standard errors again clustered by parent_code. A coefficient of 0.01 for module *m* means that chains set prices for module *m* 1 log point (1%) higher for an increase in average income of $10,000. Squares indicate pooled modules (averages of multiple products), while circles indicate individual products.

Online Appendix Figure 10. Between-Chain Relationship of Price to Income, Drug and Mass Merchandise Chains

Online Appendix Figure 10a. Drugstore Chains

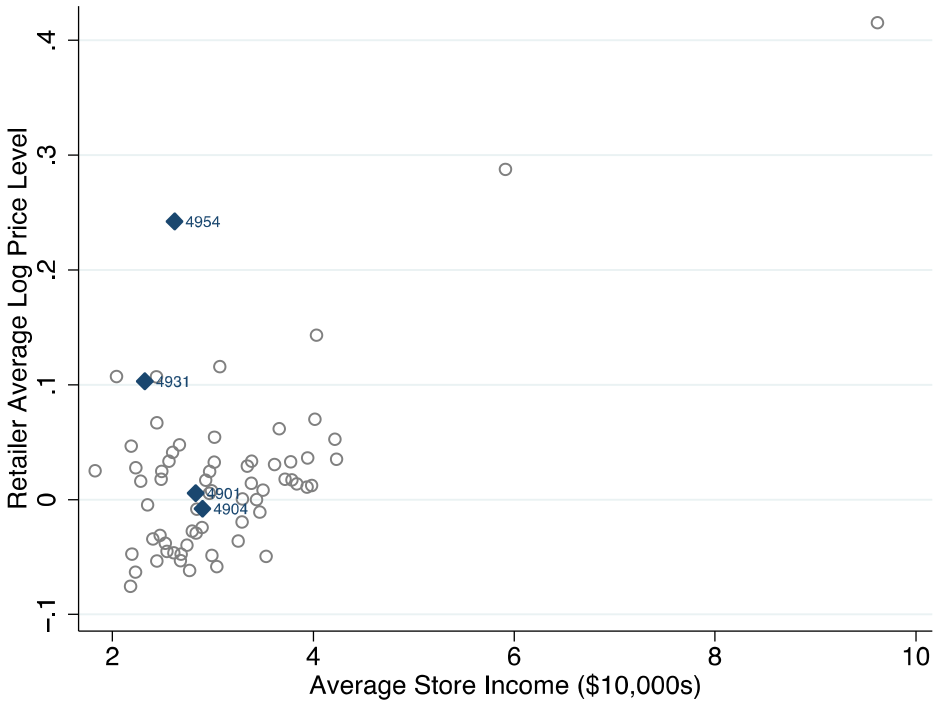

Online Appendix Figure 10b. Mass-Merchandise Chains

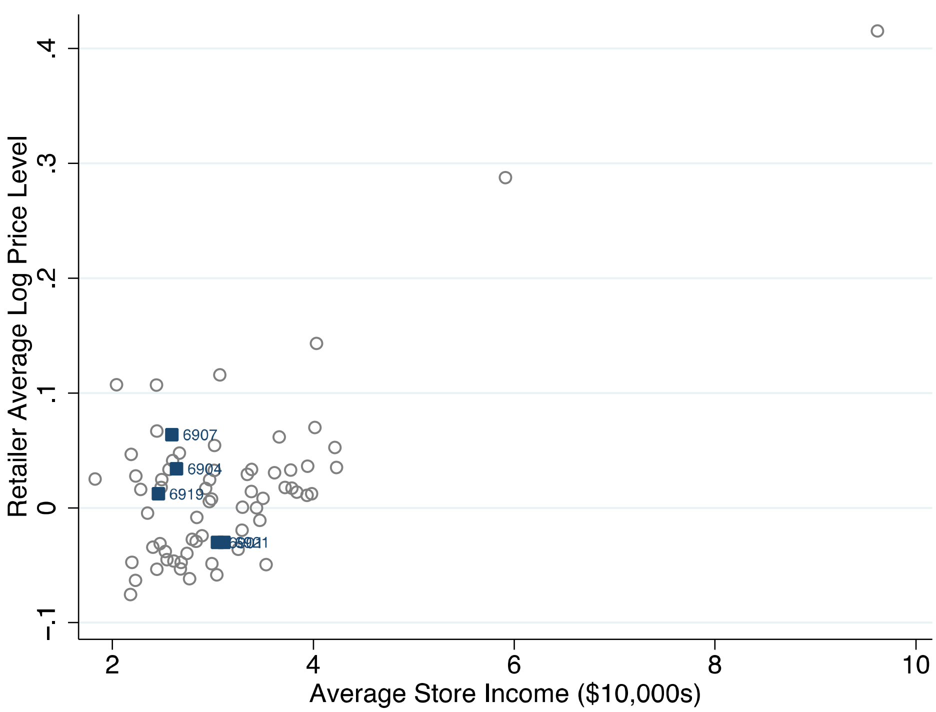

Notes: Hollow circles represent food store chains.

Online Appendix Figure 11. Price Response to Demographics, Results with Store-Level Education

O. A. Figure 11a. Price versus Education: Within-Chain

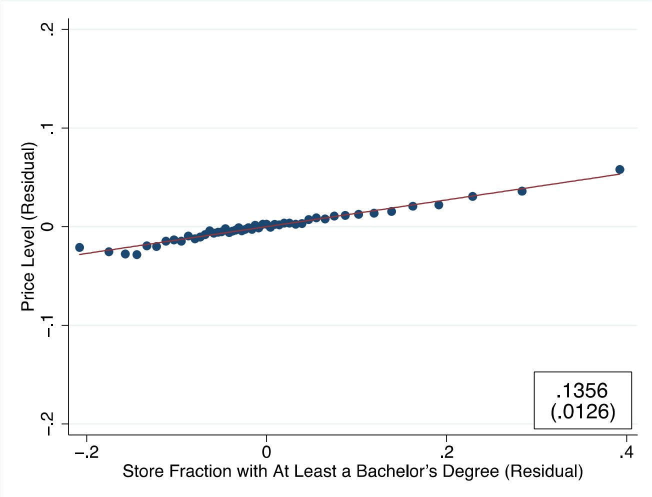

O. A. Figure 11b. Price versus Education: Between Chains (F only)

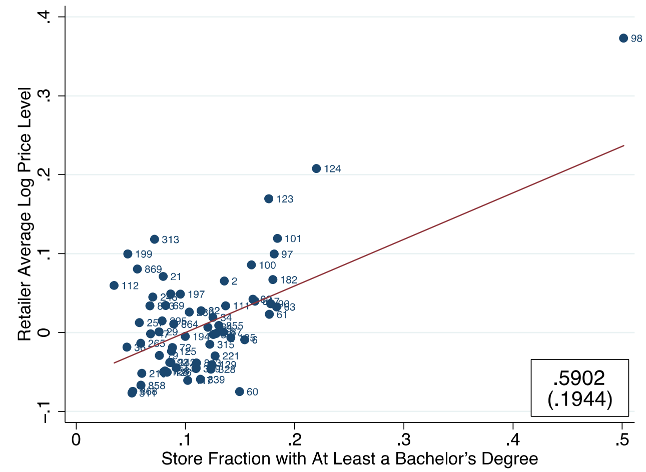

O. A. Figure 11c. Price versus Education: Within-Chain-State

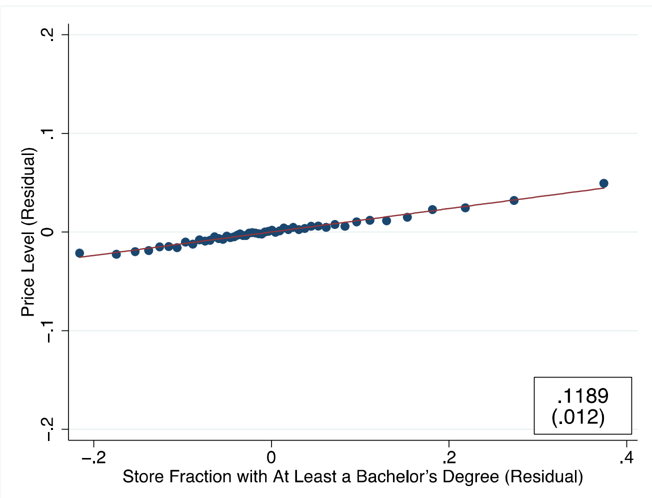

O. A. Figure 11d. Price versus Education: Between Chain-State

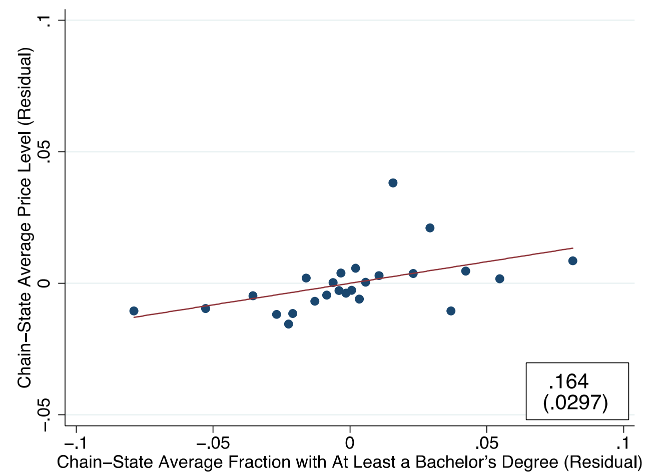

Notes: Standard errors clustered by parent_code. Axes ranges have been chosen to make the slopes visually comparable. Analytic weights equal to the number of stores in each aggregation unit are used in Figures 10c and 10d. In Figure 10a, residuals are after removing Chain FE. In Figure 10c, residuals are after removing ChainXState FE. In Figure 10b, labels indicate Chain. In Figure 10d., each observation is one of 25 bins representing 396 chain-states.

Online Appendix Figure 12. Price Response to Income: Estimated Nonsale Price Levels using Nielsen Data: Food Stores

Onl. App. Fig. 12a. Nonsale Prices: Within-Chain

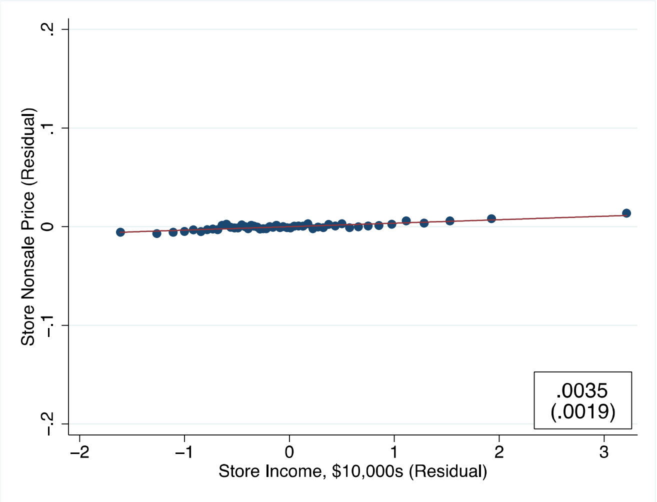

Onl. App. Fig. 12b. Nonsale Prices: Between-Chain Relationship

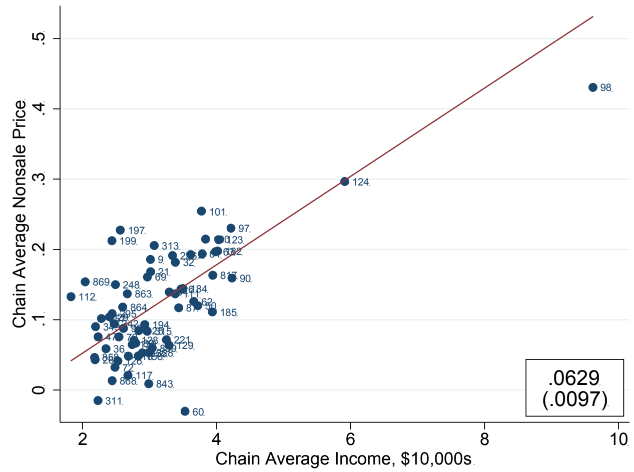

Onl. App. Fig. 12c. Nonsale Prices: Within-Chain-State

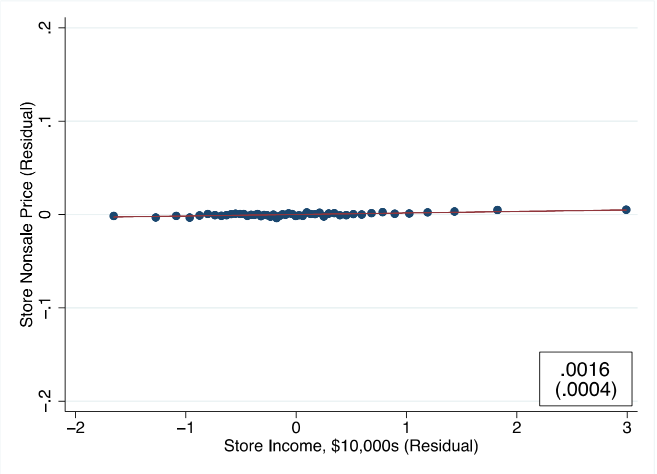

Onl. App. Fig. 12d. Nonsale Prices: Between-Zones Relationship

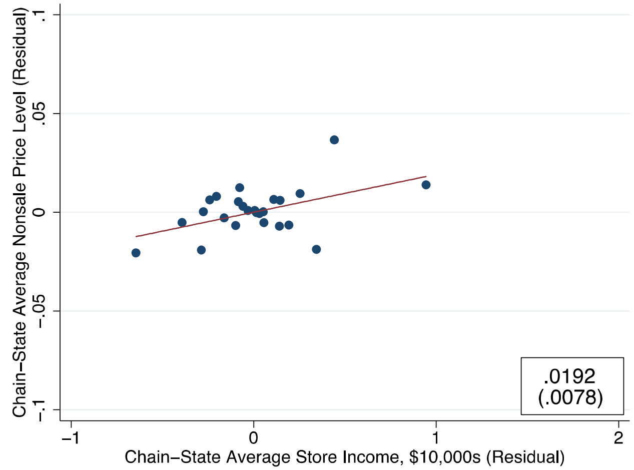

Notes: Standard errors clustered by parent_code. Axes ranges have been chosen to make the slopes visually comparable. Analytic weights equal to the number of stores in each aggregation unit are used in Figures 12b and 12d. In Figure 12a, residuals are after removing Chain FE. In Figure 12c, residuals are after removing ChainXState FE. In Figure 12b, labels indicate Chain. In Figure 12d., each observation is a bin of chain-states.

Online Appendix Figure 13. Price Response to Income: Estimated Nonsale Price Levels using Nielsen Data: Drugstores

Onl. App. Fig. 13a. Nonsale Prices: Within-Chain

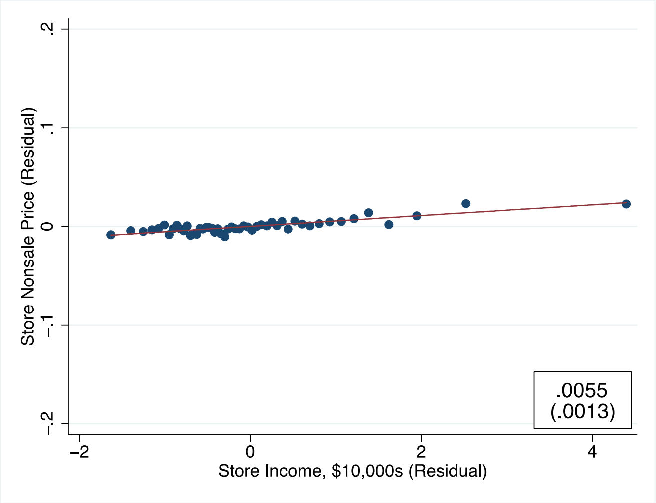

Onl. App. Fig. 13b. Nonsale Prices: Within-Chain-State

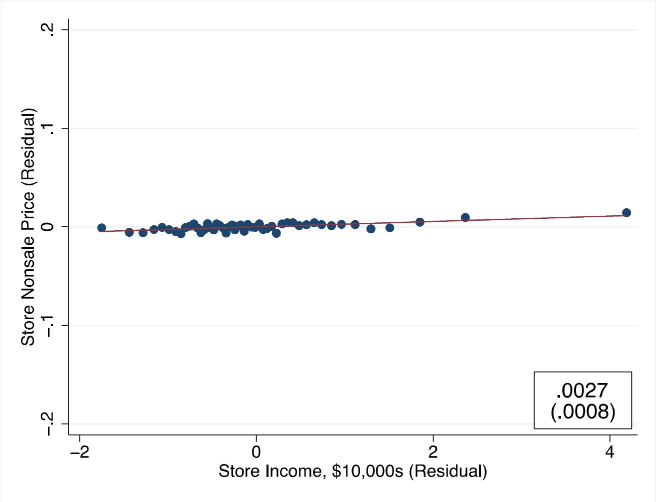

Onl. App. Fig. 13c. Nonsale Prices: Between-Zones Relationship

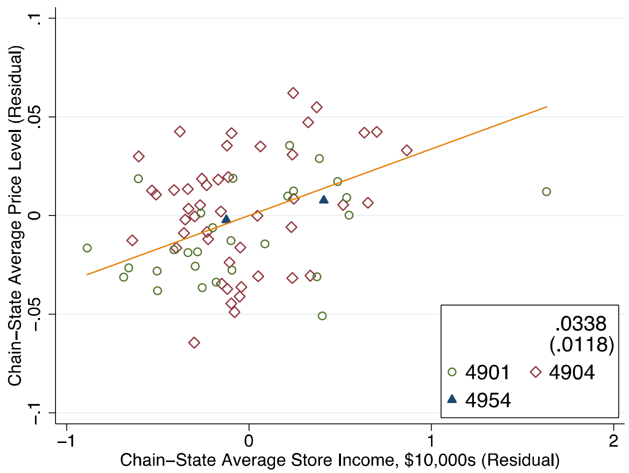

Notes: Standard errors clustered by parent_code*state. Axes ranges have been chosen to make the slopes visually comparable. Analytic weights equal to the number of stores in each aggregation unit are used in Figure 13d. In Figure 13a, residuals are after removing Chain FE. In Figure 13c, residuals are after removing ChainXState FE. In Figure 13d., each observation is a chain-state.

Online Appendix Figure 14. Price Response to Income: Estimated Nonsale Price Levels using Nielsen Data: Mass-Merchandise Stores

Onl. App. Fig. 14a. Nonsale Prices: Within-Chain

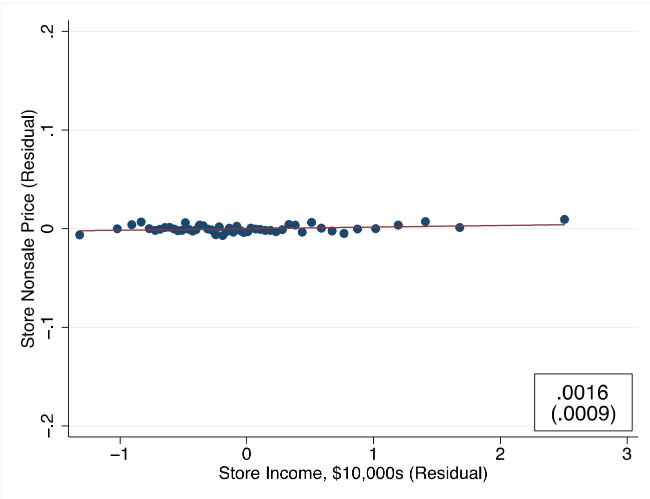

Onl. App. Fig. 14b. Nonsale Prices: Within-Chain-State

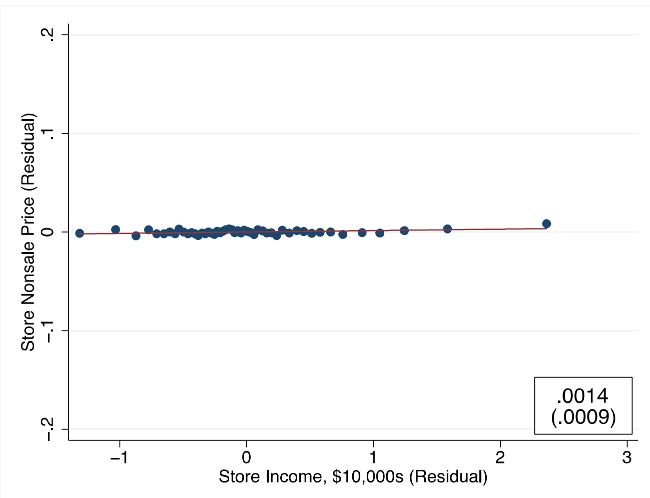

Onl. App. Fig. 14c. Nonsale Prices: Between-Zones Relationship

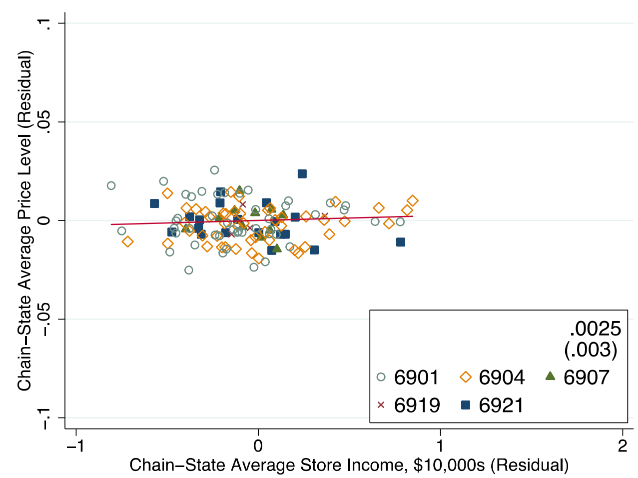

Notes: Standard errors clustered by parent_code*state. Axes ranges have been chosen to make the slopes visually comparable. Analytic weights equal to the number of stores in each aggregation unit are used in Figure 14d. In Figure 14a, residuals are after removing Chain FE. In Figure 14c, residuals are after removing ChainXState FE. In Figure 14d., each observation is a chain-state.

Online Appendix Figure 15. Shrinkage of Store-Level Elasticity

Online Appendix Figure 15a. Food Stores

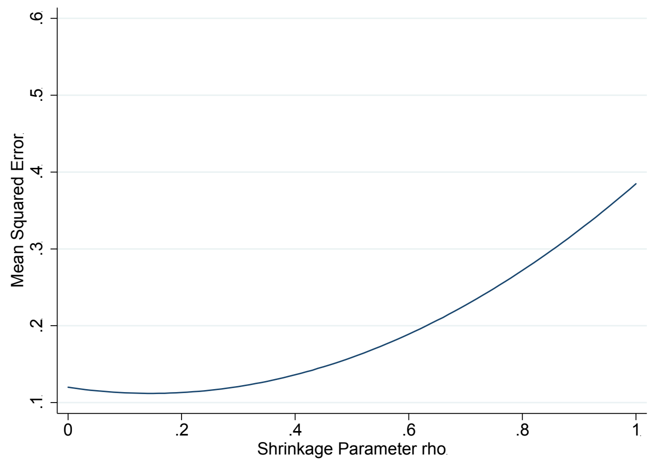

Online Appendix Figure 15b. Drugstores

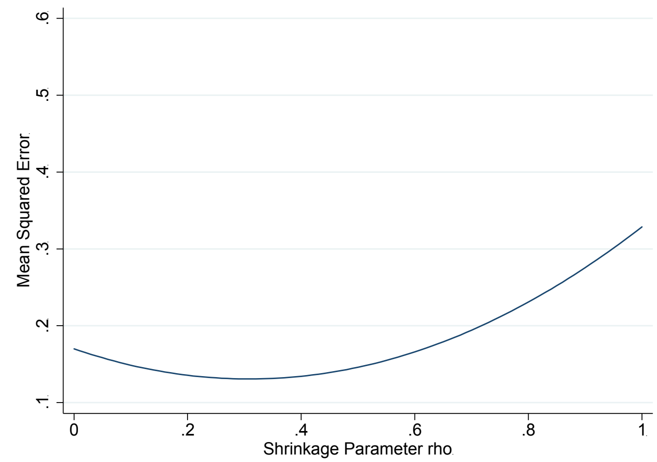

Online Appendix Figure 15c. Mass-Merchandise Stores

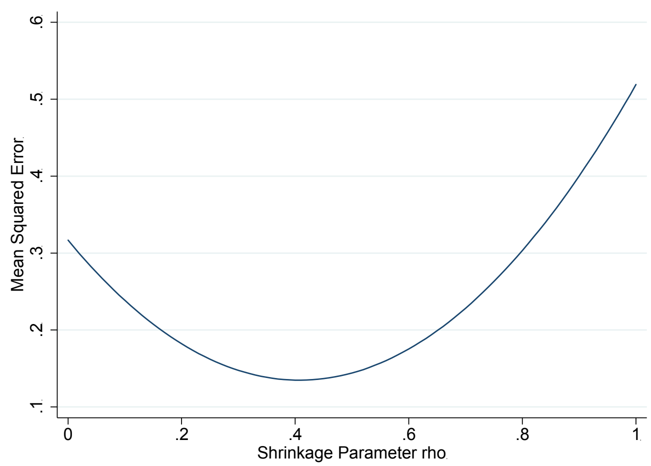

Notes: For each store type, pooled elasticities using prices and quantities from the first half of each year only are shrunk to the mean and compared to pooled elasticities using only the prices and quantities from the second half of each year. A value of zero indicates that no shrinkage is performed, and a value of one indicates that all elasticities are replaced by the mean elasticity for the first half of the year.

Online Appendix Figure 16. Additional Validation for Elasticity

Online Appendix Figure 16a. Test of Linearity: Pooled, Decomposed

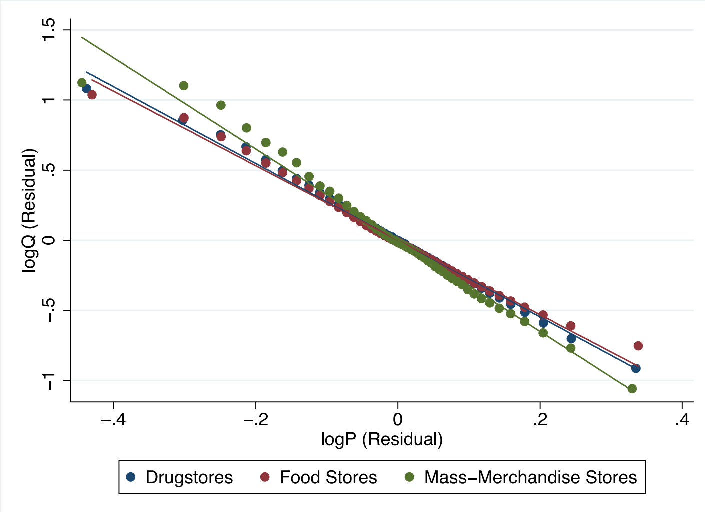

Online Appendix Figure 16b. Test of Linearity: Soda, Decomposed

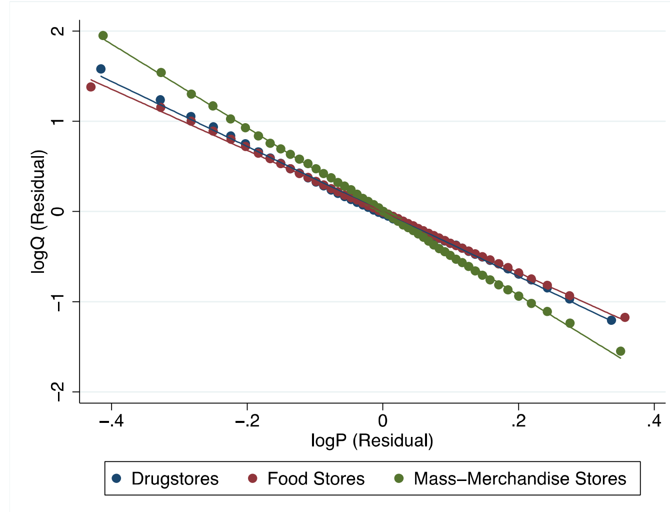

Online App. Figure 16c. Food Store Incremental R-squared

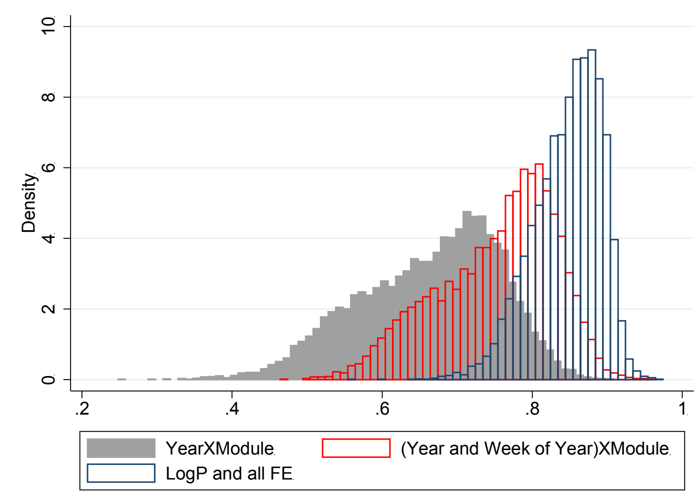

Online App. Figure 16d. Food Store Stockpiling Evidence: Lags and Leads

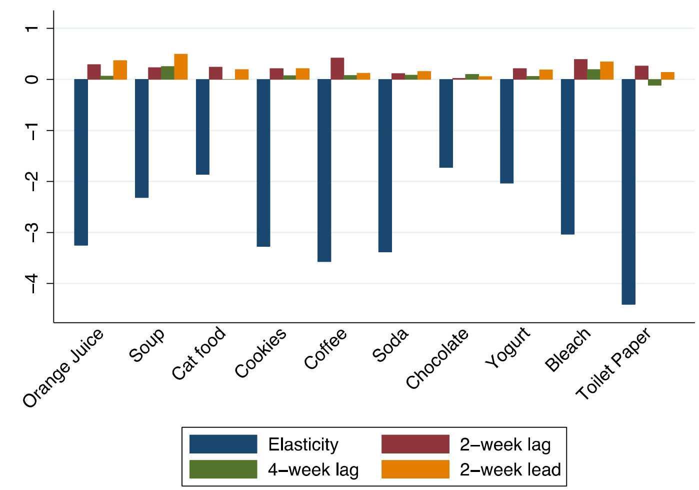

Notes: Figure 16a shows the residuals after removing yearXmodule and (week of year)* module FE for all products. Figure 16b shows the residuals after removing year and week of year FE for soda only. Each observation is a store-week of the product indicated. Figure 16c shows the distribution of R-squared using the pooled regression of all 10 products in food stores only. In Figure 16d., all food stores are regressed together with store, yearXmodule, and week-of-yearXmodule fixed effects, unlike our main specification which consists of store-by-store regressions.

Online Appendix Figure 17. Price versus Log Elasticity: Food Stores

Onl. App. Figure 17a. Within-Chain: Data vs. Log Elasticity

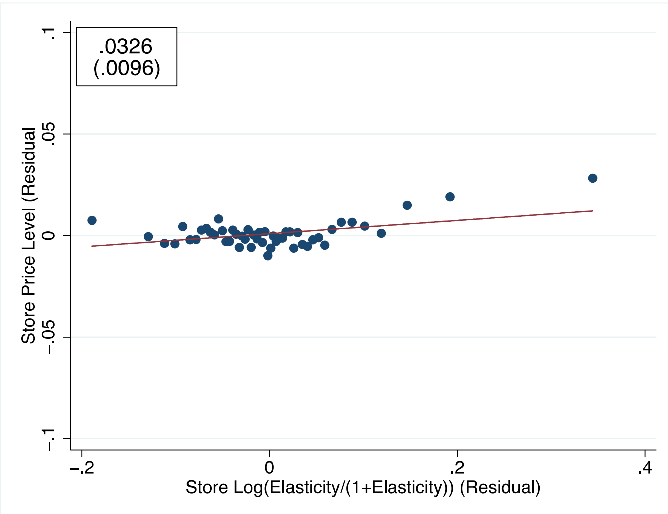

Onl. App. Figure 17b. Between-Chain: Data vs. Log Elasticity

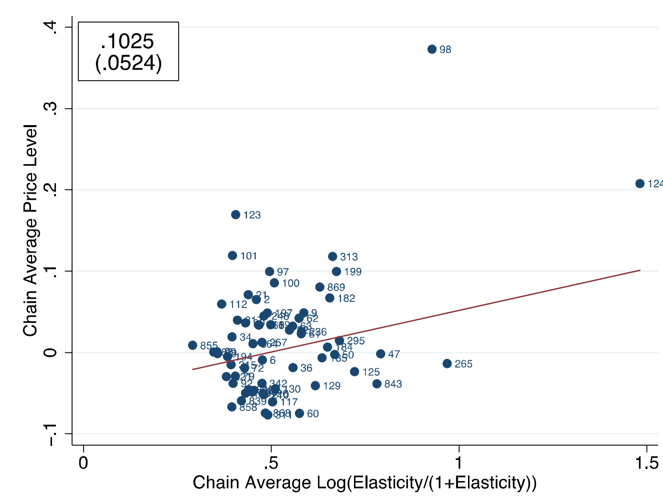

Onl. App. Figure 17c. Within-Chain-State: Data vs. Log Elasticity

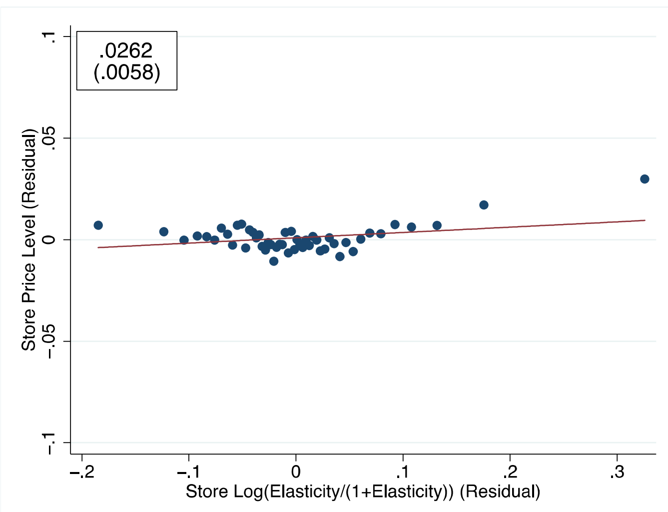

Onl. App. Figure 17d. Between-State, Within-Chain: Data vs. Log Elasticity

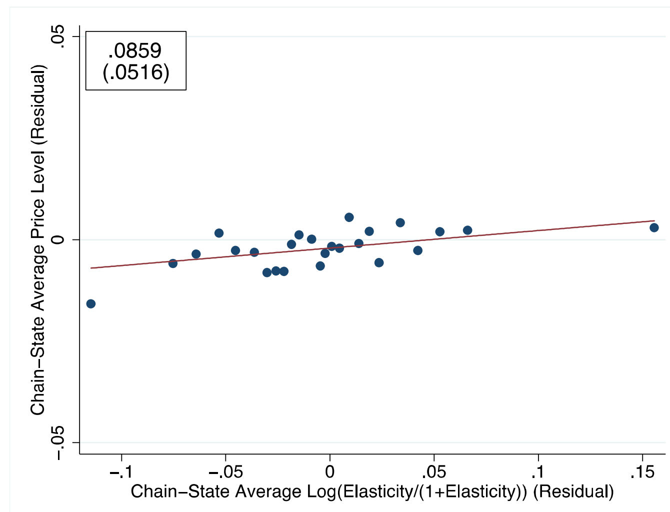

Notes: Standard errors are clustered by parent_code. Elasticities are Winsorized at -1.2. For Figures 17b and 17d, the values are the average Log(Elasticity/(1+Elasticity)). The solid red lines are the line of best fit. The dashed green line shows the model predicted price. Axis ranges were chosen so that slopes are visually comparable.

Online Appendix Figure 18. Price versus Log Elasticity: Drugstores

Onl. App. Figure 18a. Within-Chain: Data vs. Log Elasticity

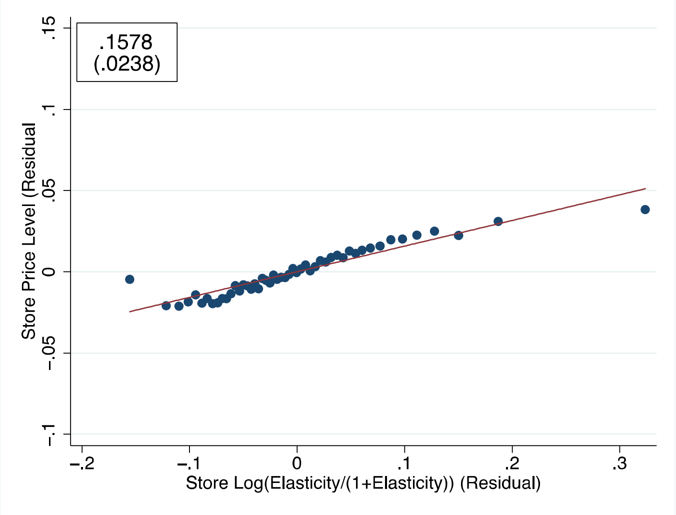

Onl. App. Figure 18c. Within-Chain-State: Data vs. Log Elasticity

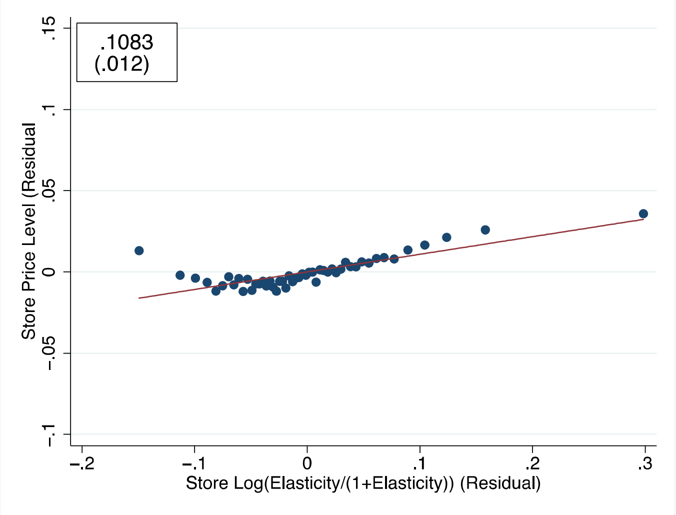

Onl. App. Figure 18d. Between-State, Within-Chain: Data vs. Log Elasticity

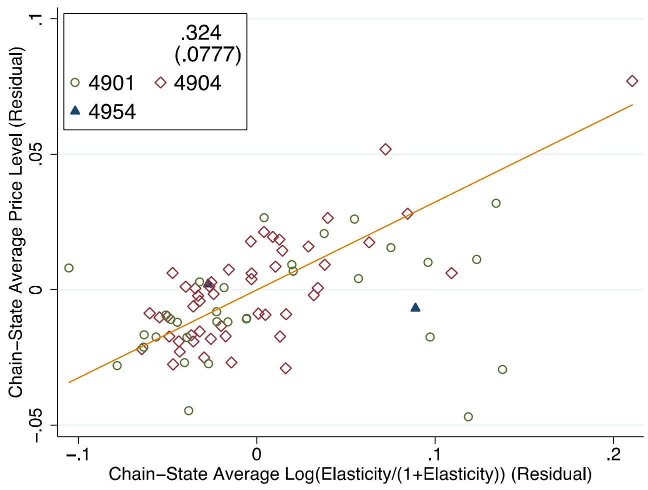

Notes: Standard errors are clustered by parent_code*state. Elasticities are Winsorized at -1.2. For Figure 18d, the values are the average Log(Elasticity/(1+Elasticity)). The solid red lines are the line of best fit. The dashed green line shows the model predicted price. Axis ranges were chosen so that slopes are visually comparable.

Online Appendix Figure 19. Price versus Log Elasticity: Mass-Merchandise Stores

Onl. App. Figure 19a. Within-Chain: Data vs. Log Elasticity

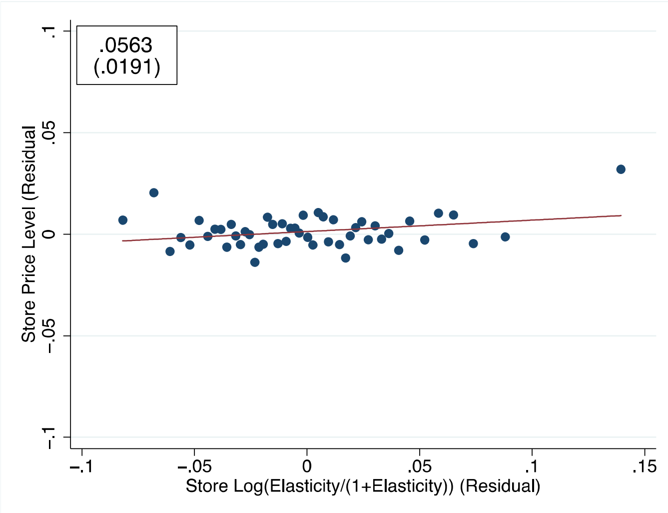

Onl. App. Figure 19c. Within-Chain-State: Data vs. Log Elasticity

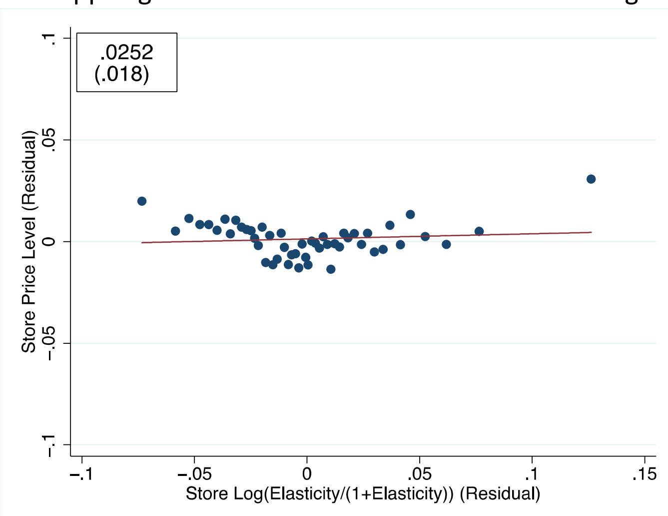

Onl. App. Figure 19d. Between-State, Within-Chain: Data vs. Log Elasticity

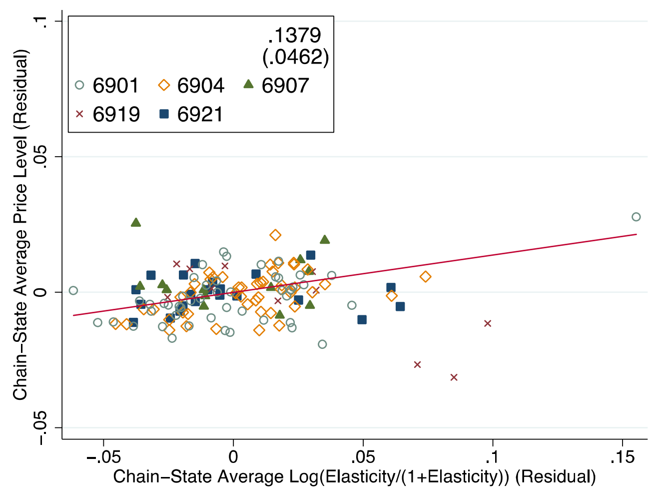

Notes: Standard errors are clustered by parent_code*state. Elasticities are Winsorized at -1.2. For Figure 19d, the values are the average Log(Elasticity/(1+Elasticity)). The solid red lines are the line of best fit. The dashed green line shows the model predicted price. Axis ranges were chosen so that slopes are visually comparable.

Online Appendix Figure 20. Robustness of Result on Price vs. Income, Different Products (Food Stores only)

Online Appendix Figure 20a-b. Top product relationships in Food Stores

Top Seller Price vs. Income, within chain

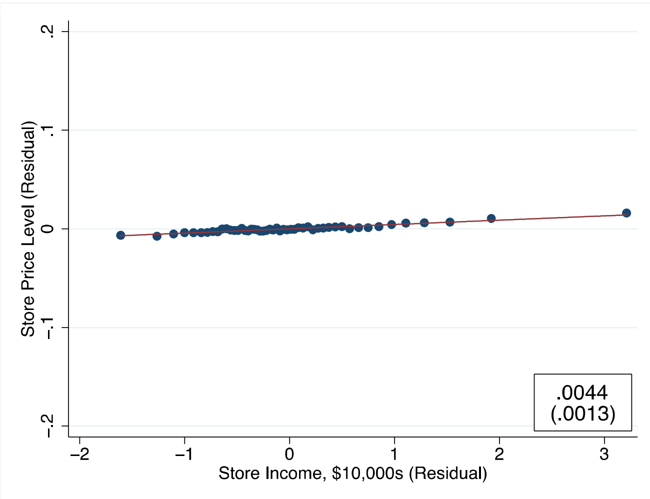

Top Seller Price vs. Income, between chain

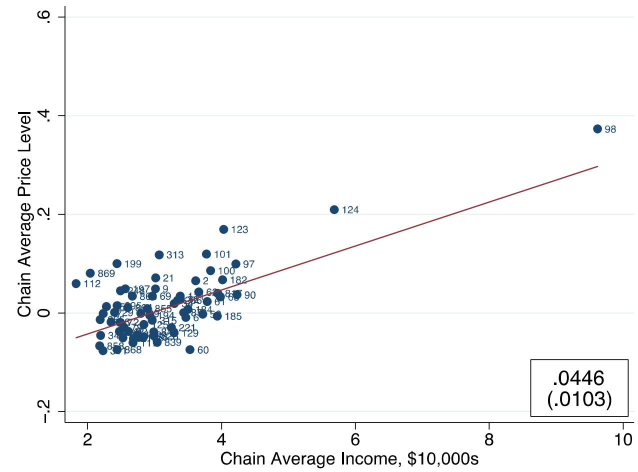

Online Appendix Figure 20c-d. Lower-Selling Products (20th Highest Selling)

20th Seller Price vs. Income, within chain

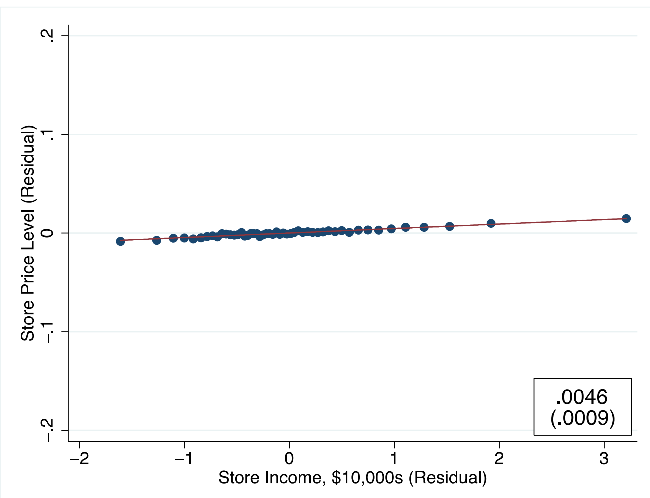

20th Seller Price vs. Income, between chain

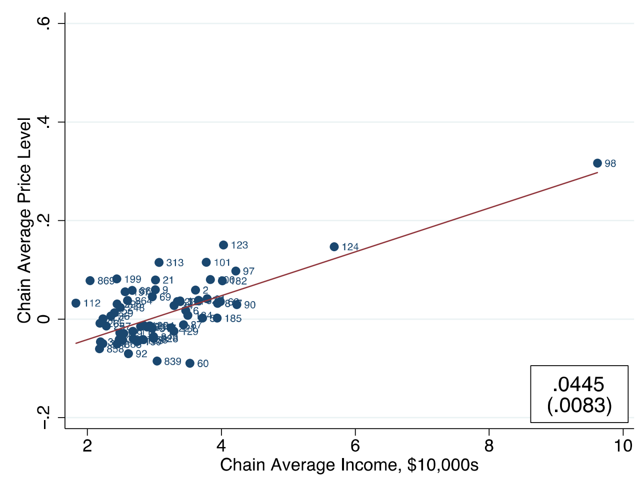

Notes: These robustness checks use the price level for the alternate products indicated. Residuals are after removing Chain FE. Standard errors are clustered by parent_code.

Online Appendix Figure 20e. Top-Selling Generic Product

Top Generic Product vs Income, within chain

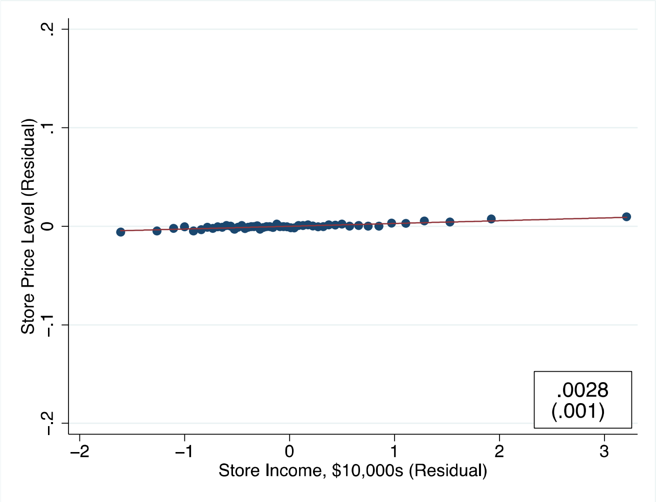

Online Appendix Figure 20f-g. Similar Generic Product Across Chains

Similar Generic Product, within chain

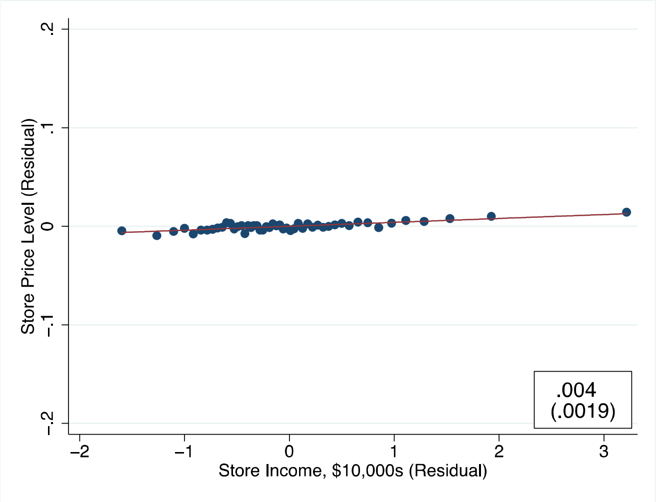

Similar Generic Product, between chain

Notes: Residuals are after removing Chain FE. Standard errors are clustered by parent_code except for Figure 20d., where they are robust. Figure 20d. does not impose any minimum number of modules to be present for a store to be included. Changing this does not affect the within-chain relationship much but steepens the between-chain relationship.

Online Appendix Figure 20h. Index Price Level

Index Price Level, within chain

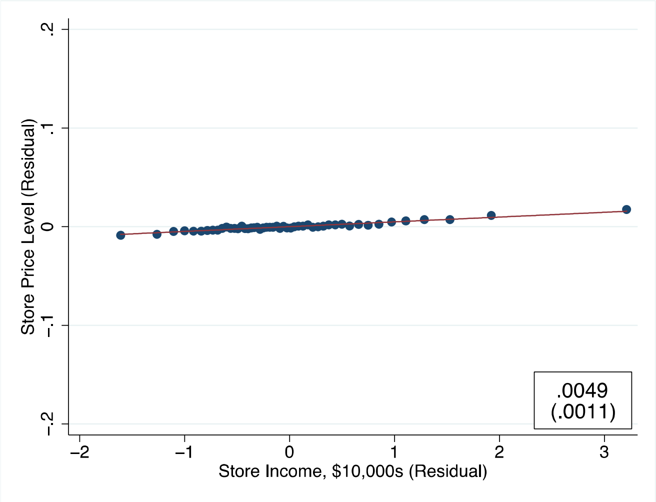

Notes: Residuals are after removing Chain FE. Standard errors are clustered by parent_code. Figure 20h is the only specification that uses more than one product per module. Since module-level indices are constructed at the chain level, between-chain comparisons are not possible.

Online Appendix Figure 21. Quarterly Elasticity Estimates and Validation (Food Stores Only)

Online Appendix Figure 21a. Quarterly Elasticity Estimates

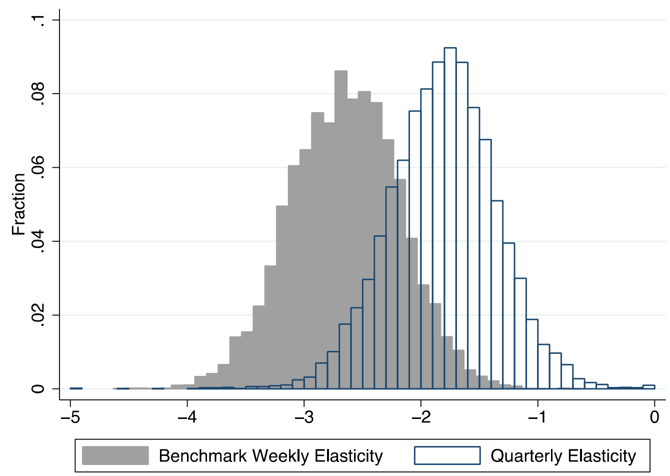

Online Appendix Figure 21b. Correlation with Benchmark Elasticity

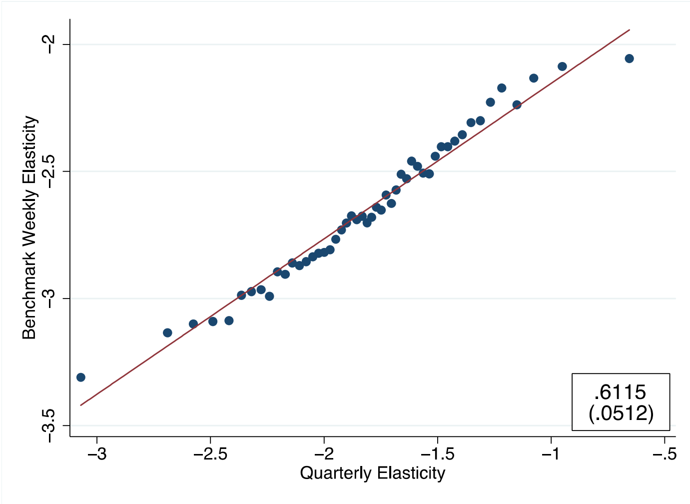

Online Appendix Figure 21c. Test of Linearity

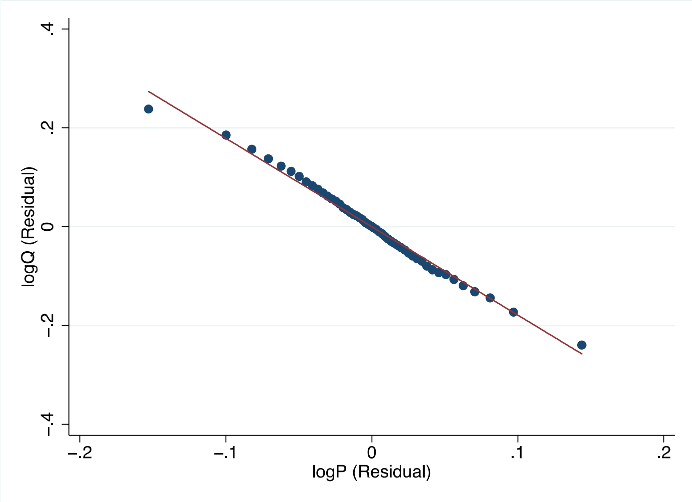

Online Appendix Figure 21d. Distribution of Standard Errors

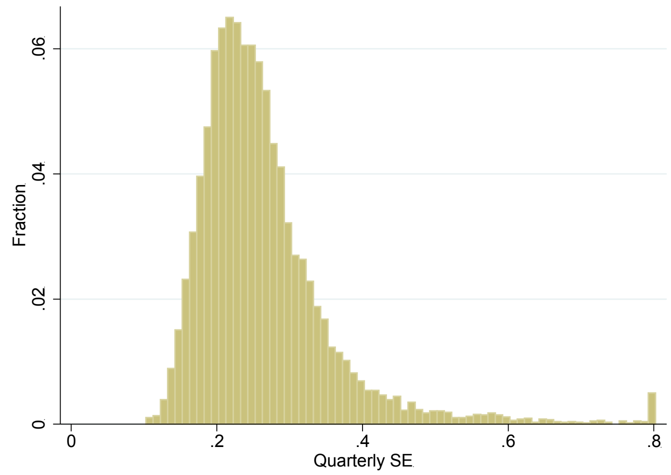

Notes: In Figure 21a., elasticities are Winsorized at 0 and -5. In Figure 14b., there are 50 bins representing 9,415 stores. Standard errors are clustered by parent_code. In Figure 21c., there are 50 bins representing 3,488,542 store-quarter-modules. Residuals are after removing module*quarter-of-year and module*year fixed effects. In Figure 21d., standard errors are clustered by quarter and are Winsorized at .8.

Online Appendix Figure 21e. Correlation with Income

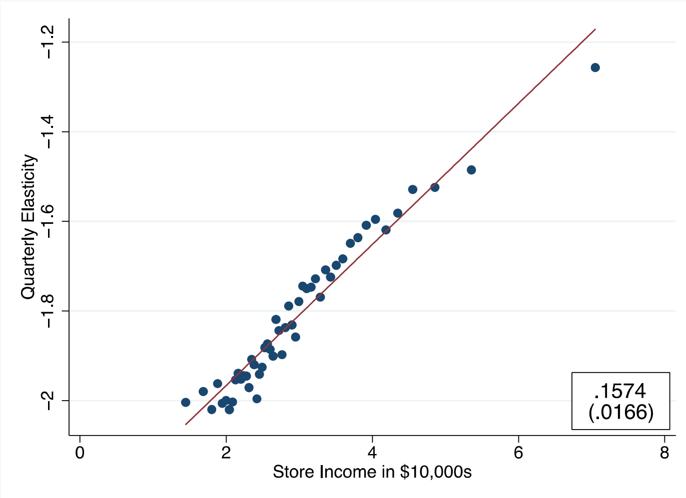

Notes: Standard errors are clustered by parent_code.

Online Appendix Figure 22. Weekly Price Index Elasticity Estimates and Validation (Food Stores Only)

Online Appendix Figure 22a. Index Elasticity Estimates

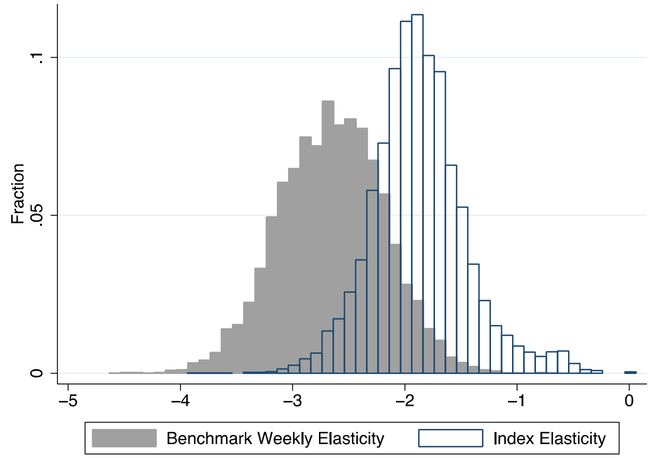

Online Appendix Figure 22b. Relation to Benchmark elasticity

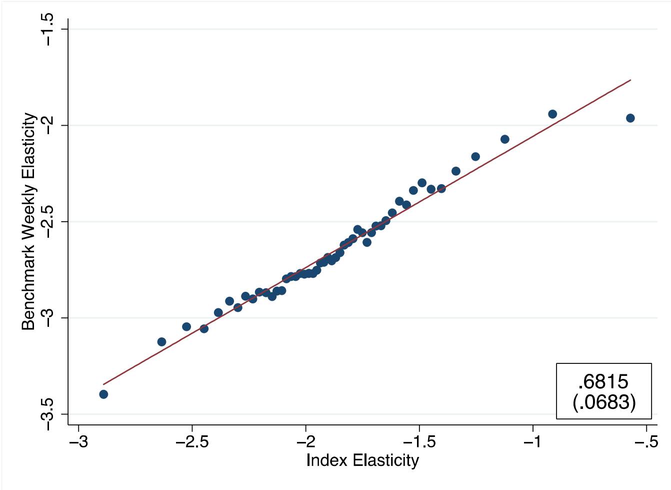

Online Appendix Figure 22c. Test of Linearity

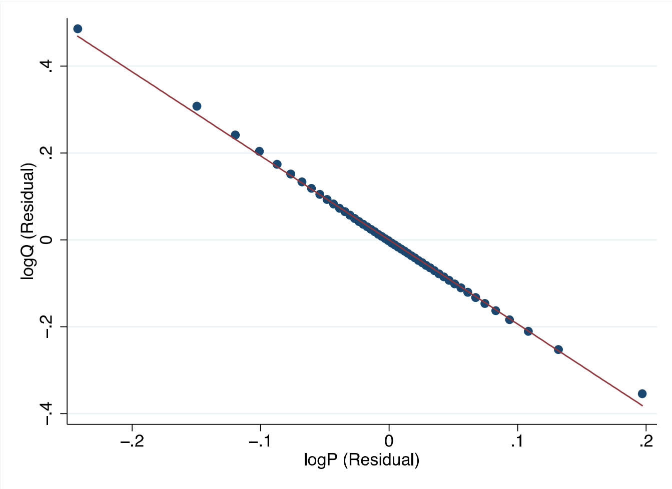

Online Appendix Figure 22d. Distribution of Standard Errors

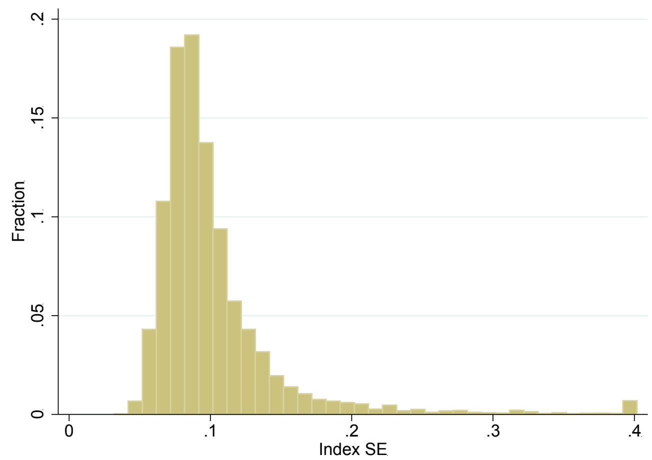

Notes: In Figure 22a., elasticities are Winsorized at 0. In Figure 22b., there are 50 bins representing 9,415 stores. Standard errors are clustered by parent_code. In Figure 22c., there are 50 bins representing 39,778,208 store-week-modules. Residuals are after removing module*quarter-of-year and module*year fixed effects. In Figure 22d., standard errors are clustered bimonthly and are Winsorized at .4.

Online Appendix Figure 22e. Index Elasticity

Elasticity vs. Income

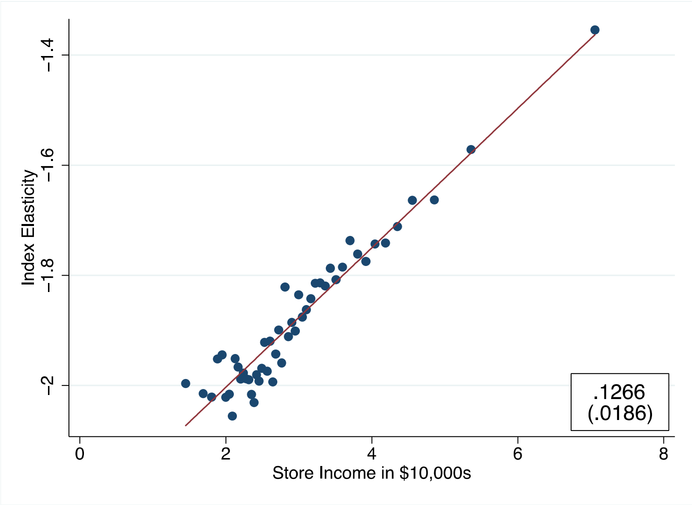

Notes: standard errors are clustered by parent_code.

Online Appendix Figure 23. Implications of Price Rigidity for Grocery Prices in Areas with Different Characteristics, Robustness

Online Appendix Figure 23a. Force Marginal Cost to be the same across all stores, benchmark elasticity

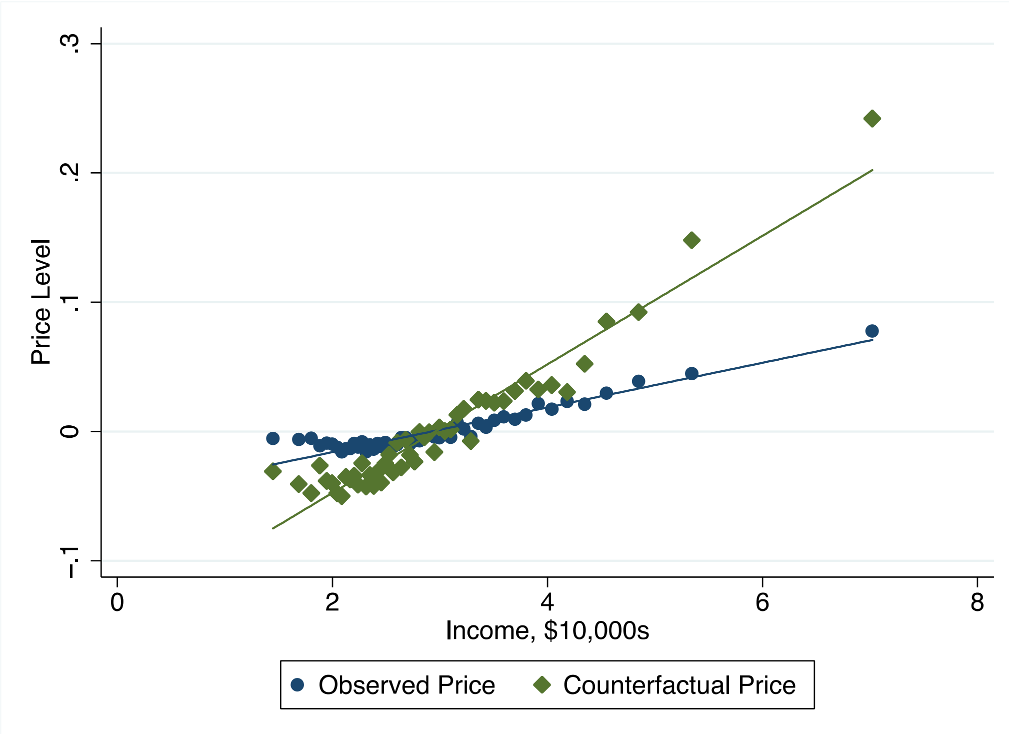

Online Appendix Figure 23b. Allow Marginal Cost to Vary by Chain, predicted elasticity given income

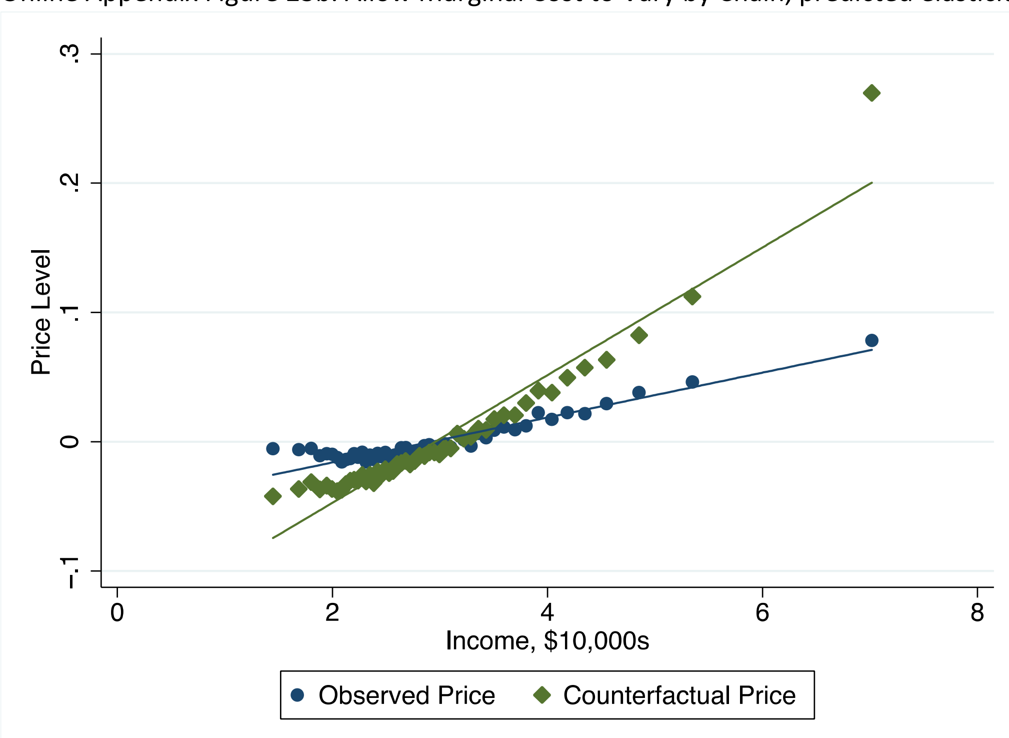

Notes: These are variations of Figure. Figure 23a. forces the marginal cost to be the same across all stores regardless of chain, while Figure 23b. allows marginal cost to vary but uses elasticities predicted with income in a manner identical to the first stage of our IV estimation.

Online Appendix Figure 24. Evidence on Tacit Collusion: Within-Chain-State Price vs. Elasticity by Number of Stores within 10 km

Online Appendix Figure 24a. Food Stores

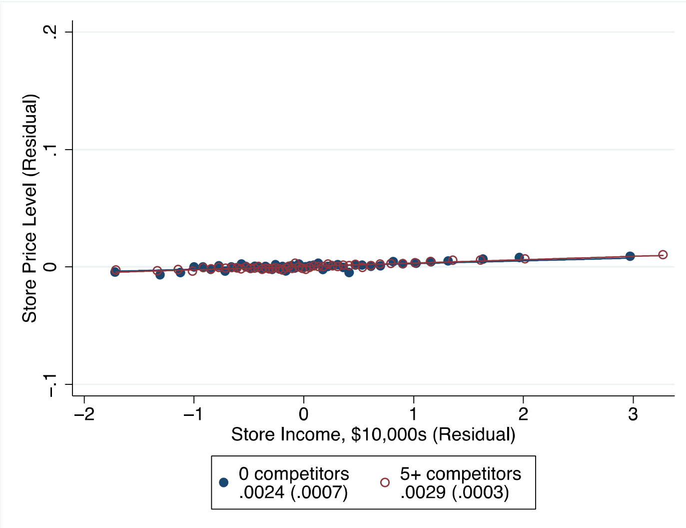

Online Appendix Figure 24b. Drugstores

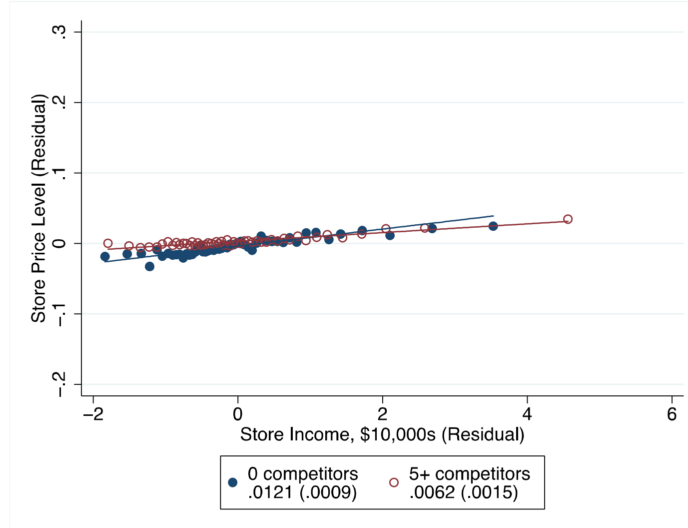

Online Appendix Figure 24c. Mass-Merchandise Stores

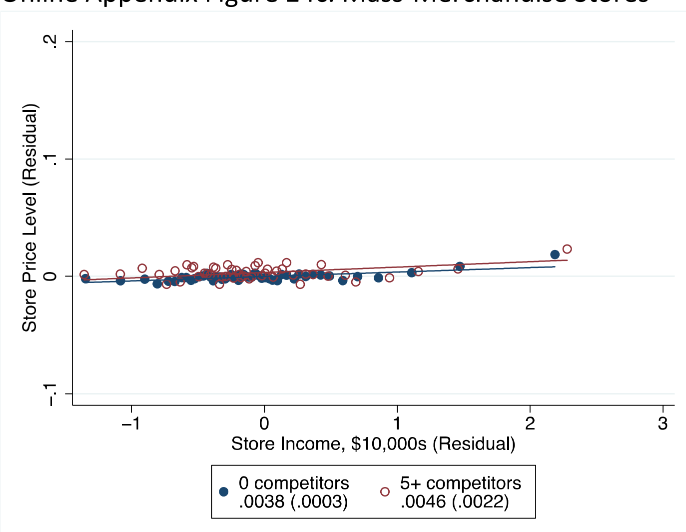

Notes: In Figure 24a-c, residuals are after removing Chain-State FE. Number of competitors is number of other stores of the same type within 10 km, including stores in the same chain.

Onl. Appendix Figure 24d. Between-Chain Average Quarterly Absolute Log Price Difference vs. Fraction of Stores with zero competitors within 10 km

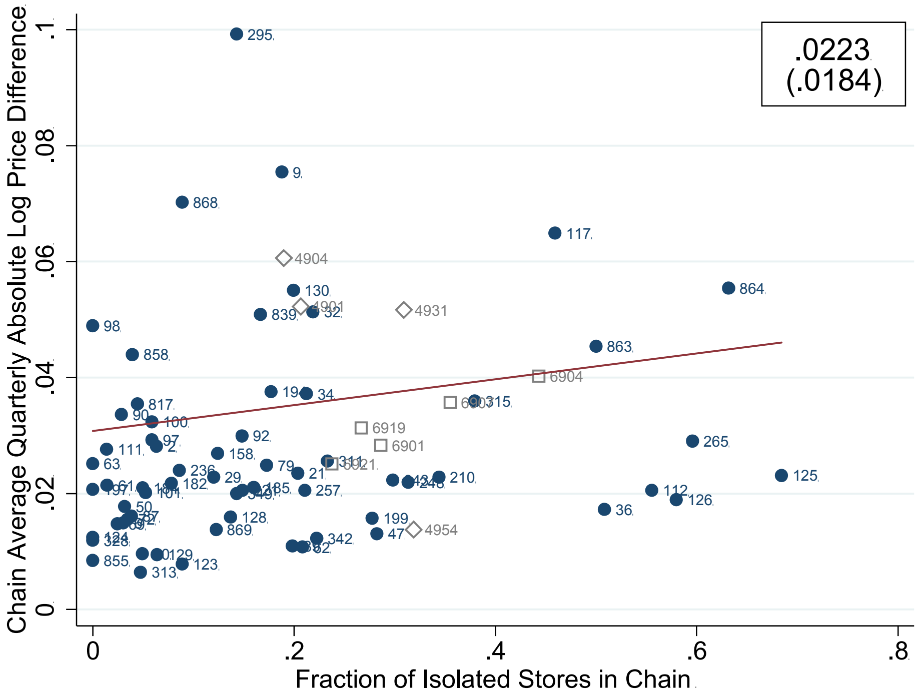

Onl. Appendix Figure 24e. Chain-Level IV Coefficient (Price on Elasticity, instrumented with Income) vs. Fraction of Stores with zero competitors within 10 km

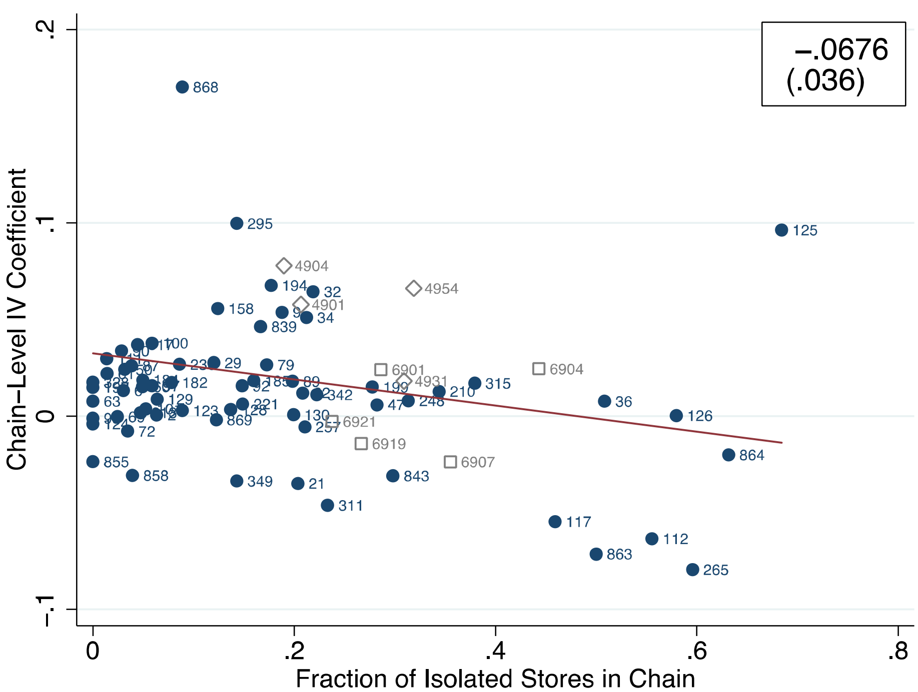

Notes: In Figures 24d. and 24e., the fraction of isolated stores in chain is the fraction of stores within each chain that have zero other stores of the same type within 10 km. The regression line fits only the food stores (solid circles). In Figure 24e., the same first stage using all stores of the type is used for all stores of each type.

Online Appendix Figure 25. Test for Advertising Constraints (Food Stores)

Online Appendix Figure 25a. Evidence of Zone Pricing at the State Level

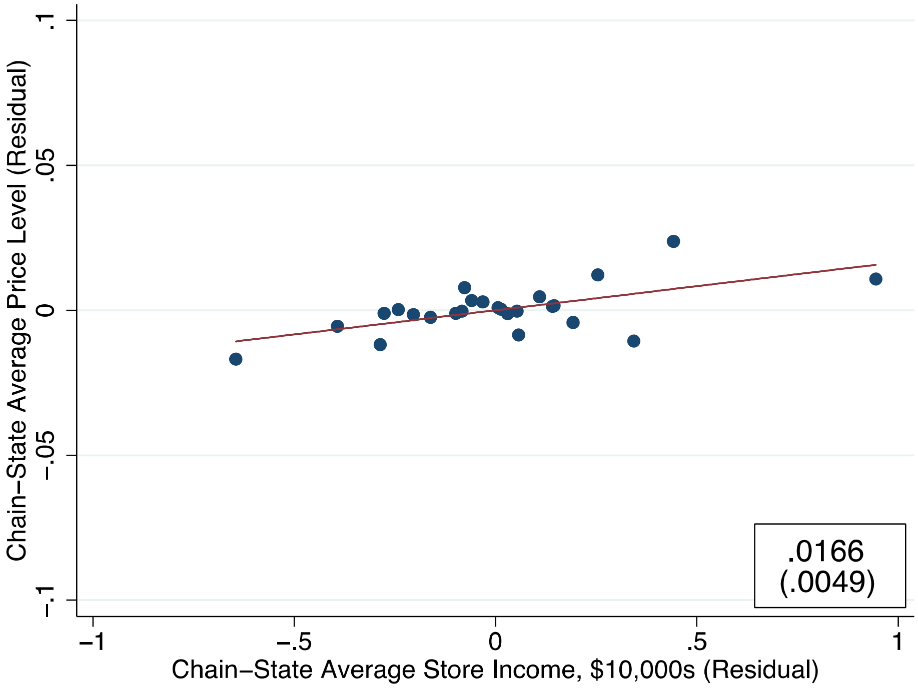

Online Appendix Figure 25b. Evidence of Zone Pricing at the DMA level

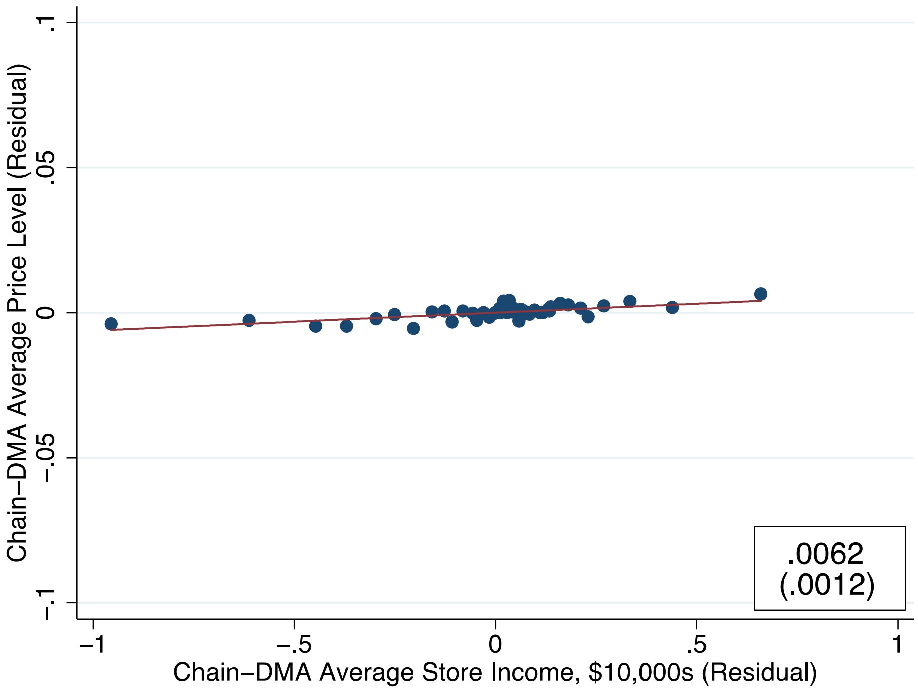

Notes: Only Food stores are included. Analytic weights equal to the number of stores in each chain-geography are used. Figure 25a: each of the 50 bins consists of chain-states. Residuals are after removing Chain FE. Figure 25b: each of the 50 bins consists of chain-DMAs. Residuals are after removing Chain-State FE.

Online Appendix Figure 26. Test for Advertising Constraints (Drugstores and Mass-Merchandise Stores)

Online Appendix Figure 26a-b. Evidence of Zone Pricing at the State Level

Drugstores

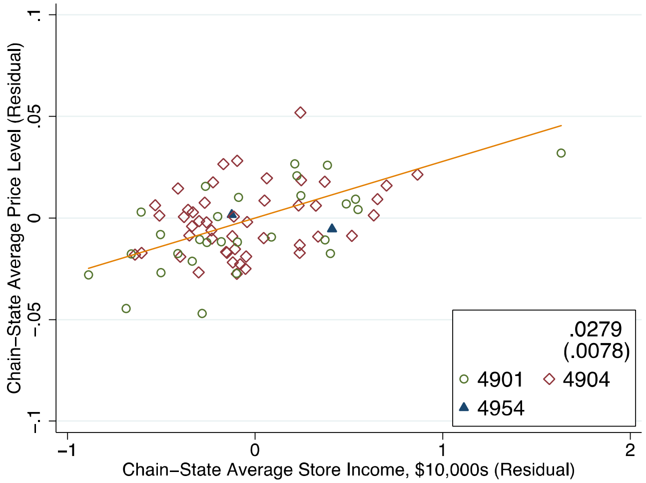

Mass-Merchandise Stores

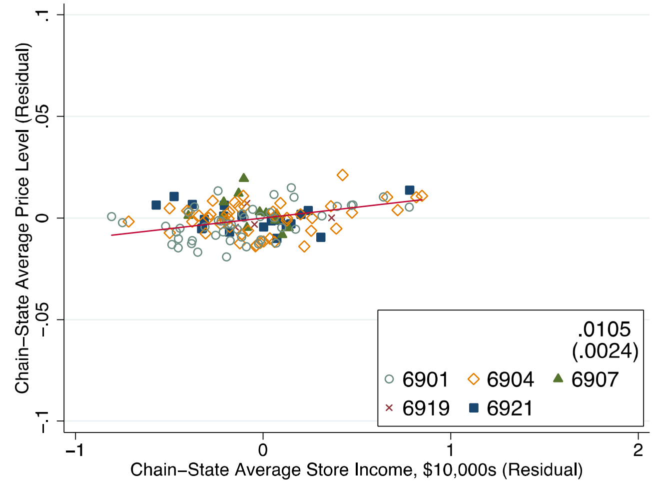

Online Appendix Figure 26c-d. Evidence of Zone Pricing at the DMA level

Drugstores

Mass-Merchandise Stores

Notes: Analytic weights equal to the number of stores in each chain-geography are used. In Panel A, residuals are after removing Chain FE. In Panel B, residuals are after removing Chain-State FE.

Online Appendix Figure 27: Learning Over Time in Food Stores

Onl. App. Fig. 27a-b. Absolute Quarterly Log Price Distance and Weekly Correlation, 2006-08 vs. 2012-14

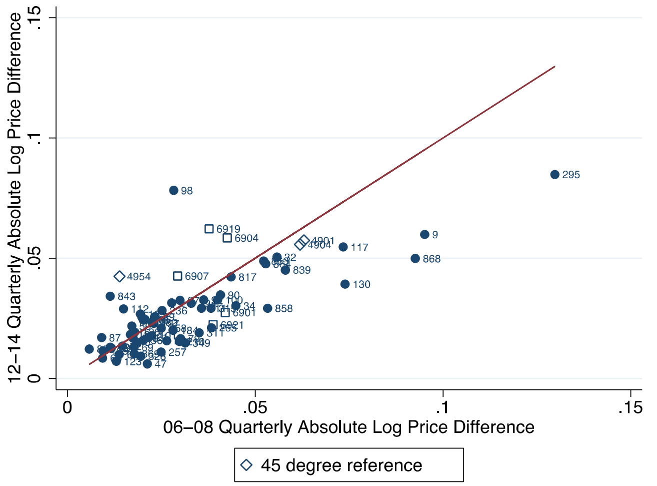

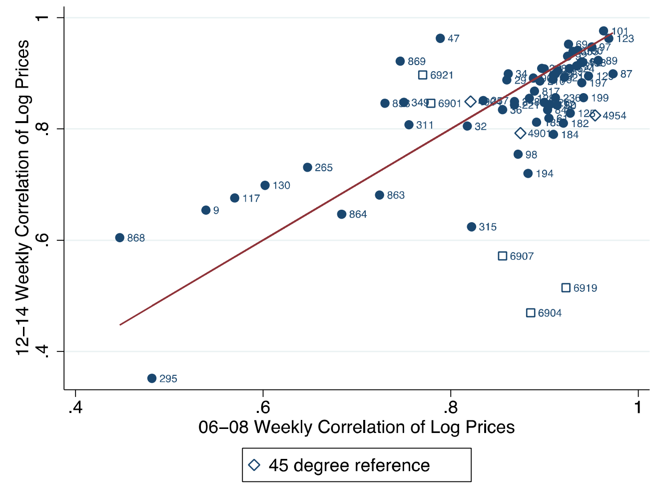

Onl. App. Fig. 27c. Price versus Income Within-Firm Coefficients

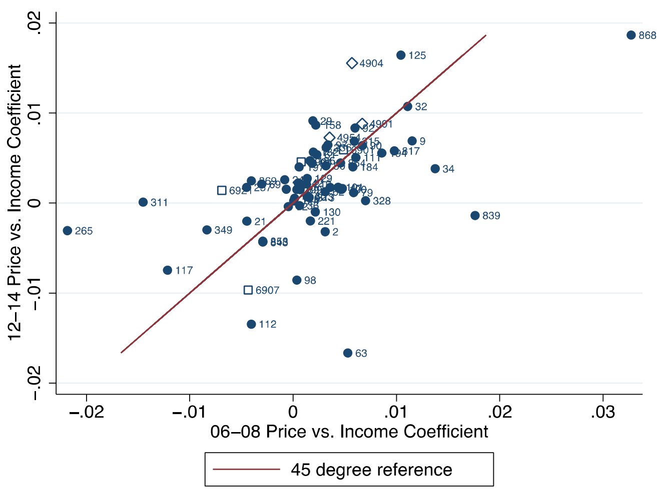

Notes: In Figure 27c., only stores in both periods are included. Chains with robust SE greater than .01 for either period are omitted. The same nine-year average income (2006-2014) is used for all figures so only price levels differ.

Online Appendix Table 1. Additional Summary Statistics

	No. of Stores	No. of Chains	No. of States
	(1)	(2)	(3)
Panel A. Main Sample by Year			
2006	19,252	64	48+DC
2007	20,311	73	48+DC
2008	21,164	73	48+DC
2009	21,564	73	48+DC
2010	21,663	73	48+DC
2011	21,666	73	48+DC
2012	21,669	73	48+DC
2013	21,331	73	48+DC
2014	20,666	70	48+DC

	Yearly Module Revenue by Store (in $)	% of Module Revenue Captured	Average Price
	(1)	(2)	(3)
Panel B. Index Characteristics, Food Stores			
Canned Soup	$55,500	34.0%	$1.15
Cat Food	$11,000	19.8%	$0.49
Chocolate	$41,500	21.3%	$0.87
Coffee	$30,000	15.5%	$5.80
Cookies	$50,500	24.5%	$2.33
Soda	$240,500	53.1%	$1.98
Orange Juice	$77,500	62.2%	$3.09
Yogurt	$80,500	41.6%	$0.82
Bleach	$7,500	41.2%	$1.84
Toilet Paper	$75,500	32.2%	$4.42

Notes: Panel A reports the number of stores and chains in our main sample for each year. In Panel B, we present summary statistics on the price index computed aggregating within each module all products (UPCs) available in at least 95% of week-store observations for that chain (details in the text). We report the average yearly revenue at the store level for all the products included in the index (Column 1), the percent of module revenue covered by the basket (Column 2) and the average price (Column 3).

Online Appendix Table 2. Price versus Income, Within Chain and Between Chain

Dependent Variable:	Price for Store *s*		Average Price for Chain-State	Average Price for Chain c		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Food Stores						
Income Per Capita	0.0044***	0.0029***	0.0166**	0.0448***	0.0505***	0.0395***
(in $10,000)	(0.0013)	(0.0003)	(0.0079)	(0.0101)	(0.0060)	(0.0119)
Fixed Effects	Chain	Chain*State				
Weighted by number of stores			X	X		X
Drop Two Outlier Chains						X
R Squared	0.920	0.958	0.948	0.312	0.542	0.209
Number of Observations	9,415	9,415	171	64	64	62
Panel B: Drug Stores						
Income Per Capita	0.0093***	0.0075***	0.0278***	NA	NA	NA
(in $10,000)	(0.0011)	(0.0008)	(0.0077)			
Fixed Effects	Chain	Chain*State	Chain			
Weighted by number of stores			X			
R Squared	0.443	0.684	0.660			
Number of Observations	9,973	9,973	83			
Panel C: Mass Merchandise						
Income Per Capita	0.0041***	0.0029***	0.0105***	NA	NA	NA
(in $10,000)	(0.0011)	(0.0010)	(0.0024)			
Fixed Effects	Chain	Chain*State	Chain			
Weighted by number of stores			X			
R Squared	0.914	0.945	0.969			
Number of Observations	3,288	3,288	142			

Notes: Standard errors are clustered by parent_code in Panel A and are clustered by parent_code*state in Panels B and C.

Online App. Table 3. Determinants of Price Elasticity, Robustness

Dependent Variable:	Store s Shrunk Estimated Price Elasticity				
	(1)	(2)	(3)	(4)	(5)
Demographic Controls					
Benchmark Income	0.143***			0.106***	0.0636***
(in $10,000)	(0.0087)			(0.0100)	(0.0223)
County Income		0.184***		0.0798***	
(in $10,000)		(0.0172)		(0.0203)	
HMS Range Midpoints			0.0368***	0.00782***	
(in $10,000)			(0.0036)	(0.0014)	
Fraction with College					0.609***
Degree (or higher)					(0.1349)
Median Home Price					0.00457**
(in $100,000)					(0.0021)
Controls for Urban Share					
Controls for Number of Competitors w/in 5km					
1-4 Other Grocery Stores					0.0304
					(0.0268)
5-9 Other Grocery Stores					0.0653*
					(0.0389)
10+ Other Grocery Stores					0.0218
					(0.0531)
Fixed Effect for Chain	X	X	X	X	X
R Squared	0.652	0.629	0.601	0.659	0.662
Number of Observations	22,660	22,660	22,660	22,660	22,660

Notes: Standard errors are clustered by parent_code. All independent variables are our estimate of store-level demographics at the zip-code level based on Nielsen Homescan (HMS) panelists' residences. Data from 2012 ACS 5-year estimates. Number of observations differ due to data availability. Fraction with College Degree (or higher) is the fraction of adults 25 and older with at least a bachelor's degree. Controls for Urban Share are a set of dummy variables for Percent Urban for values in [.8, .9), [.9, .95), [.95, .975), [.975, .99), [.99, .999), and [.999, 1].

Online Appendix Table 4. Log Prices and Store-Level Log Elasticity, OLS

Dependent Variable:	Log Prices in Store *s*		Average Price for Chain-State	Avg. Log Prices for Chain *c*
Specification:	**Within-Chain, OLS**		**Between-Chain-State, OLS**	**Between-Chain, OLS**
	(1)	(2)	(3)	(4)
Panel A. Food Stores				
Log (Elast. / (Elast.+1)) in Store *s*	0.0326***	0.0262***		
	(0.0096)	(0.0059)		
Mean Log (Elast. / (Elast.+1)) in State-Chain Combination			0.0859	
			(0.0516)	
Mean Log (Elast. / (Elast.+1)) in Chain c				0.102*
				(0.0524)
Fixed Effect for Chain	X		X	
Fixed Effect for Chain-State		X		
Number of Observations	9,415	9,415	171	64
Panel B. Drug Stores				
Log (Elast. / (Elast.+1)) in Store *s*	0.158***	0.108***		
	(0.0238)	(0.0121)		
Mean Log (Elast. / (Elast.+1)) in State-Chain Combination			0.324***	
			(0.0777)	
Fixed Effect for Chain	X		X	
Fixed Effect for Chain-State		X		
Number of Observations	9,975	9,975	83	
Panel C. Mass Marchandise Stores				
Log (Elast. / (Elast.+1)) in Store *s*	0.0563***	0.0252		
	(0.0191)	(0.0184)		
Mean Log (Elast. / (Elast.+1)) in State-Chain Combination			0.138***	
			(0.0462)	
Fixed Effect for Chain	X		X	
Fixed Effect for Chain-State		X		
Number of Observations	3,288	3,288	142	

Notes: For Panel A, standard errors are clustered by parent_code. For Panels B and C, they are clustered by parent_code*state. Log Model is log(elasticity/(elasticity+1)). Elasticities above -1.2 are Winsorized. Retailer means for the Between-Chain specification are average Log Model (not Log Model of average elasticity). Analytic weights equal to the number of stores in each group are used in columns (3) and (4).

Online App. Table 5. Log Price and Log Elasticity, Robustness II (Food Stores)

Dependent Variable:	Log Prices in Store *s*			Average Price for Chain-State	Average Log Prices for Chain *c*	
Specification:	Within-Chain, IV w/ income	Within-Chain, IV, All Variables		Between-Chain-State, IV w/ Income	Between-Chain, IV w/ income	Between-Chain, IV w/ All Vars.
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Benchmark Product						
Log (Elast. / (Elast.+1)) in Store *s*	0.0919***	0.101***	0.0643***			
	(0.0339)	(0.0349)	(0.0105)			
Mean Log (Elast. / (Elast.+1))				0.351*	0.936***	0.915***
in Chain c				(0.202)	(0.196)	(0.187)
Fixed Effect for Chain	X	X		X		
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171	64	64
Panel B. Price Index as Dependent Variables						
Log (Elast. / (Elast.+1)) in Store *s*	0.102***	0.112***	0.0749***			
	(0.0301)	(0.0311)	(0.0093)			
Mean Log (Elast. / (Elast.+1))				0.356**	NA	NA
in Chain c				(0.181)	NA	NA
Fixed Effect for Chain	X	X				
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171		
Panel C. Generic top-seller within chain						
Log (Elast. / (Elast.+1)) in Store *s*	0.0594***	0.0676***	0.0414***			
	(0.0223)	(0.0239)	(0.0051)			
Mean Log (Elast. / (Elast.+1))				0.225	NA	NA
in Chain c				(0.178)	NA	NA
Fixed Effect for Chain	X	X				
Fixed Effect for Chain*State			X			
Number of Observations	9,415	9,415	9,415	171		
Panel D. High Quality Items						
Log (Elast. / (Elast.+1)) in Store *s*	0.0985***	0.107***	0.0921***			
	(0.0112)	(0.0127)	(0.0096)			
Mean Log (Elast. / (Elast.+1))				0.184**	0.848***	0.850***
in Chain c				(0.0870)	(0.207)	(0.189)
Fixed Effect for Chain	X	X				
Fixed Effect for Chain*MSA			X			
Number of Observations	9,395	9,395	9,395	170	63	63

Notes: The same first stage (benchmark weekly elasticity on income) is used for all panels; only the second stage differs. Standard errors are bootstrapped. Bootstrap samples are clustered at the parent_code level. 100 replications are used. Elasticities are Winsorized at -1.2. Generic Products in Panel C are chosen by chain. For Panel D., twenty stores do not sell sufficient quantities of our high-quality items.

Online Appendix Table 6. Estimated Loss of Profits at Chain Level, Robustness

	Mean	25th	Median	75th	
Panel A. Food Stores					
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Raw Prices)	14.03%	7.75%	10.32%	15.97%	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Within Chain-State Variation)	8.69%	5.28%	7.17%	9.33%	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Chain-State Means)	11.13%	6.43%	8.12%	10.35%	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Yearly Average Prices)	13.24%	7.25%	8.95%	14.96%	
Panel B. Drugstores	**Chain 4901**	**Chain 4904**	**Chain 4931**	**Chain 4954**	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Raw Prices)	20.80%	13.27%	21.91%	10.35%	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Within Chain-State Variation)	9.07%	8.41%	15.86%	9.91%	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Chain-State Means)	12.89%	10.45%	18.92%	11.97%	
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Yearly Average Prices)	37.62%	11.70%	18.67%	9.13%	
Panel C. Mass Merchandise Stores	**Chain 6901**	**Chain 6904**	**Chain 6907**	**Chain 6919**	**Chain 6921**
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Raw Prices)	19.06%	19.62%	8.67%	53.07%	10.97%
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Within Chain-State Variation)	6.02%	5.16%	3.06%	2.84%	3.77%
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Chain-State Means)	6.19%	5.44%	4.62%	36.12%	5.58%
Loss of Profits Comparing Optimal Pricing to Actual Pricing (Yearly Average Prices)	20.49%	23.42%	13.52%	62.52%	14.59%

Notes: Elasticities are Winsorized at -1.2. The first row in each store type uses observed prices. The second row uses predicted prices from a regression of price on chain-state elasticity residuals and chain-state elasticity averages. The third row uses chain-state means in place of observed prices. The fourth row uses observed average yearly price paid. Since the base is different, this percentage is not directly comparable to the other measures but relative comparisons within each panel are still valid.

Online Appendix Table 7. Test of Firm Learning: Comparing 2006-08 to 2012-14

Dependent Variable:	Log Prices in Store *s*		Average Price for Chain-State	Average Log Prices for Chain *c*	Log Prices in Store *s*		Average Price for Chain-State	Average Log Prices for Chain *c*
Specification:	Within-Chain, IV		Chain-State, IV	Between-Chain, IV	Within-Chain, IV		Chain-State, IV	Between-Chain, IV
Years:	2006-2008				2012-2014			
	(1)		(2)	(3)	(4)		(5)	(6)
Panel A. Food Stores								
Log (Elast. / (Elast.+1)) in Store *s*	0.116***	0.0676***			0.0912**	0.0624***		
	(0.0349)	(0.0104)			(0.0361)	(0.0110)		
Mean Log (Elast. / (Elast.+1)) in Chain c			0.515***	0.977***			0.337	0.769***
			(0.189)	(0.197)			(0.225)	(0.232)
Fixed Effect for Chain	X		X		X		X	
Fixed Effect for Chain-State		X						
Number of Observations	8,642	8,642	167	64	8,642	8,642	167	64
Panel B. Drug Stores								
Log (Elast. / (Elast.+1)) in Store *s*	0.192***	0.167***			0.416***	0.339***		
	(0.0487)	(0.0318)			(0.0647)	(0.0513)		
Mean Log (Elast. / (Elast.+1)) in Chain c			0.458	NA			1.250***	NA
			(0.3646)				(0.328)	
Fixed Effect for Chain	X				X			
Fixed Effect for Chain-State		X				X		
Number of Observations	8,553	8,553	80		8,553	8,553	80	
Panel C. Mass Merchandise Stores								
Log (Elast. / (Elast.+1)) in Store *s*	0.127**	0.0899			0.241***	0.128**		
	(0.0558)	(0.0576)			(0.0762)	(0.0559)		
Mean Log (Elast. / (Elast.+1)) in Chain c			0.315	NA			0.820***	NA
			(0.264)				(0.238)	
Fixed Effect for Chain	X				X			
Fixed Effect for Chain-State		X				X		
Number of Observations	3,012	3,012	139		3,012	3,012	139	

Notes: Standard errors are clustered by parent_code in Panel A and are clustered by parent_code*state in Panels B and C. Elasticities above -1.2 are Winsorized. Retailer means for the Between-Chain specification are average Log Model (not Log Model of average elasticity). The same first stage with nine-year elasticities and incomes are used within each panel. Stores must be present in both the early and late periods in order to be included.

Online Appendix Table 8. Macro Shocks Under Uniform Pricing

Outcome:	Pass-Through of an Income Shock on Food Prices, For Shocks at Different Geographic Level, Assuming Impact as in Table 4, Col. 2	
Price Measure:	**Average Weekly Price**	**Average Yearly Price**
Locality of the Shock	(1)	(2)
National Shock	100%	100%
State-Level Shock	50%	57%
DMA-Level Shock	37%	46%
County-Level Shock	18%	30%

Notes: Displayed are the measured shocks given a response to a permanent 1% shock in income in each locality as a percent of the response to a nationwide shock. Since the base (price response for a 1% national shock) is different, the percentages are not comparable across columns.